



Calix E-Series (E7 OS R2.6) Engineering and Planning Guide

September 2017

#220-01049, Rev. 10



Contents

About This Guide.....	11
Chapter 1: System Overview	15
About the E7 OS Systems.....	16
Supported Technologies and Services.....	20
Supported Customer Premises Equipment (CPE)	23
xDSL Modems	23
ONTs	24
Residential Gateways (RGs).....	25
Pluggable Optics Modules.....	26
Port Summaries	28
E7-2 Platform Description.....	31
Operation Domains Within the E7-2 Modular Chassis	37
Modular Chassis Configuration	40
E7-2 Inter-Card Communication and Backplane Link	45
Mixed Copper and Optical Card Deployment.....	50
E7-2 Chassis	51
E7 LED Behavior	53
E7 Power Subsystem.....	55
E7-20 Platform Description.....	57
Operation Domains Within the E7-20.....	64
Service Intelligence Edge	70
E7-20 Chassis	72
E7-20 Fan Tray Assembly (FTA)	74
E7-20 LED Behavior	76
E3-8G Node-Based OLT Description	78
E3-8G Accessories	81
E3-48C Platform Description	84

E3-48 r2 and E3-48C r2 Platform Description.....	89
E3-48 r2 and E3-48C r2 Chassis Attributes	91
Access Compartment Connections.....	94
Cable Entry Locations.....	96
E5-48 and E5-48C Platform Description	98
E5-48C Combo Unit Overview	98
E5-48 Overlay Unit Overview.....	101
E5-48/E5-48C Chassis Attributes	103
E7 Converged Network Solutions	106
E3-8G Network Solutions.....	108
E3-48C, E5-48/C, E3-48/C r2 Network Solutions	110
 Chapter 2: System Management	 113
Client Management Interfaces.....	114
Embedded-Web Interface.....	115
Command Line Interface.....	117
Calix Management System	121
CMS Features.....	122
View and Manage Calix Equipment	123
CMS Interfaces	123
CMS Connections.....	124
CMS Desktop Views	125
CMS Security Features.....	128
Default CMS Server Ports.....	129
Default Outgoing CMS Server Ports	130
System Software Upgrades	131
(E7-20 Only) In-Service SCP-10GE to SCP2-10GE Upgrades.....	134
Compass Application Suite	136
Security Management.....	137
Configuring Secure System Login Access	137
Configuring Secure Network Access.....	138
New Key and Self-Signed Certificate	140
Fault Management and Event Monitoring.....	141
Link Layer Discovery Protocol (LLDP)	144

ERPS Topology View	146
SNMP	147
SNMPv3 User Accounts	148
Supported E-Series MIBS	149
Loading the MIBs	152
 Chapter 3: Network Topologies	 153
E7 Configurations and Deployments	154
E3-8G Configurations and Deployments	157
E3-48C, E3-48/C r2 Configurations and Deployments	160
Ethernet Topology Protocols	163
Rapid Spanning Tree Protocol (RSTP)	163
Ethernet Ring Protection Switching (ERPS)	167
G.8032v2 Ethernet Ring Protection Switching	174
Link Aggregation Group (LAG)	182
Ethernet Transport and Aggregation	188
Ethernet Ring Topologies	188
Diverse Fiber Route and RSTP Node Protection	191
Ethernet Point-to-Point Aggregation	195
Ethernet Aggregation Topology	197
Protected Uplink with Virtual Router Redundancy (VRRP)	198
Transparent LAN (TLAN)	199
Operation Domains	202
 Chapter 4: Quality of Service	 205
Quality of Service Overview	206
Traffic Classification and Queuing	212
Queuing Overview	212
Layer 3 Priority Mapping	213
VLAN Tag Actions and Traffic Classification	219
Ethernet Traffic Management	221
Ethernet Traffic-Control Objects	221
Ethernet Traffic Data Path	223
Class Map and Class Rules	228
Policy Map and Policies	230

Ethernet Port Class of Service Profile	233
VLAN Tag Actions.....	235
GPON Traffic Management	239
Overview of GPON Upstream Bandwidth Management	239
Overview of the E7 GPON Upstream Traffic Flow	241
Overview of the E7 GPON Downstream Traffic Flow.....	258
Traffic management recommendations.....	262
Applying shaping and policing	263
Traffic shaping on P-Series ONTs.....	264
Queuing Methods on ONT Ethernet Ports	265
Configuration Guidelines and Parameter Descriptions	267
xDSL Traffic Management.....	294
Ethernet Bandwidth Profile for xDSL Services.....	295
DSCP Map.....	299
Match Lists for xDSL Services	300
Service Tag Actions for xDSL Services	303
 Chapter 5: Subscriber Services	 311
Ethernet Elements	312
Ethernet Interfaces.....	312
Ethernet Ports	320
VLAN Support Overview	324
Additional multicast addressing considerations	327
VLAN Memberships for Interfaces	345
VLAN Tag Action to an Ethernet Interface	347
Service-Tag Action to a Provisioned Service	351
IPv4 Security Features	363
DHCP Snoop	363
MAC Forced Forwarding.....	365
IP Source Verify	366
DHCP Relay Option 82 and LDRA	366
(Optional) Configuring DHCPv4 Proxy.....	373
IPv6 Support	383
Lightweight DHCPv6 Relay Agent (LDRA)	383
Neighbor Discovery Protocol (NDP) Flood Control	386
Home Gateway IPv6 Support	386
MEF Compliant Services.....	387

Data Services	390
About Data Services	390
Configuration Process.....	391
PPPoE Subscriber Authentication and Service Activation	392
PPPoE Relay and Intermediate Agent.....	392
Transparent LAN Services Provisioning Model.....	396
Video Service	397
About Video Services.....	397
Configuration Process.....	402
Multicast Profile	403
MVR Profile and MVR VLAN Addresses.....	405
Multicast White List Map and Ranges.....	408
IGMP Snooping Operation.....	410
IGMP Profile	413
VoIP Service	420
About the E-Series VoIP Services Applications	420
Configuration Process.....	422
TDM Gateway Overview	423
SIP Overview	425
H.248 Gateway Overview	437
MGCP Overview	438
IP Hosts for VDSL2 Voice Services	439
IP Hosts for Voice Services on GPON ONT	440
Voice (POTS) Overview.....	444
RF Overlay Video Service	447
RF Video Overlay with RF Return Path	448
Configuration Process.....	449
T1/E1 PWE3 Service	451
E1 Bits Timing Support	454
PWE3 Timing and Clock Recovery Considerations	455
Encapsulation Methods.....	463
T1 Pseudowire Planning Guidelines	466
Configuration Example.....	468
 Chapter 6: GPON Subscriber Access	 473
Calix E7 GPON Solution Overview	474
Supported GPON Equipment.....	476
E7 GPON Configuration Guidelines	482

Home Gateway Overview	484
Home Gateway Concepts	485
Home Gateway IPv6 Support	487
HGW Services	491
Home Gateway Port Assignments	493
RG Management Modes	495
Configuration File Process Flow	499
HGW Management Interfaces	500
ONT Configuration Files for HGW Services	503
Ethernet OAM on GPON 760GX ONTs	506
Ethernet OAM Domains	508
Open Access Networks	510
HPNA Services Deployment	513
Creating GPON ONTs	514
Manually Creating an ONT (Pre-Provisioning)	514
Automatic ONT Creation (RONTA)	518
Automatically Creating an ONT - Smart Activate	520
Resetting an ONT to the Factory Default	526
Upgrading ONTs	528
ONT Ethernet Port	530
Configuring Type B PON Protection	534
ONT Power Shedding	538
Fiber Cabling	539
PON Protection	540
 Chapter 7: xDSL Subscriber Access	 543
Physical Access to Copper Loops	544
Multiple Modes of DSL Deployment and Multi-Pair Bonding	545
xDSL Bonding Groups	546
Copper Bonded Uplinks	552
Remote Powered Bonded Line Pairs	557

Multiple VC Support for ADSL2+ Fallback	560
Integrated Voice Services on the VDSL2 Units	561
G.INP Support	562
High-Speed Internet Access	563
DSL Port Noise Protection	564
VDSL2 Vectoring	569
System Level Vectoring	579
Vectoring Fallback for an SLV Target	582
Vectoring Performance	584
Third-Party Modem Interoperability	586
 Chapter 8: Reference Information	 587
System Support Capacities	588
xDSL Support Capacities	593
VDSL2 Unit Support Capacities	597
Power Consumption and Heat Dissipation	600
Protocols and Standards Supported	602
ODC Cabinet Solutions	604
ODC-100 Cabinet Solutions	607
Calix Outdoor Cabinets	610
ODC Cabinet Powering Options	611
Cabinet Mounting Options	612
Cabinet Retrofit Kits	613



About This Guide

This document provides engineering and planning information for the Calix E-Series portfolio, specifically the E7-2, E7-20, E3-48C, E5-48C, E3-8G, as well as the newly introduced E3-48/C r2. These units are all powered by Calix Ethernet eXtensible Access operating system (E7 OS). This document describes the features and capabilities of each system, and provides engineering guidelines to assist engineers and network planners effectively deploy the systems running E7 OS R2.6.

In this document, the following references are made:

- **"E7"** indicates:
 - E7-2 single-shelf
 - E7-2 modular chassis
 - E7-20
- **"E-Series"** indicates:
 - All E7 products listed above, *and*
 - E3-48C
 - E3-48 r2, E3-48C r2
 - E5-48C
 - E5-48
 - E3-8G

Intended audience

The primary audience for this guide includes network planners, engineers, and other personnel responsible for planning and engineering carrier networks. It assumes that you have an understanding of telephony, Ethernet technology including switching and routing, the Internet, video services networking principles, networking configuration, and experience in access and inter-network administration.

Related documentation

You can access Calix product documentation by logging into My Calix (www.calix.com/my-calix) and browsing the My Calix Documentation Library.

The Calix E-Series documentation set includes:

Engineering and Planning
• <i>Calix Ethernet Access Networks Engineering & Planning Guide</i> This document provides high-level engineering and planning information for building secure, reliable, resilient, and scalable Layer 2 switched Ethernet access networks using Calix products to deliver data, voice, and video services. It describes all aspects of Ethernet access network design—from physical topologies to network bandwidth requirements in the context of Calix products, describing an access network where the traffic terminates at the edge router. • <i>Calix E-Series (E7 OS) Engineering and Planning Guide</i> This document provides engineering and planning information for the Calix E7 Ethernet Service Access Platform (ESAP) and Calix E3-48C/E5-48/E5-48C/E3-8G Ethernet Service Access Nodes (ESAN). It describes the features and capabilities of each system, and provides engineering guidelines to assist engineers and network planners effectively deploy the E-Series.
Installation
• <i>Calix E7-2 Installation Guide</i> • <i>Calix E7-20 Installation Guide</i> • <i>Calix E3-48C Installation Guide</i> • <i>Calix E3-48 r2/E3-48C r2 Installation Guide</i> • <i>Calix E5-48/E5-48C Installation Guide</i> • <i>Calix E3-8G Installation Guide</i> These documents provide a general installation practice for the Calix Ethernet service access platform, including guidance for planning, power installation, cabling, and maintenance.
Initial Turn-up and Transport Configuration
• <i>Calix E-Series (E7 OS) Turn-up and Transport Guide</i> These guides are intended for initial turnup and also show you how to set up management access, system attributes, administrative tasks, and configure transport and aggregation applications.

Software Upgrade
• <i>Calix E-Series (E7 OS) System Upgrade Guide</i> This document describes how to perform software/firmware upgrades as well as database backup and restore operations. • <i>Calix E-Series (E7 OS) GPON ONT Upgrade Guide</i> For E-Series GPON systems, this document describes how to perform upgrades for compatible P-Series, GigaCenter, and T-Series GPON ONTs.
Application Provisioning
• <i>Calix E-Series (E7 OS) GPON Applications Guide</i> • <i>Calix E-Series (E7 OS) xDSL Applications Guide</i> • <i>Calix E-Series (E7 OS) Active Ethernet Applications Guide</i> These application guides show you how to provision subscriber services using specific technologies, assuming that the system is already installed and turned up.
Maintenance and Troubleshooting
• <i>Calix E-Series (E7 OS) Maintenance and Troubleshooting Guide</i> This guide includes procedures for monitoring E-Series network operation, general troubleshooting, and replacing or installing equipment.
Command Line Interface
• <i>Calix E-Series (E7 OS) CLI Reference</i> This document provides a comprehensive command reference for the E-series Command-Line Interface (CLI) and describes how to perform key system management and operational functions from the CLI. The embedded command-line interface (CLI) for system management access can be used over local or remote TCP/IP connections and local console connections.

Related Documentation

- *P-Series/T-Series ONT Software Matrix for E-Series (E7 OS) GPON*
- *Calix P-Series ONT Model/Feature Matrix*
- *Calix 800 GigaFamily Service Provider's Guide*
- *Calix 844G/854G GigaCenter User's Guide*
- *Calix T1 Pseudowire Applications Guide*
- *Calix C7 VoIP Services Guide*
- *Calix Application Note: Using the VoIP Configuration File*
- *Calix Application Note: Using the RG Configuration File*
- *Completing Residential Gateway and SIP Configuration File Intake Forms*
- *Calix P-Series VoIP Configuration File - Template*
- *Calix Residential Gateway Wi-Fi Best Practices Guide*
- *Calix GPON RF Overlay Deployment Guidelines*
- *Calix Application Note: GPON Interface Adaptor*
- *Calix xDSL Best Practices*
- *Calix Application Note: Designing an E3-48/E3-48C r2 Line-Powered Network.*
- *Calix E7 Pluggable Transceiver Module Support*
- *CAB-12-023 - Pairing Bidirectional SFPs (to Support Single-Fiber Ethernet Links)*



Chapter 1

System Overview

This chapter provides an overview of the Calix E-Series (E7 OS) systems.

Topics Covered

This chapter covers the following topics:

- About the Calix E-Series systems
- A description of the platform supported technologies and services
- Supported customer premises equipment (CPE)
- A description of the features and key components per system

About the E7 OS Systems

The Calix E-Series Multiservice Ethernet products provide cost-effective, flexible service delivery of advanced IP services across the entire access network and are designed to work in conjunction with existing Calix equipment.

- The Ethernet transport and aggregation features of the E7-20 are common to the E7-2 products, and the systems can be mixed within Ethernet transport networks.

Note: The term "E7" as used in this document and the larger set of E7 documents, refers to the E7 platform that includes the E7-2, the E7 modular chassis system, and the E7-20 system. When distinctions must be made, the specific configuration will be distinguished.

- The VDSL2 features of the E3-48C, E5-48/C, E3-48/C r2 align with the provisioning, performance, and configuration characteristics of the E7-2 VDSL2 line cards.
- The GPON features of the E3-8G align with the provisioning, performance, and configuration characteristics of the E7-2 GPON-8 line card.

Note: The term "E-Series" as used in this document and the larger set of E7 documents, refers to the platform that includes the E7-20, E7-2, E3-48C, E5-48, E5-48C, E3-48/C r2, and E3-8G products.

E7-20 Access System


The Calix E7-20 Terabit Access System is a revolutionary optical access system designed for high density, high bandwidth fiber deployments. With a 2-Terabit (2 Tbps) backplane providing 100 Gbps to each line card slot, it's ready for the next generation of technologies and the bandwidth needs of tomorrow, while cost effectively meeting service delivery needs of today. The E7-20 delivers gigabit passive optical network (GPON) and point-to-point gigabit Ethernet (Active Ethernet) services with integrated, standards based, 10-gigabit Ethernet (10GE) transport and aggregation within a high availability carrier-class system. The E7-20 allows service providers to deliver triple play and over-the-top residential services and advanced Ethernet business services from a common Unified Access chassis. When coupled with the E7-2 modular chassis ESAP, service providers now have the ability to meet their deployment and broadband capacity needs with the E7 family of access systems.

E7-2 Ethernet Service Access Platform

The Calix E7 Ethernet Service Access Platforms (ESAP) provide cost-effective, flexible service delivery of advanced IP services across the entire access network and are designed to work in conjunction with all Calix products. From 10GE transport and aggregation to VDSL2 and point-to-point Gigabit Ethernet (Active Ethernet) or GPON, Calix E-Series products deliver high capacity broadband from a large chassis (E7-20) or small form factor footprint (E7-2). The modularity and flexibility of Calix E-Series platforms result in superior operational efficiency without sacrificing deployment flexibility or service functionality.

- E7-2 1RU design is small yet can expand with an extra slot for very low first install cost
- Additional chassis are added with subscriber growth yielding a near linear cost curve
- Line cards are managed as a single chassis for operational efficiency
- Mix and match line cards can be installed in a common chassis
- Line cards can be added or removed without affecting other units or cards
- Each chassis unit can be physically added or deleted without affecting services on the other units
- Subscribers are aggregated and network resources efficiently shared across protected trunk facilities
- Resilient, hot-swappable fan tray requires less power

With the Calix E7, service providers no longer need to decide between a single service product and a high-growth chassis solution when considering a new service market. The E7 provides low cost deployments, operational efficiencies, and near linear incremental costs per subscriber. The system grows in step as the subscriber base grows, enabling Calix customers to maximize the return on their investment.



E7-2 Modular Access System

The Calix E7-2 Modular Chassis System is an evolutionary copper and optical access system designed for high density, high bandwidth copper and fiber deployments. With a backplane providing 100Gbps to each line card slot in two slot increments, it's ready for the next generation of technologies and the bandwidth needs of tomorrow, while cost effectively meeting service delivery needs of today. The E7-2 delivers Gigabit Passive Optical Network (GPON), point-to-point gigabit Ethernet (Active Ethernet), and copper pair based DSL services with integrated, standards based, 10-gigabit Ethernet (10GE) transport and aggregation within a high availability carrier-class system.

E3-48C, E3-48 r2, and E3-48C r2 Sealed Access Nodes

The Calix E3-48C, E3-48 r2, and E3-48C r2 Sealed Access Nodes (SAN) are compact, hardened, high capacity sealed systems designed for outside plant applications. The E3-48 products are designed for rapid deployment of advanced copper-based broadband services and optimized to meet the needs of higher subscriber density. The E3-48 products may be AC, DC, or Line powered and support 48 VDSL2 / ADSL2+ ports.

The E3-48 products support advanced VDSL2 as well as ADSL2+ fallback technologies with a non-blocking architecture via 3 x 1GE / 2.6GE (SFP) and 2 x 1GE / 10GE (SFP+) ports as well as supporting unit level vectoring. The E3-48 r2 products introduce new copper bonded uplink and VDSL2 system level vectoring enhancements. The E3-48C aligns with the provisioning, performance, and configuration characteristics of the E7-2 VDSL2 r1 line card while the E3-48r2 products align with the performance, and configuration characteristics of the E7-2 VDSL2 r2 line cards. The E3-48 products support 2-port ATM and PTM bonding, as well as the full suite of E7 DSL capabilities, the units may be integrated smoothly into network deployments.

The E3-48 products can be used to deliver high bandwidth residential and business services from a Remote Terminal (RT), fiber-to-the-node (FTTN), fiber-to-the curb (FTTC) or MDU location. They support environmental flexibility with a sealed enclosure, which has been specially designed for deployment in extreme climate conditions. The E3-48 products have been designed to easily fit within existing rights-of-way (ROW) using AC, DC, or Remote power.

Below is an overview of the key features and benefits of the E3-48 r2 and E3-48C r2 products:

- 48-port VDSL2/ADSL2+ ports with up to 8-port bonding for high bandwidth needs
- Unit Level and System Level Vectoring over VDSL2 to maximize the copper bandwidth
- High-capacity vector control interface port for sealed connection to a sealed E3-VCP
- High-capacity network interfaces of 3 x GE/2.5GE and 2 x GE/10GE for flexible uplink and downlink configurations
- 8-port bonded, vectored VDSL2 for copper trunk uplinks (available on copper uplink enabled models)
- Universal powering for AC, -48vDC, and Line +/- 190vDC powering
- 8-port bonded copper trunk uplinks may be used with line powered pairs
- Interoperability with E-Series and C-Series using ERPS Ethernet protection for main ring or subtended ring topologies
- Daisy chain or RSTP ring (max of 8 nodes) configurations to meet various network topology needs
- Sealed design that meets the most stringent industrial standard on outdoor deployment environment, IP-67/GR-950.

E5-48 and E5-48C

The Calix E5-48 and E5-48C are the latest additions to the E5 fixed form factor DSL products. These products are designed to extend and/or replace existing E5-110 and E5-111 for ADSL2+ deployments, as well as the E5-120 and E5-121 for VDSL2 deployments. The E5-48 provides higher subscriber density per RU and more connectivity topology options than the earlier E5-100 products. Similar to the E7-2 chassis, the E5-48 and E5-48C have physical outside plant connectors that are routed out of the rear of the E5-48 chassis. The E5-48 User Interface and configuration parameters aligns with the provisioning, performance, and configuration characteristics of the E7-2 VDSL2 line card.

E3-8G Node-Based Optical Line Terminal

The E3-8G Node-Based Optical Line Terminal (OLT) is optimized to offer cable operators a solution and form factor that eases the transition from an all HFC network to PON. The E3-8G can be used to support a variety of deployment scenarios ranging from:

- Full HFC to FTTH conversion
- HFC DOCSIS node offload of high bandwidth, data hungry customers
- Business services overlay
- MDU focused deployments
- Rural and remote line extensions

Supported Technologies and Services

The Calix E-Series systems support the following technologies and services:

Fiber subscriber access (E7-2, E7-20, E3-8G)

- Fiber-to-the premises (FTTP) subscriber access over Gigabit Passive Optical Networks (GPONs)
 - Compliant with the ITU G.984 family of standards
 - GPON OLT-to-ONT reach up to 20 km, or up to 40 km (with extended-reach optical interface modules)
 - Interfaces available on GPON ONTs: DS0 voice, T1/E1 (with PWE3 transport), Ethernet, HPNA, CATV RF
- Point-to-point GE links (Active Ethernet) over fiber to subscribers on AE ONTs
 - Compliant with the IEEE 802.3ah standard for AE technology
 - Interfaces available on AE ONTs: DS0 voice, T1/E1 (with PWE3 transport), Ethernet, HPNA

Copper subscriber access (E7-2, E5-48C, E5-48, E3-48C, E3-48 r2, E3-48C r2)

- Fiber-to-the node (FTTN), with support for the following services over copper pairs:
 - DS0 (POTS) voice-only
 - xDSL-only (bonding-capable)
 - DS0 voice and xDSL combined (bonding-capable)
- Support for the following xDSL service types: VDSL2 (G.993.2), G.DMT (G.992.1), T1.413, G.Lite (G.992.2), ADSL2 (G.992.3), ADSL2+ (G.992.5), Annex M (G.992.3 & G.992.5 Annex M), and Reach Extended ADSL2 (G.992.3 Annex L)

Network transport and aggregation

- Ethernet transport—10GE ring (ERPS or RSTP)
- Point-to-point links (RSTP or LAG)
- ITU G.8032 Ethernet Ring Protection Switching (ERPS)
- IEEE 802.1w Rapid Spanning Tree Protocol (RSTP)
- IEEE 802.3ad/802.1AX Link Aggregation
- E7-2/E7-20 systems provide aggregation and transport of Ethernet traffic from multiple Calix systems (E-Series, C-Series, B-Series, etc.) in the access network

Supported services

E-Series systems support the delivery of IP services over Ethernet over fiber, with fiber or copper access (FTTP with fiber access or FTTC with copper access). E-Series network VoIP is compatible with SIP, H.248, and MGCP soft switches, as well as traditional TDM voice switches (via C7 TDM voice gateways). In addition, the following services are also supported via E7 GPON/AE ONTs: T1/E1 services via PWE3, HPNA (712/722GE), and CATV RF overlay (GPON only).

- Supported Layer-2 Ethernet services per IEEE and DSL Forum TR-101 specifications:
 - Ethernet 802.1d/802.1ad bridging, VLAN translations, and VLAN stacking (Q-in-Q)
 - Hybrid VLAN service models (untagged and tagged/double-tagged traffic) on the same Ethernet uplink.
 - Metro Ethernet Forum (MEF) compliant services (MEF 9 and 14) over GPON and point-to-point Ethernet
 - Ethernet Virtual Connections (EVCs) over E-LINE (point-to-point), and ELAN (multipoint-to-multipoint) services over both GPON and GE networks
 - Ethernet Private Line (EPL) and Transport LAN (TLAN) business service models. (Untagged, single-tagged, or double-tagged traffic from the business.)
- Supported data service models:
 - VLAN per Service (N:1)
 - VLAN per Port (1:1)
 - Transparent LAN (TLAN) Service
 - IPv6 support for data services
 - Delivery of concurrent IPv4 and IPv6 data services
 - Lightweight DHCPv6 Relay Agent (LDRA)
 - Neighbor Discovery Protocol (NDP) Flood Control
 - DHCPv6 Snooping
- Supported voice services:
 - E7 with GPON/AE ONTs: SIP (via E7 provisioning), SIP (via TR-069 provisioning), H.248, MGCP, TDM (via C7 TDM voice gateway)
 - E7-2 VDSL2-48C, E5-48C, E5-48, E3-48C, E3-48/C r2: SIP (via E7 provisioning), H.248, TDM (via C7 TDM voice gateway)
 - E7-2 VDSL2-48: TDM (via pass-through of existing POTS lines)
- Supported video services:
 - Internet Protocol Television (IPTV), broadcast and video on demand (VOD)
 - Tested compatibility with leading IPTV middleware ecosystems (Myrio, Minerva, Microsoft Mediaroom, etc.)

- Multiple IPTV video + data service models supported, including:
 - Video + data without an RG, edge QoS control
 - Video + data with RG, edge or core QoS control
 - Service level control per subscriber (number of streams and multicast VLAN registration)
- Over-the-top (OTT) video as a data service

Note: For supported xDSL and GPON ONT troubleshooting features, including supported statistics and PM data, see the *Calix E-Series (E7 OS) Maintenance and Troubleshooting Guide*.

Supported Customer Premises Equipment (CPE)

xDSL Modems

Required for copper access lines, and commonly known as the ADSL Termination Unit-Remote (ATU-R), the xDSL modem provides the DSL termination on WAN side and an Ethernet LAN interface on the customer side (for PCs and STBs).

Note: CMS supports the ability to manage CPE devices using a Calix-supported TR-069 service provider, such as ClearAccess. The TR-069 Customer Premises Equipment (CPE) WAN Management Protocol employs a common method for CPE devices to communicate with an Auto Configuration Server (ACS). See "Managing CPE Devices" in the *Calix Management System (CMS) Guide*.

The E-Series supports interoperability with standards-based VDSL2 and ADSL2+/ADSL modems.

Third Party Modem Interoperability

Calix continues to test VDSL2 and ADSL2+ modems against the different DSL Service Platforms supported. These modems must be commercially available and candidates for service deployment by Calix customers. Extensive rate / reach and data throughput performance testing is performed against all of modems. Qualification testing is conducted for physical layer, HIS / data as well as IPTV / Video service performance. Customer Advisory Bulletins (CAB) documents are available from Calix summarizing completed testing; new versions of the document will provide updates to tests that are in progress. Please refer to up to date CABs from Calix for the latest list of qualified and supported modems, which can be located on the Calix documentation portal.

For the initial release of E7 software R2.1 Calix has fully tested interoperability with the following modems:

- Comtrend NL-3101 – This modem covers:
 - Single-port ADSL2+ ATM Fallback
 - Single-port ADSL2+ PTM Fallback
 - ADSL2+ ATM bonding
 - ADSL2+ PTM bonding
 - VDSL2 bonding,
 - Single-port VDSL2

- Comtrend CT-5374 – This modem covers:
 - Single-port ADSL2+ ATM Fallback
 - Single-port ADSL2+ PTM Fallback
 - Single-port VDSL2
- Actiontec Q1000 – This modem covers:
 - Single-port ADSL2+ ATM Fallback
 - Single-port ADSL2+ PTM Fallback
 - Single-port VDSL2

ONTs

Required for fiber access lines, the ONT terminates the Passive Optical Network (PON) delivered services at the customer premises. The ONT typically provides the following end-user service terminations:

- Voice: RJ-11 interfaces for POTS
- Data: An RJ-45 10/100BaseT interface for data and IP video
- Analog Video: A 75-ohm coaxial Radio Frequency (RF) interface capable of supporting up to four RF devices
- Home Phoneline Networking Alliance (HPNA) (Calix 712, 722, and 761 ONTs)

Supported ONT Models/Families

For supported ONTs and additional information, see *Supported GPON Equipment* (on page [476](#)).

Independent OLT/ONT upgrades are also supported.

See the following matrices for information on the available models of Calix ONTs:

- *Calix Software Matrix*
- *Calix P-Series ONT Model/Feature Matrix*
- *Calix P-Series ONTs: Minimum Platform Software*
- *GigaFamily Firmware Matrix*

Residential Gateways (RGs)

The E7 supports the delivery of services to a subscriber via a residential gateway (RG), which is defined as follows:

Residential Gateway: A residential gateway (RG) is a type of broadband router that connects devices in the home or business to the Internet or some other WAN, and may be integrated with a modem or ONT. RGs typically provide all of the following features:

- An IP layer demarcation, allowing the LAN ports to operate as a single Ethernet bridge/switch, supporting value-added services such as media sharing only within the home.
- For IPv4, Network Address Translation (NAT), which allows all the devices in the home to receive unique private addresses while sharing a smaller number of public addresses* on the WAN. For IPv6, the ability to advertise one or more subnets* to the home, which are assigned by the service provider.
- Advanced QoS control between the home and the WAN, to prioritize and classify multiple services.
- Advanced firewall capability, providing enhanced security and protection between the home and the WAN.

*** Note:** Usually only one (1) public IP address for IPv4, or one (1) service provider subnet for IPv6.

The E7 supports RGs that are connected to or integrated with modems or ONTs. For GPON and AE subscriber access, the E7 is compatible with the Calix 836GE RSG, 700GE, and GigaCenter ONTs with the Calix Home Gateway (HGW) software enabled. For more information, see *Home Gateway* (on page [484](#)).

Calix ONTs

See the following matrices for information on the available models of Calix ONTs:

- *Calix Software Matrix*
- *Calix P-Series ONT Model/ Feature Matrix*
- *Calix P-Series ONTs: Minimum Platform Software*
- *GigaFamily Firmware Matrix*

Pluggable Optics Modules

Pluggable optics

The E7 platform supports industry standard pluggable modules for all service and network interfaces, including ITU G.984 compliant GPON, SFP GE, XFP 10GE ports, and SFP+ 10GE ports.

Note: The 10GE SFP+ port on the E7-20 SCP card does not support GE SFP modules.

See the following Calix Quick Tip bulletins on the Calix website for information on the support for SFP modules and direct-attach cables:

- *Calix E7 Pluggable Transceiver Module Support*
- *CAB-12-023 - Pairing Bidirectional SFPs (to Support Single-Fiber Ethernet Links)*

Module restrictions

The following module restrictions apply:

- Unless otherwise specified, SFP sockets in all Calix products accept any SFP module.
- GE-12 cards (E7-2) support any Calix-keyed 1GE SFP module in the SFP sockets.
- Copper SFP modules should be purchased from Calix to ensure correct operational behavior.
- The 10GE, 2.5GE, GPON OIM, and CSFP modules must always be purchased from Calix.

Module specifications

Refer to the *Calix Optical Transceiver Modules Datasheet* for a complete list of modules and specifications.

CSFP Option 2	1GE optical dual-port Compact Small Form-factor Pluggable (CSFP) Option 2 modules
SFP	1GE and 2.5GE optical and copper Small Form-factor Pluggable (SFP) modules
SFP+	10GE optical Enhanced Small Form-factor Pluggable (SFP+) modules
Direct Attach	Multi-rate copper Small Form-factor Pluggable (SFP/SFP+) cables
XFP	10GE optical Small Form-factor Pluggable (XFP) modules
GPON OIM	2.5Gbps GPON (Class B+ ODN with minimum 28dB link budget, up to 1:64 splits)
ER-GPON OIM	2.5Gbps Extended Reach GPON (up to 40 km with 1:8 split) SFP GPON/C+ 60Km 1490nm SC/UPC CO G.984.2

Notes:

- 10GE SFP+ pluggable transceivers and Direct Attach cables must be purchased directly from Calix. The use of 10GE SFP+ pluggable transceivers and Direct Attach cables not purchased directly from Calix is not supported and will void all product warranties covering the Calix equipment to which such third-party materials are connected.
- SFP modules may also be used in SFP+ sockets at 1GE rate.
- Copper Direct Attach cables can operate in SFP and SFP+ sockets at 1GE, 2.5GE, and 10GE data rates as supported by the card type.

Pluggable modules' CLEI code

Telcordia CLEI codes have been added to all Calix pluggable modules (SFP, CSFP, SFP+, XFP, and direct attach cables). With Release 2.0 and above, operators can retrieve the CLEI code via the E7 user interface without removing the SFP from the socket and disrupting service.

2.5G Ethernet on SFP Ports

The E7-2 supports 2.5 Gbps Ethernet interfaces on the SFP ports of the 10GE-4, GPON, GE-24, and VDSL2 cards. This software-only feature does not require any new E7 line cards. However, 2.5GE pluggable modules are required to support the 2.5G interfaces. E7-2 GE-12 line cards, E7-20 GE-24x and SCP cards do NOT support 2.5 Gbps Ethernet interfaces.

E7 2.5GE Planning Guidelines

- E7 10GE-4, GE-24, GPON, and VDSL2 cards support 2.5 Gbps Ethernet interfaces on SFP ports.
- 2.5GE interfaces are NOT supported on SFP+ ports.
- E7-2 GE-12 line cards, E7-20 GE-24x and SCP cards do NOT support 2.5Gbps Ethernet interfaces.
- 2.5GE interfaces support equivalent functions and networking protocols as GE and 10GE interfaces.
- 2.5GE modules must be purchased from Calix. 2.5GE modules are new to the market and purchase from Calix ensures correct operation.

Port Summaries

The following tables provide summaries of the features and functions supported on each port type.

1G Fiber SFP in SFP/CSFP Port

Valid speed settings:	auto, 1G, Module Rate
Speed settings triggering alarm:	2.5G (not E5-400), 100M, 10M
Speed settings rejected by system:	10G
Valid duplex settings:	Any
Valid flow control settings:	Any

2.5G Fiber SFP in SFP Port

This is supported on GPON, GE-24 and 10GE-4 cards, only.

Valid speed settings:	auto, 2.5G, Module-Rate
Speed settings triggering alarm:	1G, 100M, 10M
Card types triggering alarm:	not GPON-4, not 10GE-4
Speed settings rejected by system:	10G
Valid duplex settings:	Any
Valid flow control settings:	Any

Copper SFP in SFP/CSFP Port

This applies to GE-24 card odd-numbered ports, only.

Valid speed settings:	auto, 1G, 100M, 10M, Module-Rate
Speed settings triggering alarm:	2.5G
Speed settings rejected by system:	10G
Valid duplex settings:	Any
Valid flow control settings:	Any

Direct Attach Cable in SFP/CSFP Port

Valid speed settings:	auto, 1G, Module-Rate
Speed settings triggering alarm:	100M, 10M
Speed settings rejected by system:	10G
Valid duplex settings:	Any
Valid flow control settings:	Any

SFP+ in SFP/CSFP Port

This is not supported and will trigger an alarm regardless of port settings.

CSFP in SFP Port

This is not supported and will trigger an alarm regardless of port settings.

SFP/CSFP/ Direct Attach Cable in CSFP Port

This applies to GE-24 card, only.

Valid speed settings:	auto, 1G, Module Rate
Speed settings triggering alarm:	100M, 10M
Speed settings rejected by system:	10G, 2.5G
Valid duplex settings:	Any
Valid flow control settings:	Any

1G Fiber SFP in SFP+ Port

Valid speed settings:	auto, 1G, Module-Rate
Speed settings triggering alarm:	10G
Speed settings rejected by system:	2.5G (not E5-400), 100M, 10M
Valid duplex settings:	Any
Valid flow control settings:	Any

2.5G Fiber SFP in SFP+ Port

This is not supported and will trigger an alarm regardless of port settings.

Copper SFP in SFP+ Port

Link will come up if link partner supports 1000BASE-T. E5-400 and Link speeds less than 1G are not supported.

Valid speed settings:	auto, 1G, Module-Rate
Speed settings triggering alarm:	10G
Speed settings rejected by system:	2.5G, 100M, 10M
Valid duplex settings:	Any
Valid flow control settings:	Any

10G SFP+ in SFP+ Port

Valid speed settings:	auto, 10G, Module-Rate
Speed settings triggering alarm:	1G
Speed settings rejected by system:	2.5G (not E5-400), 100M, 10M
Valid duplex settings:	Any
Valid flow control settings:	Any

Direct Attach Cable in SFP+ Port

Valid speed settings:	auto, Module Rate, 10G, 1G
Speed settings triggering alarm:	None
Speed settings rejected by system:	2.5G (not E5-400), 100M, 10M
Valid duplex settings:	Any
Valid flow control settings:	Any

XFP in XFP Port

Valid speed settings:	auto, 10G
Speed settings triggering alarm:	None
Speed settings rejected by system:	2.5G (not E5-400), 1G, 100M, 10M, Module-Rate
Valid duplex settings:	full, auto
Duplex settings rejected by system:	Half
Valid flow control settings:	Any

Provisioning Checks

The Provisioning Manager performs the following checks when an Ethernet port is provisioned:

Provisioning Checks for 10GE-4, GE-24 and GPON Cards

Port	Allowed	Disallowed	Default
SFP	10M, 100M, 1G, 2.5G (note E5-400), Auto, Module-Rate	10G	Auto
SFP+	1G, 10G, Auto, Module-Rate	10M, 100M, 2.5G (not E5-400)	Module Rate
XFP	10G	10M, 100M, 1G, 2.5G (not E5-400), Auto, Module-Rate	10G

Provisioning Checks for All Other E7-2 and E7-20 Cards

Port	Allowed	Disallowed	Default
SFP/CSFP	10M, 100M, 1G, Auto, Module-Rate	10G, 2.5G	Auto
SFP+	1G, 10G, Auto, Module-Rate	10M, 100M, 2.5G	Module Rate
XFP	10G	10M, 100M, 1G, 2.5G, Auto, Module-Rate	10G


E7-2 Platform Description

The Calix E7-2 Ethernet Service Access Platform (ESAP) integrates IP service delivery and Ethernet transport into a compact, high availability (HA) and carrier-class, modular chassis that delivers extensible network solutions for service providers.

The E7-2 ESAP delivers the following applications within a single 1RU, 2-slot chassis:

- Gigabit Passive Optical Network (GPON)
- Point-to-point Gigabit Ethernet (GE) and Active Ethernet (AE) services
- Copper-based DSL services
- Integrated, standards based, 10-Gigabit Ethernet (10GE) transport and aggregation

The E7-2 allows service providers to deliver differentiated triple play services, advanced Metro Ethernet Forum (MEF) compliant business services, and mobile backhaul from a single converged network. When coupled with the E7-20 multi-terabit chassis, service providers now have the ability to meet their deployment and broadband capacity needs with the E7 family of Ethernet Service Access Platforms.



Front View of E7-2 Chassis with Two Line Cards



Rear View with Power, Alarm, Management, Timing and RF-21 Copper Connectors

Proprietary Information: Not for use or disclosure except by written agreement with Calix.

Calix E7-2 features

Calix E7-2 features are provided by a wide array of Ethernet, VDSL2, and GPON line cards.

E7-2 Line cards:

10GE-4 line card

- Four 10-Gigabit Ethernet (10GE) ports
 - Twelve Gigabit Ethernet (GE) ports
- The Calix E7-2 10GE-4 card provides high speed Ethernet transport with integrated aggregation of lower speed Ethernet devices. The E7-2 10GE-4 line card can be installed into one or both of the two universal slots within a Calix E7-2 shelf to create a compact, high availability (HA) Ethernet transport switch for aggregation and delivery of IP services across the access network. The E7-2 10GE-4 supports a full set of Ethernet services and network topology protocols and can be used interchangeably with other E7-2 line cards to create a redundant system configuration.
- Note:** E7-2 supports 2.5 Gbps pluggable module interfaces on the SFP ports of the 10GE-4 cards. The 2.5GE interfaces support equivalent functions and networking protocols as GE and 10GE interfaces.

GPON-4 r2 line card

- Four ITU G.984-compliant Gigabit Passive Optical Network (GPON) interfaces
 - Eight Gigabit Ethernet (GE) interfaces
 - Four ports of integrated 10-Gigabit Ethernet
- The Calix E7-2 GPON-4 r2 line card can be installed into one or both of the two universal slots within a Calix E7-2 shelf. Each PON supports up to 64 Optical Network Terminations (ONTs) for a total of 512 GPON and 16 point-to-point Ethernet subscribers per 1RU E7-2 chassis. The E7-2 GPON-4 r2 card supports a full set of Ethernet services and network topology protocols on the Ethernet ports and can be used interchangeably with other E7-2 line cards to create a redundant system configuration.
- Note:** E7-2 supports 2.5 Gbps pluggable module interfaces on the SFP ports of the GPON-4 r2 cards. The 2.5GE interfaces support equivalent functions and networking protocols as GE and 10GE interfaces.

GPON-8 line card

- Eight ITU G.984-compliant Gigabit Passive Optical Network (GPON) interfaces
 - Four Gigabit Ethernet (1GE/2.5GE) interfaces
 - Two ports of integrated 10-Gigabit Ethernet
- The Calix E7-2 GPON-8 line card can be installed into one or both of the two universal slots within a Calix E7-2 shelf. Each PON supports up to 64 Optical Network Terminations (ONTs) for a total of 512 subscribers per card. The E7-2 GPON-8 card supports a full set of Ethernet services and network topology protocols on the Ethernet ports and can be used interchangeably with other E7-2 line cards to create a redundant system configuration.

GE-24 line card

- Twenty-four (GE) interfaces
 - Two 10-Gigabit Ethernet (10GE) ports
- The Calix E7-2 GE-24 has 12 pluggable optics sockets that support CSFP (Compact SFP) 'dual port,' single fiber Gigabit Ethernet (GE) modules that are mechanically compatible with SFP modules. Each CSFP module provides two fully independent GE interfaces, aggregating up to 48 Point-to-Point / Active Ethernet Optical Network Terminations (ONTs) per E7-2.
- Note:** E7-2 supports 2.5 Gbps pluggable module interfaces on the SFP ports of the GE-24 cards. The 2.5GE interfaces support equivalent functions and networking protocols as GE and 10GE interfaces.

GE-12 line card

- Twelve Gigabit Ethernet (GE) interfaces
 - Two ports of integrated 10-Gigabit Ethernet (10GE)
- The Calix E7-2 GE-12 line card can be installed into one or both of the two universal slots within a Calix E7-2 shelf to support up to 24 point-to-point Ethernet Optical Network Terminations (ONTs) per 1RU E7-2 chassis. The GE-12 card can also be used along with the E7-2 10GE-4 or E7-2 GPON-8 cards to provide additional point-to-point Ethernet service drops in mixed-application configurations.

VDSL2-48 r2 line card

- 48x DSL Overlay ports with corresponding 900-Ohm integrated splitters
- Four 1GE / 2.5GE SFP MSA sockets, and two 1GE / 10GE SFP+ MSA sockets
Integrated splitters provide a means to extract analog voice from the service loop and terminate it on an existing DLC or analog POTS port on a switch line bay. One E7-2 VDSL2-48 line card can be plugged into a Calix E7-2 shelf to create a compact, high density DSL overlay node, with Ethernet aggregation and transport, ideal for copper based delivery of IP services leveraging existing POTS infrastructure, across the access network.

VDSL2-48C r2 line card

- 48x Combo DSL and POTS ports, two 1GE / 2.5GE SFP MSA sockets
- Two 1GE/10GE SFP+ MSA sockets
The E7-2 VDSL2-48C line card can be plugged into one or both of the two universal slots within a Calix E7-2 shelf to create a compact, very high density DSL node of 96 ports in a 1RU high chassis, for DSL and VOIP access.

VDSL2-48D r2 line card

- 48x VDSL2/ADSL2+ Fallback data-only ports
- Two 1GE / 2.5GE SFP MSA sockets
- Two 1GE/10GE SFP+ MSA sockets
The E7-2 VDSL2-48D line card can be plugged into one or both of the two universal slots within a Calix E7-2 shelf to create a compact, very high density DSL node with Ethernet aggregation and transport, for copper based delivery of IP services across the access network.

VCP-192 line card

- Two SFP sockets at 1GE rate (redundant ports, MGMT GE, for out-of-band management only)
- Two 20GBps Vector Control Interface ports supporting two line cards each
The Calix E7 VDSL2-Vector Control Processor (VCP) line card is a dedicated VDSL2 vector control engine that optimizes the VDSL2 signaling for up to four 48-port E7-2 VDSL2 r2 line cards. The VCP-192 is a two-slot line card that operates in an E7-2 standalone shelf and connects to each of the E7-2 VDSL2 r2 Combo, Overlay, or Data Only line cards through a dedicated vector control interface port, providing 192 ports of vectored controlled VDSL2 density. In-band management of the VCP-192 is not supported.

Maximum interface count per 1RU E7 chassis

The E7 delivers a wide array of high-performance applications, including 10GE Ethernet transport, delivery of high density residential triple play services over GPON and point-to-point Ethernet, Metro Ethernet Forum (MEF) compliant business services, mobile backhaul, and protected GE aggregation of other E7 chassis as well as Calix C7 and E5 platforms.

Maximum Interface Count Per 1RU E7 Chassis	
10GE Transport	8 ports 10GE with 24 ports GE aggregation
GE Aggregation	48 ports GE with 8 ports 10GE transport and uplink
GPON Service	16 GPON OLT ports, 32 GE ports for service or aggregation, and 10GE transport and uplink
xDSL Service	96 xDSL ports, 96 voice (POTS), two GE/2.5GE ports for service or aggregation, and two 10GE transport and uplink
Point-to-Point (Active) Ethernet	48 ports GE with up to 4 ports 10GE uplink

Pluggable optics

The E7 platform supports industry standard pluggable modules for all service and network interfaces, including ITU G.984 compliant GPON, SFP GE, XFP 10GE ports, and SFP+ 10GE ports.

Note: The 10GE SFP+ port on the E7-20 SCP card does not support GE SFP modules.

See the following Calix Quick Tip bulletins on the Calix website for information on the support for SFP modules and direct-attach cables:

- *Calix E7 Pluggable Transceiver Module Support*
- *CAB-12-023 - Pairing Bidirectional SFPs (to Support Single-Fiber Ethernet Links)*

Module restrictions

The following module restrictions apply:

- Unless otherwise specified, SFP sockets in all Calix products accept any SFP module.
- GE-12 cards (E7-2) support any Calix-keyed 1GE SFP module in the SFP sockets.
- Copper SFP modules should be purchased from Calix to ensure correct operational behavior.
- The 10GE, 2.5GE, GPON OIM, and CSFP modules must always be purchased from Calix.

E7-2 Fan tray assembly

Note: The first and second-generation E7-2 Fan Tray Assemblies (FTA1 and FTA2) are both referred to as FTA in the text below. FTA2 is backward compatible with FTA1, but provides more granular software fan speed control with release 2.2+, operating at preset fan speeds depending on the ambient temperature of the E7-2 shelf.

The FTA design includes four individual variable speed fans that maintain system cooling even with one fan failure. Airflow is from right to left (toward the line cards).



Calix E7-2 Fan Tray Assembly Front Panel

Each FTA is hot-swappable and can be quickly replaced by unlocking the sliding latch, removing the failed unit and sliding in the new FTA. Each FTA includes a field replaceable fan filter that slides in from the top of the FTA once it is removed from the E7-2 shelf.

Note: When the E7-2 is operated in a remote terminal that includes integrated cabinet filters, the FTA fan filter should be removed from the FTA to increase airflow within the E7-2.

System status indicators

The E7-2 system includes multiple visual status indicators located on the front of the E7-2 FTA. The FTA status indicators will remain dark until at least one card is inserted in the E7-2 shelf.

- Critical (CR) Alarm – RED LED indicates a critical alarm is present in the system
- Major (MJ) Alarm – RED LED indicates a major alarm is present in the system
- Minor (MN) Alarm – AMBER LED indicates a minor alarm is present in the system
- System Controller (MGT) – GREEN LED indicates E7-2 shelf has an active shelf controller
- 7-segment LCD display – not used in R1.0

An Alarm Cut-Off (ACO) button is also integrated on the front on the FTA.

System management ports

The E7-2 provides the following user interfaces for system management:

- **Command-line interface (CLI)** - Embedded CLI that supports local or remote access.
- **Web interface** - Embedded graphical user interface that operates on a standard web browser for local or remote access.

The E7-2 web and command line interfaces, along with Calix Management System (CMS) allow you to manage services while understanding their relationship to the network infrastructure.

The E7-2 has the following ports available for device management:

- **MGT-1** located on the FTA front panel for local craft access to management user interfaces. The connector is 10/100 Base-T Ethernet (RJ-45).
- **MGT-2** can be configured within any VLAN on any E7-2 Ethernet port interface for an in-band management interface.
- **MGT-3** located on the E7-2 rear panel for permanently connecting to the network back office management systems. The connector is 10/100 Base-T Ethernet (RJ-45).
- **MGT-4** located on the FTA front panel for local craft management. The connector is RS-232 serial: 38.4 kbps. (Supports the Calix command line interface (CLI) only.)

Power inputs: The E7-2 is equipped with dual power inputs to support redundant –48VDC power feeds (A and B), switching between them when one source fails.

Alarm inputs and outputs: The E7-2 supports eight external alarm input/output (I/O) positions via wire-wrap pins located on the E7-2 rear panel. The eight external alarm positions include seven inputs and one output position.

BITS Timing: The E7-2 can be synchronized to a local traceable clock using the Building Integrated Timing Supply (BITS) composite clock timing wire-wrap pins located on the back of the E7-2 shelf. Up to 10 shelves may be connected serially. The BITS interface is designed to support both DS1 and E1 inputs. See *T1/E1 PWE3 Service* (on page [451](#)).

Operational LEDs: The E7-2 is equipped with the following LED operational indicators:

- **Line card LEDs** - Located at each card slot on the front panel to indicate when at least one port is active (green) and when a card is in standby (orange).
- **Port status LEDs** - Each GPON Optical Interface Module (OIM), GE-SFP, 10GE-XFP module sockets have a combined link status/activity LED below the socket.
 - Green LED (Ethernet ports) is solid green when a link is established and blinks at a variable speed to indicate traffic load.
 - Green LED (GPON ports) blinks steadily when the first ONT is ranging on a GPON port and remains lit when at least one ONT is in service.
 - Red LED indicates a module is inserted into a socket which is directed to the backplane (only x2 or x4).
- **Alarm LEDs** - Located on the fan module front panel; indicate critical, major, and minor alarms.

E7-2 Dual-line card operation

The E7-2 line cards are designed to operate at full port count and capability when deployed in a single-card system, and also operate as a protected pair when deployed in a dual-card system.

Line cards and system controller status

In dual-line card systems, one card is designated as the system controller which manages alarms, configuration, and performance monitoring. A "*" is shown next to the card label in the web interface and in the CLI show card command results. The other card is in a standby status.

Backplane ports

For an E7-2 system with more than one card installed, at least one 10-Gigabit Ethernet port must be used as an inter-card data link. When an Ethernet port is used for this link, its front-panel connector is disabled.

The selection of backplane ports within the E7-2 system affects the backplane ports on both cards, ensuring that the settings are compatible. You can choose to direct the x2 or x4 port to the backplane.

When a front-facing port is reassigned as a back-facing port, its port settings are no longer controlled by provisioning. When the port is reassigned to front-facing, its attributes revert to the default (all auto) setting. See *Configuring Backplane Links*.

Operation Domains Within the E7-2 Modular Chassis

The following diagram shows the logical partitioning of the E7-2 modular chassis cards into two distinct domains:

- **The Modular Chassis Controller (MCC) Cards** participate in the domain that supports the following:

Primarily:

- Uplinks to the network
- Connects to transport rings
- Aggregates other Calix access nodes (for example, a C7 providing TDM voice gateway services)


Secondarily:

- Allows subscribers to access services

Port support:

- 1GE and 10GE Network ports; support protection protocols (ERPS, LAG, RSTP):
 - ♦ Transport; ERPS or RSTP protection
 - ♦ Uplinks; RSTP and/or LAG protection
 - ♦ Aggregation; protected downlinks w/RSTP or LAG
- SFP sockets:
 - ♦ 1GE or 2.5GE modes
 - ♦ Stacking ports at 2.5GE rate for MC
- SFP+ sockets:
 - ♦ 1GE or 10GE modes
 - ♦ Stacking ports at 10GE rate for MC
- Support for subscriber DSL drops; Access
- **The Modular Chassis Expansion (MCE) Cards** only participate in the domain that supports the following:
 - Allows subscribers to access services
 - **Port support:**
 - 1GE and 10GE Access ports; no support for protection protocols
 - ♦ Aggregation; unprotected point-to-point links on SFP and SFP+ sockets
 - SFP sockets:
 - ♦ 1GE or 2.5GE modes
 - ♦ Stacking ports at 2.5GE rate for MC

- SFP+ sockets:
 - ♦ 1GE or 10GE modes
 - ♦ Stacking ports at 10GE rate for MC
- Support for subscriber DSL drops; Access



Modular Chassis Controller cards

The MCC cards provide the centralized processing and switching functions for the entire E7-2 modular chassis. Because of the integrated nature of the MCC, it is the core component of both the Management and Networking operations. The cards in this shelf also have integrated, network facing transport / uplink ports, as well as the option of being subscriber ports. The two MCC cards are in shelf #1 of the modular chassis.

The E7-2 MCC supports a flexible set of network topology protocols for use in aggregation, ring-based transport, and uplink applications:

- ITU G.8032 Ethernet Ring Protection Switching (ERPS)
- IEEE 802.1w Rapid Spanning Tree Protocol (RSTP)
- IEEE 802.3ad/802.1AX Link Aggregation (LAG)
- VLAN and MAC Switching
- Tag Actions
- Traffic management and QoS
- Ingress rate limiting
- Egress Queuing and scheduling

Modular Chassis Expansion cards

The function of the E7-2 MCE cards is as follows:

- Provide the fiber interfaces to the customer premises demarcation points
- Operate as the untrusted subscriber service delivery access domain
- Support 10Gbps full duplex connections to each SCP-10GE

Interface roles

Each GE and 10GE interface in the E7 has one of the following configuration-role types:

Trunk: A port connecting to other equipment belonging to the service provider or to another service domain with consistent VLAN tagging levels. These ports may also be referred to as Network ports or Provider ports in industry standards. These ports support outer VLAN tag plus MAC switching.

Examples of trunk ports are 10G ERPS transport ports and GE uplinks. Trunk ports can be configured for link aggregation, RSTP, or ERPS.

Edge: A port facing customer equipment or facing reduced functionality devices, alternative administrative domains, or managed CPE. Generally, this E7 Ethernet port interface is where all classification is first performed on ingress traffic (if customer facing). The E7 Ethernet interface is also expected to add, replace, or remove one or more VLAN tags on edge traffic. RSTP, and LAG networking protocols are supported.

Examples of an edge port would include GE ports to managed CPE, a GE/10GE port to external equipment which may use different tagging levels, or GPON ports.

Access: A port facing untrusted customer equipment or other devices serving subscribers. Generally, this port interface is where individual subscriber services are defined and enforced (bandwidth limits, security, multicast profiles). Networking protocols are not supported.

Examples of an access port would be point-to-point connections to subscribers or other devices serving subscribers.

VLAN support

By default, every E7 Ethernet port with a trunk or edge interface on the unit is a member of VLAN 1, the Native VLAN. The Native VLAN is available to pass any untagged traffic. You can provision an E7 Ethernet port interface to use a different existing VLAN as the Native VLAN.

To forward untagged traffic on E7 Ethernet ports with an access interface, an add-tag action must be applied to untagged frames, assigning the traffic to a designated VLAN.

Tagged traffic that does not match any of the tagging criteria is dropped.

For modular chassis nodes, any VLAN created on the system is automatically mapped to the Stacking Ports. The remaining port interfaces in the system must be a VLAN member for traffic to pass on the VLAN through the interface.

Interface role configuration guidelines

	Trunk	Edge	Access
E7-2 MCC	X	X	X
E7-2 MCE			X
ONT Ethernet Port			X
Tag Actions		X	X
Native VLAN	X	X	
Networking Protocol (RSTP, ERPS*, LAG)	X	X	


*Edge ports only support RSTP and LAG networking protocols, NOT ERPS.

Modular Chassis Configuration

The Calix E7-2 modular chassis Ethernet Service Access Platform (ESAP) is an optical access platform designed for high density, high bandwidth fiber deployments. With a backplane providing 100 Gbps to each line card slot in two slot increments, it's ready for the new technologies and the bandwidth needs of tomorrow, while cost effectively meeting service delivery needs of today.

Up to ten* E7 shelves can operate as a single Modular Chassis (MC) system. The MC node is intended to provide additional subscriber drop shelves, known as Modular Chassis Expansion (MCE) shelves, for high density GPON and point-to-point GE applications, where the MC operations and protocol support within the MC can be optimized for the subscriber direct aggregation application.

***Note:** Only 5 shelves are recommended for MC nodes populated with VDSL2-48C cards in the MCC shelf and MCE shelves.



The MC node shares unified management interfaces (CLI, Web GUI, CMS, SNMP) where all shelves, slots, cards, and ports are visible from a single management session. New service configurations are automatically distributed across the MC stacking ring and stored within a shared database of service profiles. For example, all VLANs created on any interface in the E7 MC are automatically distributed to all other shelves in the MC.

E7 Modular Chassis Components:

- **Modular Chassis Controller (MCC)** – system control for the MC, designated shelf “1”
- **Modular Chassis Expansion (MCE)** – a non-controller shelf in the MC (shelf 2-10)
- **MC Shelf ID** – an index number from 1-10 identifying the shelf number within the MC
- **Stacking Ring** – a closed ring connecting all E7 shelves. When optical card(s) populate any of the slot(s) in the MC, 10GE links in SFP+ ports must be used as Stacking Ports. When only copper cards populate the slots in the MC, 2.5GE links in SFP ports can be used as Stacking Ports (10GE links in SFP+ ports are also an option for all copper E7-2 MCs).
- **Stacking Port** – the port on which the stacking ring is enabled on each E7 line card
- **Transport Ring** – ERPS or RSTP ring connecting the E7 MC to other nodes. Network protocols supporting the transport ring and redundant uplink from the MC are supported only in the MCC.

See "Turning Up a Modular Chassis System" in the *Calix E7 User Guide* for information on initially configuring several E7s into a modular chassis and see the *Calix Method of Procedure (MOP): Migrating Standalone E7 Systems to an E7 Modular Chassis* for instructions on how to migrate multiple standalone E7 systems into a unified E7 modular chassis.

Configuration guidelines

Shelves and cards within the Modular Chassis participate in the Control, Network, and Access domains of the system.

- The MCC shelf provides the redundant Control domain for the Modular Chassis and also provides the Network domain connectivity.
- The MCC and MCE shelves span the Access domain to deliver services.
- The MCE shelves only participate in the Access domain. Ethernet interfaces in the Access domain are intended to be tied to the back-office system; part of flow-through service provisioning.

The following guidelines must be adhered to when configuring a modular chassis system:

- A Modular Chassis is managed through a single IP address.
- The MCC must have two-cards defined in it, with exception of the VDSL2-48 Overlay line card which occupies two slots (double-wide) in each individual E7-2 shelf. When the VDSL2-48 card is deployed in the MCC, the MCC will not have a standby controller and is not able to support database redundancy.
- MCE shelves may have one or two cards, as required, where the ports are considered "subscriber ports" and cannot run link protection protocols such as RSTP, LAG, or ERPS.
- When a VDSL2-48 Overlay line card is in the MCC or MCE shelf, an inter-card data link is not required because it requires no backplane communication to a second card in the same E7-2 chassis.

- When a GE-24 (AE card) or GPON-8 line card is in the MCC or MCE shelf, an inter-card link is not required because these cards have internal virtual interfaces through the backplane, leaving the SFP+ ports available for the stacking ring and uplink, or transport connections.
- Each shelf in the MC is identified by a shelf number, from 1 to 10 where shelf 10 is identified as "0" on the fan-tray assembly and the MCC shelf ID is "1." The shelves may be numbered in any order around the Stacking Ring.

Note: For E7-2 modular chassis with VDSL2-based MCCs in an ERPS ring, the maximum VDSL2 combo subscribers/node and nodes/ring are lower than in other configurations. See *xDSL Support Capacities* (on page [593](#)) for more detailed information.

- Any given port or interface in an MC is identified as *shelf/slot/port*.
- One port in each line card in the MCC and MCE shelves is used for the Modular Chassis inter-shelf control path (Stacking Ring).
- The stacking ring can be implemented in one of the following speeds:
 - At 2.5GE rate using SFP sockets when an MC has only VDSL2 cards
 - At 10GE rate using SFP+ sockets when an MC has either all optical cards or a mix of optical and VDSL2 cards, where the VDSL2 cards are not in the MCC slots

Note: The stacking ring does not support a 1GE data rate, or a mixed data rate (2.5GE/10GE).
- **For 10GE stacking ring connectivity:**
 - Only E7 10GE ports (SFP+ X3 and X4 ports) may be configured as Stacking Ports. If required by the network design, use 10GE-4 or GPON-4 cards in the shelf that will be designated as the MCC to provide additional 10GE ports for the network uplink or 10GE transport. One Stacking Ring is supported per MC where each shelf has two 10GE stacking ports:
 - For two-card shelves, use port X3 on each line card.
 - For one-card shelves, use ports X3 and X4 on the line card.
 - The X4 port can only be used as the stacking port on the MCC shelf if the backplane links are configured for 10G-A, because 10G-B and 20G use port X4 as a backplane port. The exception to this rule is for the GPON-8 or GE-24 cards because these cards have internal virtual interfaces through the backplane, leaving the SFP+ ports available for the stacking ring and uplink, or transport connections.
 - A Calix 10GE Direct Attach copper cable must be used for stacking port connections.
- **For 2.5GE stacking ring connectivity:**
 - When the MCC has VDSL2 card(s) installed and they must provide a 10GE port for a transport ring interface or uplink, a 2.5GE stacking ring is required.
 - The E7-2 supports 2.5 Gbps pluggable module interfaces in the SFP ports of the VDSL2-48C and VDSL2-48 cards.

- For 2.5GE stacking rings, the default stacking interfaces on VDSL2 Linecards will be the first SFP interface designated as “G1”. For a single VDSL2 Linecard, the Secondary Default Stacking port will be the second SFP interface designated as “G2.”
- A Calix Direct Attach copper cable must be used for 2.5GE stacking port connections.
- Although the GPON-4 and 10GE-4 cards support 2.5GE modules in the SFP sockets, including the use of Direct Attach cables, these ports cannot be used for stacking rings.
- All network connections (ERPS, RSTP, LAG, DHCP servers, multicast routers, IGMP-enabled video servers, network-facing routers) are made on the Modular Chassis Controller (MCC) shelf.
- All outer VLANs for each service defined on the MC are automatically added to the Stacking Ring domain.
- To retain the E7 IP address for the MCC upon resetting the database, use the "keep-craft-fe" option.
- The Stacking Ring must be a closed ring during normal operations.
- Port mirroring source and destination ports must be on the same E7 line card.
- Split Horizon is only enforced between ports on the same line card.
- RSTP Node Protection pair is not supported with a E7-2 MCC.
- An IGMP Querier (multicast source) connection to an MCE port is not supported.
- A complete copy of the provisioning database is maintained on every card within the MC, providing full redundancy across all shelves.
- An MCE shelf establishes communications with the MCC over the stacking ring when the following conditions are met:
 - The MCE is provisioned through shelf creation at the MCC.
 - A Calix-brand Direct Attach copper cable is connected.
- The Craft FE ports can be enabled and configured on MCC shelves, however, the Craft FE ports on MCE shelves are disabled when MCE shelves are in communication with an MCC shelf. If MCE shelves lose communication with the MCC shelf, the Craft FE ports on the MCE shelves become enabled.
- The E7 uses the First Reserved VLAN ID for the stacking ring control VLAN, as configured in the system settings (the default is 1002). Changing this value after a stacking ring has been configured is service affecting for subscribers in the MCE shelves.
- The MCE shelf is used for GE and GPON subscriber port expansion.
- Access Domain Ethernet interface parameters (MCE shelves):
 - The Ethernet interface Role parameter is limited to Access.
 - The DHCP Trust parameter is limited to disabled or N (untrusted).
 - Split Horizon Forwarding parameter is disabled.
 - The BPDU Guard parameter is enabled, by default.


- The Native VLAN is not supported, however tag actions can be used to match untagged traffic to a VLAN ID.
- The MCE Ethernet interfaces cannot be designated as a VLAN Router Interface. That is, a static multicast router ('mrouter') interface is not supported as a multicast destination on MCE shelves. Multicast routers are not allowed to source multicast traffic into the network on an Access domain Ethernet port.
- A subtended E3 is supported from an MCE Access port for simple point-to-point connections, as the E3 is seen as just another version of an MDU ONT. If it is necessary to have RSTP or LAG configured to the E3, then the E3 must be connected to the MCC.

Interface role configuration guidelines


	Trunk	Edge	Access
E7-2 MCC	X	X	X
E7-2 MCE			X
Tag Actions		X	X
Native VLAN	X	X	
Networking Protocol (RSTP, ERPS*, LAG)	X	X	

*Edge ports only support RSTP and LAG networking protocols, NOT ERPS.


MC with 10GE stacking ring and redundant controller



MC with 2.5GE stacking ring and redundant controller



MC with 2.5GE stacking ring and non-redundant controller



E7-2 Inter-Card Communication and Backplane Link

(E7-2 and Modular Chassis systems only)

This topic shows you how to configure the Ethernet ports that are dedicated to inter-card communication within the E7.

When an E7 system has more than one card installed, at least one 10-Gigabit Ethernet port must be used as an inter-card data link, with the exception of the GE-12, GE-24, GPON-8, and VDSL2 r2 cards which use a virtual interface through the backplane. When an Ethernet port is used for this link, its front-panel connector is disabled.

The selection of backplane ports within the system affects the backplane ports on both cards, ensuring that the settings are compatible. You can choose between the following settings:

- **10G-A** (default setting)
- **10G-B**
- **20G**
- **None** (only to be used for single-card systems)

Backplane connections for E7 line cards

Setting	Ports Used	Cards
10G-A	XFP 2	• 10GE-4 • GPON-4 r2
	None	• GE-12 • GE-24 • GPON-8 • VDSL2-48C r2 • VDSL2-48D r2
	SFP+2	• VDSL2-48C r1 • VDSL2-48D r1
10G-B	SFP+2	• 10GE-4 • GPON-4 r2 • GE-12
	None	• GE-24 • GPON-8 • VDSL2-48C r2 • VDSL2-48D r2
20G	XFP 2 and SFP+ 2	• 10GE-4 • GPON-4 r2
	SFP+ 2	• GE-12
	None	• GE-24 • GPON-8
None	None	All

Configuration guidelines:

- For 10GE-4 and GPON-4 r2 cards, the 10G-A setting redirects the XFP 2 port to the backplane. The 10G-B setting redirects the SFP+ 2 port to the backplane. The 20G setting redirects both the XFP 2 and SFP+ 2 ports to the backplane.
- For the GE-12, the 10G-A setting is established using an internal port and therefore does not use an external 10GE port. The 10G-B setting redirects the SFP+ 2 port to the backplane. The 20G setting uses an internal port and redirects the SFP+2 port to the backplane.
- For the VDSL2-48C r1 and VDSL2-48D r1, the 10G-A setting redirects the SFP+ 2 port to the backplane. The 10G-B and 20G settings are not supported for these cards.
- For the GE-24, GPON-8, VDSL2-48C r2, and VDSL2-48D r2 all settings are established using internal ports and therefore these cards do not use any external 10GE ports.

- Any change to the backplane link affects the backplane ports on all E7-2 and Modular Chassis line cards.
- For those 10GE Ethernet ports that are directed to the backplane, the Ethernet port settings are no longer controlled by provisioning. When the port is reassigned to the front-facing port, the port attributes revert to defaults.
- A change to the backplane link assignment can only be made if the port that is to be directed to the backplane meets the following conditions:
 - The port is not provisioned as an ERPS port.
 - The port interface is not part of a VLAN membership.
 - The card where the port is located is up and running with no existing alarms.
- When a backplane-link setting is changed, all cards in the E7 reset and then operate with the new setting.
- The E7 supports an ERPS domain spanning the backplane link, configured as a single ERPS domain with one interface on each card. Any VLANs associated to the ERPS domain are carried on the backplane link.
- For modular chassis nodes, any VLAN created on the system is automatically mapped to the Stacking Ports and the Backplane Ports. The remaining port interfaces in the system must belong to a VLAN membership for traffic to pass on the VLAN.

Backplane Port Use Cases

The following use cases illustrate how backplane configuration works for various scenarios. The 10GE port numbers in these examples presume the cards are 10GE-4 or GPON-4 cards. For GE-12, GE-24, GPON-8, and VDSL2 cards the port numbering would be different, as described in Backplane Port Configuration. The following use cases are considered:

- Case 1. Default for single card in a shelf
- Case 2. Send all 10GE ports to front (single card system)
- Case 3. Adding a second card to a shelf
- Case 4. Changing which ports are backplane ports
- Case 5. Using two backplane ports with 20 Gbps bandwidth
- Case 6. ERPS domain spanning across two cards
- Case 7. Mis-configuration during commission of a second card

Case 1: Default for single card in a shelf

When you commission a new E7 and create a default database, the backplane port configuration defaults to “First Ports.” When a card object is created in the user interface, port X2 is configured as a backplane port. The remaining three 10GE ports are available for association with an interface or link aggregation group. You may not change port X2’s interface association directly, as validation prevents this based on conflict with the system-level provisioning.

Case 2: Send all 10GE ports to front (single-card system)

When you change the system backplane port configuration to “No Ports,” the interface association of port X2 updates to the matching Ethernet interface. No ports are configured as backplane ports. You may then add VLAN associations to Ethernet interface X2 to add services to this port.

Case 3: Adding a second card to a shelf

When you add a second card to the E7 shelf, the second card’s X2 port defaults to a backplane port, if the system backplane port is set to the default configuration of “First Ports.” The two cards are connected by the backplane trace. As you provision services on the two cards, the system determines which VLANs are common to both cards and switches these VLANs across the backplane.

If the system backplane port is set to the configuration of “No Ports” before adding the second card, then no 10GE ports on the second card are assigned to the backplane. The system raises an alarm indicating there is no data path between the cards.

Case 4: Changing which ports are backplane ports

When you change which ports are directed to the backplane interface by changing the system backplane port configuration, the following transitions are possible:

- a. “First Ports” to “Second Ports”** changes the Backplane port on each card from X2 to X4. The X2 port on each card is re-set to forward facing Ethernet interface X2. The X4 port on each card is configured as a backplane port.
 - If the X2 port had previously been associated to its Ethernet interface, any services that are provisioned against the Ethernet interface remain, but are disabled.
 - If the X2 port had previously been associated to a LAG Interface, any services provisioned against that LAG interface may be degraded (due to reduced bandwidth) or disabled.
- b. “Second Ports” to “First Ports”** changes the Backplane port on each card from X4 to X2. With this transition the X4 ports are returned to their corresponding Ethernet interfaces and the X2 ports are configured as backplane ports.
- c. From Any Setting to “No Ports”** reassigns all 10GE ports to front facing, with no ports connected to the backplane traces.
 - If more than one card is provisioned, the system raises an alarm indicating No Data Path between the cards.
 - If any ports were configured as backplane ports, they are re-assigned to their corresponding Ethernet interface.

- d. **“No Ports” to Any Other Setting** re-configures one or both ports from front facing to backplane ports.
 - If the X2 or X4 ports had previously been associated to its Ethernet interface, any services that are provisioned against the Ethernet interface remain, but are disabled.
 - If the No Data Path alarm had been raised in the system, it is cleared by this change.
- e. **From Any Setting to “Both Ports”** configures both available 10GE ports on each card to backplane ports. If the X2 or X4 ports had previously been associated to its Ethernet interface, any services that are provisioned against the Ethernet interface remain, but are disabled.
- f. **From “Both Ports” to “First Ports” or “Second Ports”** re-configures one of the two backplane ports on each card to front facing. No alarm is raised, as the system continues to have a backplane data path between the cards. The ports which are no longer configured as backplane ports are re-configured to their Ethernet interfaces.

Case 5: Using two backplane ports with 20Gbps bandwidth

When you direct two ports on a card to the backplane, a Link Aggregation (LAG) is automatically implemented between the ports to avoid routing loop issues.

Case 6: ERPS domain spanning across two cards

When you configure an ERPS domain with one interface on each card, the domain spans the backplane links. Any VLANs associated to the ERPS domain are carried on the Backplane link. The ERPS provisioning is sent to the SIT on each card, which configures ERPS between the specified card local interface and the Backplane Interface. This effectively creates two functional ERPS domains (one on each card), acting as a single ERPS domain across the pair of cards.

Any VLANs associated to the ERPS domain are automatically carried across the Backplane Interface on each card. Also, any VLANs that are not associated with the ERPS domain yet have a presence on both cards are automatically added to the backplane interface.

Case 7: Mis-configuration during commission of a second card

When a system has a backplane port configuration of “No Ports” and a second card is provisioned, either by your action or by card arrival, a system alarm is raised (for example, “No Data Path”). The presence of this alarm does not prevent provisioning of services on the shelf, since you may wish to move services from a 10GE port on the first card onto a port on the second card before switching the first card’s port to the backplane. This alarm remains until you change the backplane port configuration setting to another value, or until the provisioning for the second card is deleted.

Mixed Copper and Optical Card Deployment

When a VDSL2-48C or VDSL2-48D card is deployed in an E7-2 with any other line card, the card “SFP+ 2” port (right-hand SFP+ port on the faceplate, X4 in the user interface) is used for inter-line-card connectivity via the chassis’ backplane, thus only “SFP+ 1” port (left-hand SFP+ port on the faceplate, or X3 in the user interface) remains available for external chassis/shelf connectivity. Accordingly, a dual line card E7-2 chassis, deploying one or two VDSL2-48C or VDSL2-48D line cards, have a fixed 10Gbps backplane capacity.

For mixed copper and optical cards deployment within the same E7-2 chassis, where one slot is populated with a VDSL2-48C or VDSL2-48D card, the behavior of the backplane connectivity is dictated by the type of card deployed in the 2nd slot.

- An optical / transport line card (10GE-4 or GPON-4); as already mentioned, the “SFP+ 2” port is used on the VDSL2-48C or VDSL2-48D card for backplane connectivity, the ‘XFP 2’ from the optical / transport line card is used, leaving one XFP port (‘XFP 1’, left-hand side XFP port, or X1 in the user interface) and the two SFP+ ports on the card available for external chassis / shelf connectivity.
- An optical / access line card (GE-12 or GE-24); the same “SFP+ 2” port is used on the VDSL2-48C or VDSL2-48D card for backplane connectivity, and since no XFP ports are available, the backplane connection is established without redirecting any ports to the backplane. Due to the 10GE backplane connection available to a VDSL2-48C or VDSL2-48D, a “Second Ports” choice for the GE-12 or GE-24 (X4 within the user interface) is not allowed, leaving the two SFP+ ports on the card available for external chassis / shelf connectivity.

The above choices for intra-chassis, backplane connectivity among optical and copper E7-2 line cards provide:

- Wider commonality and consistency among different line cards coexisting with VDSL2-48C or VDSL2-48D within an E7-2 chassis.
- XFP and SFP+ options for optical / transport line card (10GE-4 or GPON-4) within E7-2 chassis coexisting with VDSL2-48C or VDSL2-48D.
- Always availability of SFP+ ports for inter-chassis stacking ports in modular chassis deployments.
- Minimize faceplate port loss when mixing VDSL2-48C or VDSL2-48D and GE-24 or GE-12 line cards in an E7-2 chassis.


E7-2 Chassis

The E7-2 two-slot chassis has two rack mounting options and the following features on the rear panel:

- Two power connectors
- Frame ground lugs
- Alarm Interface Module (AIM)
- BIT timing interface
- Ethernet management port
- Copper subscriber interfaces

Rack mounting options

The E7 is designed to be rack mounted in either 19" or 23" equipment racks. The following graphic shows both a 19" and a 23" rack mount bracket, that ship with the unit. The service provider should choose the bracket that works best for their installation requirements. In addition, the shelf has several mounting holes along the side, allowing it to be flush or mid mounted. This is most important in cabinet retrofit applications where airflow may be blocked if the unit is mounted in the forward position.




Calix E7-2 Mounting Brackets

In some 19" racks, it may be necessary to flip the 19" mounting brackets and protrusion mount the shelf if mid-mount is not an option.

Power and ground connections


The E7 supports two redundant DC power inputs located on the back of the unit that can be connected to A and B power sources for automatic switchover. Equipment ground is also located on the back of the unit.

The E7 supports an input power range of -42.5 VDC to -72 VDC. The following diagram shows the location of the power and ground connections on the rear panel.



Alarm I/O

The E7 supports 7 input alarms and 1 output alarm, located on the back of the unit as wire wrap pins.



The alarm interface pins are mapped as follows:


Alarm input / output pin mapping							
OUTPUT	INPUT						
OUT+	AL1+	AL2+	AL3+	AL4+	AL5+	AL6+	AL7+
OUT-	AL1-	AL2-	AL3-	AL4-	AL5-	AL6-	AL7-

Alarm input (AL1+/- through AL7+/-): pins detect contact closure only and will raise an alarm.

Alarm output (OUT+/-): normally open, this pair of pins can be configured to interface with various alarm management devices such as lights or horns located in the central office.

BITS timing and Ethernet management port

The E7 can be synchronized to a local traceable clock using the BITS timing wire warp pins located on the back of the shelf. The following diagram shows the BITS external clock input / output pin field and the location of the Ethernet management port (MGMT 3).



The BIT pins are mapped as follows:

External BITS Timing Input / Output Pin Mapping			
In (A)	SHLD	TIP+	RNG-
In (B)	SHLD	TIP+	RNG-
TIP+	RNG-	TIP+	RNG-
Out (A)		Out (B)	

Up to ten shelves can be connected together to share a local BITS input. A BITS Chaining Cable available from Calix includes connectors to attach the redundant BITS output pins from one shelf to the redundant BITS input pins on a second shelf.

Copper Connectors

The E7 shelf includes two standard 25-pair RJ-21 connectors per slot. These connectors may be used by any card installed in the system.

E7 LED Behavior

Line card LEDs

- Service (SRVC) indicates at least 1 port has been provisioned on this slot
- Control (CTRL) green LED indicates active controller, amber LED indicates standby controller
- Fault (FAIL) red LED indicates the card is in fault
- Cards have additional LED states at boot time (see table below)

Port LEDs

- Green LED (Ethernet) that stays on indicates a link with no activity, flashes during activity.
- Green LED (GPON port) that stays on indicates at least 1 ONT is ranged, blinks when the first ONT is ranging.
- Green LED (DSL port) that stays on indicates at least one DSL subscriber port is synched up and operating correctly (the system is up, operational, and ready to or able to pass traffic).

- Green LED (POTS port) that stays on indicates, indicates at least one POTS subscriber port is operating correctly (the system is up, operational, and is off-hook).
- Red LED (10GE) indicates a module is inserted into a port that is redirected to the backplane.
- Amber LED (10GE) indicates a stacking ring port for modular chassis systems.
- Inserting a supported module into a port causes the LED to blink green 3 times, indicating module recognition.
- If no blinking occurs on module insertion, the module is not supported/recognized; An alarm is present in this situation.

Line card LED boot sequence

CTRL LED	SRVC LED	FAIL LED	State Description
Active Card Boot Sequence			
Green	Yellow	Red	Power on
off	Off	Red	NB Execution
1 short green	Off	Red	UB Execution
2 short green	Off	Red	Booting Kernel
3 short green	Off	Red	Application Loading
2 short, 1 long green	Off	Off	Application Initializing
1 short, 1 long green	Off	Off	Database Loading
Green 1 Short Off	Off	Off	Database Activation
Green	Off	Off	Application Startup complete, no services defined
Green	Green	Off	Application startup complete, services defined
Standby Card Boot Sequence			
Green	Yellow	Red	Power On
Off	Off	Red	NB Execution
1 Short Yellow	Off	Red	UB Execution
2 Short Yellow	Off	Red	Booting Kernel
3 Short Yellow	Off	Red	Application Loading
2 Short 1 Long Yellow	Off	Off	Application Initializing
1 Short 1 Long Yellow	Off	Off	Database Loading
Yellow 1 Short Off	Off	Off	Database Activation
Yellow	Off	Off	Application Startup complete, no services defined
Yellow	Green	Off	Application startup Complete, services defined
Other States			
Off	Off	All Short Red	Equipment Mismatch
Green blinking	Yellow blinking	Red blinking	E7-20 card inserted in E7-2

CTRL LED	SRVC LED	FAIL LED	State Description
Yellow rapid blinking	Yellow rapid blinking	Red rapid blinking	E7-20 SCP card is partially inserted in chassis. Push the card fully into the chassis where the card reboots.
Pattern & Color Reflects Card Status	All Short Yellow	Off	Flash Write in Progress (database or program update)
Off	Off	2 Short Red	Application Initiated Shutdown
Green	Pattern & color Reflects Service and Flash Write	Red	Equipment Failure for Active Card
yellow	Pattern & color Reflects Service and Flash Write	Red	Equipment Failure for Standby Card
1 Short Green	Off	Off	No Database in Flash on Active Card
1 Short Yellow	Off	Off	No Database in Flash on Standby Card
Off	Off	2 Short 1 Long Red	Waiting for card to cool down before loading application

E7 Power Subsystem


This topic describes how the E7 power subsystem operates.

- The E7-20 has three power zones with (A/B) feeds for each.
- The E7-2 has a single power zone per shelf with (A/B) redundant power connections.

Each power source is independently monitored and the system issues an alarm if power is lost on either the A or B source. The E7 does not have centralized system-level power supplies.

- Each E7 card and fan tray assembly has connections to both power feeds A and B, with simple diode circuits to combine the A and B feeds into a single power source for the card.
- All -48V DC input power is down-converted to the appropriate voltage and current on the card.
- Power conversion is done in multiple locations on the card for optimal power management targeted to specific active components on the card (for example: CPU, GPON MAC, optics modules).
- The DC power converters are simple circuits and among the most reliable devices on the card. There has never been an in-field E7 power converter failure.

- The redundancy of the system can be tested by turning off the circuit breaker feeding the A or B side of each power zone.



E7-20 Platform Description

The Calix E7-20 Ethernet service access platform is a Large Form Factor (LFF) chassis that delivers high-performance Ethernet-based transport, aggregation, and FTTP services from a controlled environment deployment.

The wide array of high-performance applications delivered by the E7-20 includes 10GE Ethernet transport, high-density residential triple play services over GPON, AE, and point-to-point Ethernet, Metro Ethernet Forum (MEF) compliant business services, mobile backhaul, and protected GE aggregation.

The Calix E7-20 system consists of the following components:

- 20-slot Chassis
- Removable Fan Tray Assembly — version 1 (FTA) or version 2 (FTA2)


Note: Unless otherwise notes, "FTA" refers to both versions. Systems operating with FTA2 must be running system software version R2.5 or later.

- Fan Filter
- Switch Control Processors — version 1 (SCP) or version 2 (SCP2)

Note: Unless otherwise notes, "SCP" refers to both versions.

Note: Systems operating with SCP2 cards must use the FTA2.

- GE-24x Line Cards
- GPON-4x, GPON-8x, and GPON-16x Line Cards
- Removable Fiber-Management System



Note: SCP version 1 (SCP) cards are shown in the image above. SCP2 cards provide eight (8) 10GE ports.

The vertical size of the LFF chassis permits placement of 3 LFF chassis in a 7-foot rack, with at least 4RU height left over for power panels or other equipment.

E7-20 Front

The Calix E7-20 line cards, SCP cards, and fan tray install into the front of the E7-20 shelf. When equipped with line cards, the E7-20 front panel also terminates all system line interfaces (fibers). The E7-20 fan tray is equipped with Craft Ethernet and serial ports for management and features several system status indicators. The front of the chassis also includes a removable fiber management system.

E7-20 Rear

The E7-20 power inputs (A/B), ground, Craft management, and BITS timing interfaces are located on the rear of the E7-20 shelf.

Rack mounting options

The E7-20 provides the following mounting options:

- Standard US Telco with 5-inch maximum protrusion from the front of the rack mounting rails to farthest forward point of the equipment, all cables included.
- Mid-mount with 10-inch maximum protrusion from the front of the rack mounting rails to farthest forward point of the equipment, all cables included.

E7-20 ETSI rack mounting kit to provide a 75 mm forward offset limitation from the rack mounting rail.

Common control cards

The SCP-10GE cards provide the centralized processing and switching functions for the entire E7-20 Chassis. Because of the integrated nature of the SCP, it is the core component of both the Management and Networking operations. These cards also have integrated, network facing transport / uplink ports. The two common control slots in an E7-20 chassis are located between universal line card slots 10 and 11. The E7-20 chassis is designed with a dual-star backplane topology, where each common control card slots has a direct 100 Gbps high-speed backplane connection every universal line card slot, as shown in the diagram above. Each SCP-10GE provides 10Gbps full duplex connections to each line card slot, with one SCP-10GE in active mode and the second SCP-10GE in stand-by mode.

The E7-20 SCP supports a flexible set of network topology protocols for use in aggregation, ring-based transport, and uplink applications:

- ITU G.8032 Ethernet Ring Protection Switching (ERPS)
- IEEE 802.1w Rapid Spanning Tree Protocol (RSTP)
- IEEE 802.3ad/802.1AX Link Aggregation (LAG)
- VLAN and MAC Switching
- Tag Actions
- Traffic management and QoS
- Ingress rate limiting
- Egress Queuing and scheduling

Universal line cards

The primary function of the E7-20 chassis is to provide the aggregation of high density fiber services, which is accomplished through 20 universal line card slots. The GPON-8x and GE-24x service access card functions are as follows:

- Provide the fiber interfaces to the customer premises demarcation points
- Operate as the untrusted subscriber service delivery access domain
- Support 10Gbps full duplex connections to each SCP-10GE

Interface roles

Each GE and 10GE interface in the E7 has one of the following configuration-role types:

Trunk: A port connecting to other equipment belonging to the service provider or to another service domain with consistent VLAN tagging levels. These ports may also be referred to as Network ports or Provider ports in industry standards. These ports support outer VLAN tag plus MAC switching.

Examples of trunk ports are 10G ERPS transport ports and GE uplinks. Trunk ports can be configured for link aggregation, RSTP, or ERPS.

Edge: A port facing customer equipment or facing reduced functionality devices, alternative administrative domains, or managed CPE. Generally, this E7 Ethernet port interface is where all classification is first performed on ingress traffic (if customer facing). The E7 Ethernet interface is also expected to add, replace, or remove one or more VLAN tags on edge traffic. ERPs, RSTP, and LAG networking protocols are supported.

Examples of an edge port would include GE ports to managed CPE, a GE/10GE port to external equipment which may use different tagging levels, or GPON ports.

Access: A port facing untrusted customer equipment or other devices serving subscribers. Generally, this port interface is where individual subscriber services are defined and enforced (bandwidth limits, security, multicast profiles). Networking protocols are not supported.

Examples of an access port would be point-to-point connections to subscribers or other devices serving subscribers.

Interface role configuration guidelines

	Trunk	Edge	Access
E7-20 SCP	X	X	
E7-20 GE-24			X
ONT Ethernet Port			X
Tag Actions		X	X
Native VLAN	X	X	
Networking Protocol (RSTP, ERPS, LAG)	X	X	

VLAN support

By default, every E7 Ethernet port with a trunk or edge interface on the unit is a member of VLAN 1, the Native VLAN. The Native VLAN is available to pass any untagged traffic. You can provision an E7 Ethernet port interface to use a different existing VLAN as the Native VLAN.

To forward untagged traffic on E7 Ethernet ports with an access interface, an add-tag action must be applied to untagged frames, assigning the traffic to a designated VLAN.

Tagged traffic that does not match any of the tagging criteria is dropped.

E7-20 Line Cards

The Calix E7-20 is equipped with 20 universal line card slots, supporting a flexible array of high capacity line cards. Service access cards can be mixed in any of the 20 universal line cards slots within the E7-20 chassis without restriction. All E7-20 cards use pluggable transceiver modules for optical and copper-based interfaces, including industry standard SFP, SFP+, CSFP, and OIM modules. Each port on each card has an LED to indicate an established link and data traffic activity.


E7-20 SCP-10GE Common Control Card

Interfaces:

- Four (GE SFP) interfaces
- One (10GE SFP+) interface
- One (10GE XFP) interface

The Calix SCP-10GE card can be installed into two of the common control slots. The E7-20 SCP-10GE switch control processor card provides centralized management control, switching, and network connectivity for the E7-20. Two 10-Gigabit Ethernet ports on each card provide the primary high speed network connectivity for the system, while four GE ports on each card provide alternative support for NxGE LAG uplink interfaces or Ethernet aggregation.

Note: The 10GE SFP+ port on the SCP card does not support GE SFP modules.



E7-20 GPON-4x Line Card

Interfaces:

- Four ITU G.984-compliant Gigabit Passive Optical Network (GPON) interfaces

The Calix E7-20 GPON-4x service access card provides multiservice capability over four GPON OLT ports, used to subtend up to 64 ONT's per OLT port, with a card capacity of 256 GPON ONT's. The GPON-4x card supports single-fiber Class B+ (28 dB link margin) optical distribution network:

- Standard Reach OIM
- Extended Reach OIM

E7-20 GPON-8x Line Card

Interfaces:

- Eight ITU G.984-compliant Gigabit Passive Optical Network (GPON) interfaces

The Calix E7-20 GPON-8x service access card provides multiservice capability over eight GPON OLT ports, used to subtend up to 64 ONTs per OLT port, with a card capacity of 512 GPON ONTs. The GPON-8x card supports single-fiber Class C+ (32 dB link margin) optical distribution network and forward error correction (FEC):

- Maximum 64:1 split per PON

ONTs can be distributed on any PON or card in any number as long as those two limits are not exceeded.

E7-20 GPON-16x Line Card

Interfaces:

- Sixteen ITU G.984-compliant Gigabit Passive Optical Network (GPON) interfaces

The Calix E7-20 GPON-16x service access card provides multiservice capability over sixteen GPON OLT ports, used to subtend up to 128 ONTs per OLT port, with a card capacity of 1024 GPON ONTs. The GPON-16x card supports single-fiber Class C+ (31 dB link margin) optical distribution network and forward error correction (FEC):

- Maximum 128:1 split per PON

ONTs can be distributed on any PON or card in any number as long as those two limits are not exceeded.

E7-20 GE-24x Line Card

Interfaces:

- Twelve (GE CSFP) interfaces that accommodate 24 GE ports

The Calix E7-20 GE-24x service access card provides up to 24 GE ports using pluggable SFPs (12 1GE / 2.5GE ports, single GE port per optical module) or CSFPs (Mode 1 and Mode 2, 24 GE ports, two GE ports per optical module). Each CSFP module supports two independent bidirectional transceivers (1490nm Tx / 1310nm Rx), each capable of operating at a rate of 1Gbps bi-directional. High-density, dual-port CSFP modules have a maximum range of 20km, while single port SFPs can be used to reach AE ONTs up to 60 km.

Note: The CSFP modules are mechanically compatible with SFP sockets but CSFP modules cannot be used in equipment that only supports SFP sockets. Further, CSFP Mode 2 modules cannot be used in equipment that supports CSFP Mode 1 (pre-standard) modules.

Note: E7-2 supports 2.5 Gbps pluggable module interfaces on the SFP ports of the GE-24 cards. The 2.5GE interfaces support equivalent functions and networking protocols as GE and 10GE interfaces.

Blank Cards

The E7-20 “Blank” cards are designed to fill all empty access card and SCP card slots. A common E7 BLANK card is used in both the E7-2 and E7-20 and must be placed in any empty slot to ensure proper airflow and cooling for normal operations. The blank card keeps “forced air” from escaping out the front of the shelf and is required to comply with Electromagnetic Interference (EMI) standards.

Note: Blank cards must be ordered separately.



ALERT! Failure to use Blank cards will significantly alter the airflow through the E7-20 chassis and may result in insufficient airflow and cooling of active cards. Lack of proper airflow may result in a combination of high temperature alarms, system shutdown, and premature failure of system components.

Pluggable optics

The E7 platform supports industry standard pluggable modules for all service and network interfaces, including ITU G.984 compliant GPON, SFP GE, XFP 10GE ports, and SFP+ 10GE ports.

Note: The 10GE SFP+ port on the E7-20 SCP card does not support GE SFP modules.

See the following Calix Quick Tip bulletins on the Calix website for information on the support for SFP modules and direct-attach cables:

- *Calix E7 Pluggable Transceiver Module Support*
- *CAB-12-023 - Pairing Bidirectional SFPs (to Support Single-Fiber Ethernet Links)*

Module restrictions

The following module restrictions apply:

- Unless otherwise specified, SFP sockets in all Calix products accept any SFP module.
- GE-12 cards (E7-2) support any Calix-keyed 1GE SFP module in the SFP sockets.
- Copper SFP modules should be purchased from Calix to ensure correct operational behavior.
- The 10GE, 2.5GE, GPON OIM, and CSFP modules must always be purchased from Calix.

ONT Interoperability

The E7-20 operates with the Calix 700 Family of ONTs and a select set of standards-based, third-party GPON ONTs. Calix is actively involved in the Full Service Access Network (FSAN) Group and Broadband Forum, which are pursuing GPON interoperability among the world's leading telecommunications services providers, independent test labs, and equipment suppliers.


Calix 700GX and 700GE ONTs support auto-detect technology, enabling support for GPON and GE from the same ONT, allowing service providers to manage service changes through software from a NOC / CO without requiring a truck roll to the residential or business subscriber location to replace the ONT or perform a manual ONT firmware upgrade.

A common set of ONT-delivered services are supported on the GPON-8x and GE-24x service access cards. Service access cards can be mixed in any of the 20 universal line cards slot within the E7-20 chassis without restriction.

Operation Domains Within the E7-20

The following diagram shows the logical partitioning of the E7-20 cards into two distinct domains:

- **The SCP Common Control Cards** participate in the domain that supports the following:
 - Uplinks to the network
 - Connects to transport rings
 - Aggregates other Calix access nodes (for example, a C7 providing TDM voice gateway services)
- **The GPON and GE-24x Universal Line Cards** only participate in the domain that supports the following:
 - Allows subscribers to access services



Common control cards and the Network Domain

The SCP-10GE cards provide the centralized processing and switching functions for the entire E7-20 Chassis. Because of the integrated nature of the SCP, it is the core component of both the Management and Networking operations. These cards also have integrated, network facing transport / uplink ports. The two common control slots in an E7-20 chassis are located between universal line card slots 10 and 11. The E7-20 chassis is designed with a dual-star backplane topology, where each common control card slots has a direct 100 Gbps high-speed backplane connection every universal line card slot, as shown in the diagram above. Each SCP-10GE provides 10Gbps full duplex connections to each line card slot, with one SCP-10GE in active mode and the second SCP-10GE in stand-by mode.

In the E7-20 system, either SCP-10GE (in a dual-card configuration) can be the active system controller. The active system controller manages the fan tray assembly (FTA) and all management interfaces within the system. The first SCP installed and booted up in the system becomes the system controller unless the system has been otherwise configured by the service provider.

By default, both SCP cards are configured for Controller candidate = YES. This setting indicates there is no preference in the system for which SCP becomes the system controller. If an operator configures one of the SCP cards to Controller Candidate = NO, that card does the following:

- Pass system control to the other SCP (revertive) in the event of a power cycle or replacement of the original system controller
- Block user forced controller switches and software upgrades that require control to transfer to the non-preferred SCP

To preserve services and maintain connectivity to the system, the E7-20 system software will override the controller candidate selection process in the event of a failure in the designated SCP.

The E7-20 SCP supports a flexible set of network topology protocols for use in aggregation, ring-based transport, and uplink applications:

- ITU G.8032 Ethernet Ring Protection Switching (ERPS)
- IEEE 802.1w Rapid Spanning Tree Protocol (RSTP)
- IEEE 802.3ad/802.1AX Link Aggregation (LAG)
- IEEE 802.11 VLAN tags
- IEEE 802.1AD Q-in-Q VLAN tags
- VLAN and MAC Switching
- Tag Actions
- Traffic management and QoS
- Ingress rate limiting
- Egress Queuing and scheduling

Network Domain VLAN and MAC switching

The following VLAN Provider Bridge capability is supported:

- IEEE 802.11 VLAN tags
- IEEE 802.1AD Q-in-Q VLAN tags
- A mix of single tagged, double tagged (Q-in-Q), and “hidden tags” on the same trunk. Hidden tags may be used by a business subscriber and are transparent to the Q-in-Q VLAN tags used by a service provider to deliver end-to-end service.
- VLAN IDs 1-4093 are supported for VLAN+MAC switching. The E7 reserves 4 VLAN values for internal use. The default value for these VLANs is 1002, 1003, 1004 and 1005, but they can be changed to some other set of 4 consecutive VLANs if required.
- 32k MAC addresses system capacity
- No restrictions on VLAN ID for any service
- Up to 9000 bytes MTUs are supported on the E7-20 10GE and GE ports
- Default native VLAN ID is 1

Network Domain traffic Management and QoS

The general goal of traffic management within the network domain is to preserve the individual service bandwidth grants established in the Access Domain. As such, the E7-20 supports network ingress rate limiting to protect the E7-20 services in the downstream direction and network egress aggregate service COS rate limiting to preserve the committed information rate (CIR) bandwidth for each COS.

- **Ingress rate limiting:** Ingress bandwidth can be limited on a per flow basis, an aggregate basis, or as an “all in one” service. Up to 1500 combinations of S+C+P-bit can be identified by a class map and policed independently.
- **Egress Queuing and scheduling:** The E7-20 supports eight Class of Service (COS) priority queues mapped to 802.1p values, where p-bit 0 is the lowest priority and p-bit 7 is the highest priority. Queue congestion is managed using a simple tail drop mechanism. Egress traffic can be scheduled and shaped in multiple methods:
 - **Port level** – aggregate bandwidth rate of the port is limited to a service provider defined value
 - **CoS level** – all egress packets are placed into one of the eight queues and rate limited as a service class. Each service class queue can have a minimum and maximum rate, thus providing fairness for all traffic classes by ensuring each has a minimum bandwidth allocation.
 - **Port+CoS level** – egress ports can use port and CoS egress bandwidth limiters together

Universal line cards and the Access Domain

In general, access services are deployed from the E7-20 access domain, with the universal line card slots delivering subscriber-facing services. As such, the access domain is designed to be managed through the network domain, which securely connects to a service provider’s back-office systems to enable flow-through provisioning, automate provisioning through the Calix Management System (CMS), or direct provisioning (CLI or web GUI) through the local MGT ports on the E7-20 chassis. In the networking domain, services are managed in the aggregate. In the the access domain, services are managed as individual layer 2 flows. At the edge of the network, individual services are classified and rate limited, and may receive additional layer 3 service processing. (for example, DHCP Relay Agent information insertion, subscriber service isolation, ARP proxy, etc.)

The primary function of the E7-20 chassis is to provide the aggregation of high density fiber services, which is accomplished through 20 universal line card slots. The GPON-4x and GE-24x service access card functions are as follows:

- Provide the fiber interfaces to the customer premises demarcation points
- Operate as the untrusted subscriber service delivery access domain
- Support 10Gbps full duplex connections to each SCP-10GE

Interface roles and networking topology protocols

The E7-20 SCP supports a flexible set of network topology protocols for use in aggregation, ring-based transport, and uplink applications. Each GE and 10GE interface in the E7 has one of the following configuration-role types:

Trunk: A port connecting to other equipment belonging to the service provider or to another service domain with consistent VLAN tagging levels. These ports may also be referred to as Network ports or Provider ports in industry standards. These ports support outer VLAN tag plus MAC switching.

Examples of trunk ports are 10G ERPS transport ports and GE uplinks. Trunk ports can be configured for link aggregation, RSTP, or ERPS.

Edge: A port facing customer equipment or facing reduced functionality devices, alternative administrative domains, or managed CPE. Generally, this E7 Ethernet port interface is where all classification is first performed on ingress traffic (if customer facing). The E7 Ethernet interface is also expected to add, replace, or remove one or more VLAN tags on edge traffic. RSTP, and LAG networking protocols are supported.

Examples of an edge port would include GE ports to managed CPE, a GE/10GE port to external equipment which may use different tagging levels, or GPON ports.

Access: A port facing untrusted customer equipment or other devices serving subscribers. Generally, this port interface is where individual subscriber services are defined and enforced (bandwidth limits, security, multicast profiles). Networking protocols are not supported.

Examples of an access port would be point-to-point connections to subscribers or other devices serving subscribers.

Interface role configuration guidelines

	Trunk	Edge	Access
E7-20 SCP	X	X	
E7-20 GE-24			X
ONT Ethernet Port			X
Tag Actions		X	X
Native VLAN	X	X	
Networking Protocol (RSTP, ERPS*, LAG)	X	X	

*Edge ports only support RSTP and LAG networking protocols, NOT ERPS.

ITU G.8032 Ethernet Ring Protection Switching (ERPS)

Two independent ERPS domains are supported on the SCP card and can be associated with either 10GE or GE ports. All ports in a common ERPS domain must be of the same speed. ERPS is used predominantly to create aggregation and transport rings between Calix Unified Access platforms, with the ERPS rings supporting a mixture of different EXA Powered systems (for example, C7, E7-20, E7-2, etc.).

IEEE 802.1w Rapid Spanning Tree Protocol (RSTP)

A single RSTP bridge domain spans all SCP 10GE and GE interfaces. RSTP has been around for many years and has achieved a very high degree of interoperability across a wide variety of products. RSTP is often used to create redundant links from the access network to an edge router or to connect subtended network elements from multiple vendors. RSTP is enabled by default on all SCP interfaces and will block a port if a forwarding loop is created inadvertently.

IEEE 802.3ad/802.1AX Link Aggregation (LAG)

The SCP supports LACP-based LAG on all SCP ports. The following configuration guidelines apply to LAG on the E7-20:

- LACP-based link aggregation protection (Active / Standby) is supported across two SCP cards
- LAG on the same SCP card can be configured for protection (Active / Standby) or bandwidth load sharing (Active / Active)
- Up to 8x GE ports are allowed in an individual link aggregation group. Across two SCP cards, the maximum supported configuration is 4:4 LAG protection
- Up to 4x 10GE ports are allowed in an individual link aggregation group. Across two SCP cards, the maximum supported configuration is 2:2 LAG protection

RSTP can be used in conjunction with LAG interfaces on two SCP cards to create RSTP-protected NxGE interfaces to an edge router or other subtended device. This is similar to N:N cross-card LAG protection, but has the further advantage of supporting ring topologies.

VLAN support

By default, every E7 Ethernet port with a trunk or edge interface on the unit is a member of VLAN 1, the Native VLAN. The Native VLAN is available to pass any untagged traffic. You can provision an E7 Ethernet port interface to use a different existing VLAN as the Native VLAN.

To forward untagged traffic on E7 Ethernet ports with an access interface, an add-tag action must be applied to untagged frames, assigning the traffic to a designated VLAN.

Tagged traffic that does not match any of the tagging criteria is dropped.

Service Intelligence Edge

ONT-based services and the service intelligence edge

The high capacity E7-20 ESAP is designed for deployed in the Central Office, with the ONT providing the GPON or AE termination point at the untrusted edge of the access domain. Calix GPON and AE ONTs, as well as other access nodes like the Calix E3, all provide service intelligence for subscriber services. The E7-20 plays an aggregation and transport role in the delivery of services when these devices are used since they normally act as the service intelligence edge. Calix ONTs acting as the service intelligence edge allow the E7-20 to deliver support for features like DHCP Relay Agent (Option 82) circuit ID insertion on the ONT's management and VoIP VLANs.

Calix GPON and AE ONTs provide a common set of services as summarized below:

- VLAN and MAC switching
 - TR-101 1:1 per-subscriber and N:1 per-service VLAN models
 - IEEE 802.11 VLAN tags; IEEE 802.1AD Q-in-Q VLAN tags
 - Service provider VLAN tag range of 1-4093
 - Customer VLAN tag range of 1-4094
 - No restrictions on VLAN ID for any service
 - 9000 byte MTU max on the E7-20; MTU size may be limited by the ONT or node
- Tag Actions (predominantly occur on the ONT)
 - Match criteria: tagged, untagged, priority tagged, MAC OUI
 - Add-tag, change-tag, add-Add, add-change
 - Native VLAN is not supported (all untagged traffic must have an explicit tag action)
- Bandwidth profiles defined with committed information rate (CIR) and peak information rate (PIR)
- Traffic engineering and QoS
 - Bi-directional, per-service rate limiting
- Networking protocol support
 - RSTP, LACP, and ERPS are not supported interfaces on the ONTs
 - Layer 2 Control Protocol (L2CP) filtering is supported
- Subscriber isolation and denial of service (DOS) protection
 - DHCP snooping and agent relay (option 82) on subscriber services
 - MAC forced forwarding and IP source verification
 - Broadcast and multicast traffic rate limiters
 - DHCP servers and multicast sources blocked

- Other services
 - Video – IGMP proxy
 - Video – Multicast VLAN registration (MVR)
 - Voice services as supported by the ONT
 - T1 services as supported by the ONT

Note: Refer to the *Calix GPON Applications Guide* and *Calix AE Applications Guide* for additional information.

Access domain and the service intelligence edge

In some applications, the device providing the physical service provider edge is not also the service intelligence edge. Two examples in point-to-point network include:

- Simple “media converter” ONUs that act as the WAN-to-LAN interface. All traffic that comes in the WAN port is switched to one or more LAN ports without any higher level service intelligence.
- Business service Ethernet network interface devices (“E-NID”) may be used to provide advanced performance monitoring for the service provider, but they do not otherwise provide rate limiting or service intelligence.

Given the inherent intelligence needed to deploy point-to-multipoint services, and the requirements established by FSAN and the Broadband Forum, GPON ONTs have generally not followed the simplified, “no or limited intelligence” deployment models. (i.e. no rate limiting, no IGMP, etc.)



To facilitate the application of service intelligence, the E7-20 provides native access domain support for point-to-point Ethernet interfaces on the GE-24x service access card. Access domain interfaces are not supported on the E7-20 SCP 10GE or GE ports, which are used exclusively for the network domain. The E7-2 ESAP also has the ability to natively support access domain service intelligence for point-to-point GE ports. This allows the E7-20 to provide service intelligence for media converters, AE ONTs, and other limited functionality Ethernet access node devices.

E7-20 Chassis

Cabinet deployments are not supported unless they are cooled to a typical Central Office environment, using air conditioning, not a heat exchanger.

Rack mounting options

The E7-20 provides two mounting options:

- Standard US Telco with 5-inch maximum protrusion from the front of the rack mounting rails to farthest forward point of the equipment, all cables included.
- Mid-mount with 10-inch maximum protrusion from the front of the rack mounting rails to farthest forward point of the equipment, all cables included.

The chassis is shipped with mounting brackets attached that allow for installation in a 19-inch rack using the standard US Telco mounting option, allowing for a 5-inch maximum protrusion from the front of the rack mounting rails. The installation kit includes an additional set of mounting ears, to adapt the already attached brackets for the chassis installation in a 23-inch rack. The two sheet metal tabs on each mounting bracket are not used.

E7-20 chassis dimensions

The Calix E7-20 chassis exterior dimensions follow:

- Height: 22.63 inches (57.49 cm)
- Width: 19.00 inches (48.26 cm) with the chassis flanges (default)
- Depth: 15.60 inches (39.62 cm) with the fibers and comb

Power and ground connections

The E7 supports redundant DC power supply via three power input zones located on the back of the unit that can each be connected to A and B power sources for automatic switchover. Equipment ground is also located on the back of the unit.

The E7 supports an input power range of -42.5 VDC to -72 VDC. The following diagram shows the location of the power and ground connections on the rear panel.

E7-20 Ethernet management ports

The Calix E7-20 is equipped with front and rear panel out-of-band 10/100 Ethernet management ports (RJ-45 connectors).

- The front panel Ethernet management port (**MGT-1**) is for temporary connection to the E7-20 for initial turnup and troubleshooting and is located on the E7-20 fan tray assembly.

- The rear panel shared Ethernet management ports (**MGT-3A, MGT-3B**) are for permanent connection to the OSS, providing management access to the E7-20. One of the following methods is used with these ports, depending on the service provider's intraoffice network design:
 - **Single Ethernet management port operation** is the most common configuration, where a single physical management interface is connected to the service provider's back office network. This interface is active regardless of which SCP card is active.
 - **Redundant Ethernet management port operation** is common with larger carriers, where two physical management interfaces are connected to redundant back office networks (typically two different Ethernet switches). The network element has a single IP address and only one interface is active at a time. **Note:** This function will be supported in a future release.

Note: Use a standard 'straight-through' Ethernet patch cable to connect to the E7-20 Ethernet management ports.

Note: When you connect a disabled MGT port on the E7 to an Ethernet switch, the link status lights on the switch light up (yellow, and then green after 30 seconds). A link is established between the E7 and the switch, but the Management interface on the E7 is not enabled and no traffic will pass over the link.

Note: The E7-20 rear panel management ports (MGT-3A, MGT-3B) act independently and must be configured on unique subnets.

To connect to the front Ethernet management port

1. Get a 'straight-through' Ethernet patch cable with RJ-45 connectors on both ends.
2. Connect the cable to the E7-20 front Ethernet management port (labeled **MGT-1**, located on the E7-20 fan tray assembly).




3. Connect the cable's other end to the Ethernet port on your PC.

Note: The front Ethernet management port is intended for direct local management access, and should remain available for this purpose. If you require permanent out-of-band management connectivity to the E7-20, Calix recommends using the rear Ethernet management port.

To connect to the rear Ethernet management ports

1. Get two Ethernet patch cables with RJ-45 connectors on both ends.
2. Connect the cables to the E7-20 rear Ethernet management ports (labeled **MGT-3A** and **MGT-3B**, located on the E7-20 rear panel).



3. Connect the cables' other end to a LAN Ethernet hub or switch.

E7-20 Fan Tray Assembly (FTA)

The E7-20 FTA resilient design includes 8 individual variable speed fans that maintain system cooling even with one fan failure. Airflow is from the bottom of the chassis (toward the line cards).

Note: Two versions of the Fan Tray Assembly are available, FTA and FTA2.

The FTA2 module is required for operation with SCP2 cards, and is backward compatible with the FTA module. The FTA2 only accepts -48 VDC, while the FTA accepts -48 or -60 VDC. Before replacing an FTA with FTA2 module, confirm that the input power supply is -48V, not -60V.



Calix E7-20 Fan Tray Assembly Front Panel

Each FTA is hot-swappable and can be quickly replaced by unlocking the latches, removing the failed unit and sliding in the new FTA. Each FTA includes a field replaceable fan filter that slides underneath the fan assembly.

Note: Do not leave a powered E7-20 operating for more than one minute without a fan module installed, to avoid overheating the system.



System management ports

The E7 web and command line interfaces, along with Calix Management System (CMS) allow you to manage services while understanding their relationship to the network infrastructure.

E7-20 LED Behavior

System status indicators

The Calix E7-20 fan module has three LEDs to indicate system alarm and operational status.

LED	Name	Color	Status	Description
CR	Critical	Red	On	A critical alarm is present in the system
			Off	No critical alarms are present in the system
MJ	Major	Red	On	A major alarm is present in the system
			Off	No major alarms are present in the system
MN	Minor	Amber	On	A minor alarm is present in the system
			Off	No minor alarms are present in the system

Card status indicators

Each Calix E7-20 line card has three LEDs to indicate the card's operational status.

LED	Name	Color	Status	Description
FAIL	Failure	Red	On	A fault has occurred that should be addressed
			Off	Normal card operation
CTRL	Control	Green	On	This card is the active E7-20 shelf controller
		Amber	On	This card is the standby E7-20 shelf controller
SRVC	Service	Green	On	One or more ports are enabled and may carry services
			Off	No ports on this card are enabled

Note: LED behaviors are shown for the card's primary operational states. The card LEDs exhibit additional behaviors associated with boot up cycles and other activities. Refer to the *Calix E7-20 User Guide* for full details.

Port status indicators

Each interface port on an E7-20 line card has an LED located beside its module socket to indicate port status. (E7-20 cards contain a mix of some or all port types listed below)

LED	Port	Color	Status	Description
XFP (1, 2)	10GE	Green	On	Indicates that an Ethernet link has been established
			Blinking	After module insertion, blinks (3) times to indicate the inserted module is recognized and allowed to operate Blinks at variable speeds to indicate the link is transmitting/receiving traffic
SFP+ (1, 2)			Off	Socket is vacant or an invalid module is inserted
		Red	On	Indicates a module has been inserted into a 10GE socket that is directed to the backplane for connectivity between two cards (Applies to XFP 2 and SFP+ 2 only)
SFP (1–12)	GE	Green	On	Indicates that an Ethernet link has been established
			Blinking	After module insertion, blinks (3) times to indicate the inserted module is recognized and allowed to operate Blinks at variable speeds to indicate the link is transmitting/receiving traffic
			Off	Socket is vacant or an invalid module is inserted
GPON (1–4)	GPON	Green	On	Indicates that at least one ONT is in service on the PON
			Blinking	After module insertion, blinks (3) times to indicate the inserted module is recognized and allowed to operate Blinks steadily while the first ONT on the port is ranging
			Off	Socket is vacant or an invalid module is inserted

E3-8G Node-Based OLT Description

Description

The E3-8G Node-Based Optical Line Terminal (OLT) is optimized to allow cable operators to cost-effectively deliver symmetrical gigabit services today, while preparing for 10 gigabits per second (Gbps) and multi-wavelength based PON service offerings in the future. Leveraging 60/90VAC cable plant power and a strand or pole mountable node based design, the E3-8G seamlessly delivers advanced broadband services to subscribers using proven Gigabit Passive Optical Network (GPON) technology.

Similar to HFC optical nodes, the E3-8G can be installed in a wide variety of outside plant deployment conditions: strand mount, attached to a pole, within a pedestal, or on a wall. Constructed with a ruggedized aluminum housing, the E3-8G can withstand extreme environmental conditions without dependency on external cooling or shelter. The “lid” of the E3-8G contains the GPON and transport interfaces as well as accommodates the fiber management for pre-connectorized “node” cable. The “tub” of the E3-8G contains the power supply and additional electronic components associated with the unit.

For additional product details, see the *Calix E3-8G Node-Based OLT Product Datasheet*, available from the My Calix Documentation Library.



Dimensions

Width: 9 in (22.9 cm)

Height: 9 in (22.9 cm)


Depth: 19 in (48.3 cm)

Access compartment connections and cable entry locations

Cable plant power (from coax) - upper right plug or upper left plug


Note: From the factory, the left-side coax connector is internally wired to the power supply. This may easily be changed to the right side.

Outdoor optical fiber cable assembly - lower right plug



Port locations (interior)

See the following diagram for the location of the GPON ports, 10GE ports, and management ports.




Note: On the actual circuit board, the GPON ports are labeled 0—7, and the 10GE (XFP) ports are labeled 0—1. However, the numbering shown in the diagram is consistent with the user interface.

Management ports

The E3-8G provides the following two (2) management ports:

- Out-of-band 10/100 Ethernet management port (RJ-45 connector).
- RS-232 serial port for console connections to the E3-8G CLI (RJ-11 connector). The pin-out of this RJ-11 connector is identical to the same port on the E7-2 (E7-2 MGT-4 port).



E3-8G Accessories

Powering Option

The E3-8G is designed to be powered via a coaxial 60V/90VAC input. For customers without this available power network, we have worked with Alpha to provide the initial coax powering support. The XM2-300HP is a high efficiency 300 Watt, 63/87VAC power supply that features a small footprint and is engineered for use in multiple dwelling units (MDU) and business parks as well as deep fiber deployments. The XM2-300HP (017-877-20) can be ordered with a hardened OSP LPE enclosure (0290021-004) that is designed to protect the main power supply from the harsh outdoor environments and also provide a standby battery backup (1810078). The LPE enclosure can be pad mounted but also comes with a bracket designed to support pole, or H-Frame installations where a power meter has been located. The XM2-300HP can also be used stand-alone in an indoor environment, such as a lab or office, when separated from the LPE enclosure.

Please note that this initial Alpha power offering does not support an environmental alarm output compatible with the E3-8G and therefore will not provide power related alarm indications. This is being reviewed and once a solution is identified, the PPG will be update to reflect the change.

The following manufacturer part numbers are readily available from Alpha and should be ordered direct from Alpha. Please send all orders to Alpha Technologies Ltd. by emailing them at orderpo@alpha.ca or by fax at (604) 638-8698. Also, please reference quote number 31504-0002 on the purchase order (P.O.) to receive the agreed to pass through pricing for the product.

The Alpha part numbers are as follows:

017-877-20	XM2-300-HP,120Vac,60Hz, 12VDC HP w/ dual RTS
0290021-004	LPE Enclosure, Mounting Bracket, BBX, SPI, BCK, LA-P
1810078	Standby Battery Backup



Alpha XM2-300-HP Power Supply



Alpha Enclosure with Power Supply and Battery

Fiber Patch Cables

Fiber patch cables may be provided by the network operator, or obtained from Calix. Calix offers three fiber patch cables in varying lengths of 10, 30 and 100 feet. Each cable contains nine fibers with connectors for use with the GPON OIMs, and five fibers with connectors for use with the 10GE XFP modules. The 10 and 30 foot cables have connectors on one end, whereas the 100 foot cable provides connectors on both ends.

The Calix part numbers are as follows:

100-04597	Fiber patch; 9 x for GPON OIMs, 5 x for 10GE XFP modules,10ft, Pigtail, OSP rated
100-04605	Fiber patch; 9 x GPON OIMs, 5 x for 10GE XFP modules,30ft, Pigtail, OSP rated
100-04606	Fiber patch; 9 x GPON OIMs, 5 x for 10GE XFP modules,100ft, Jumper, OSP rated

E3-48C Platform Description

The Calix E3-48C is a compact, hardened, high capacity sealed Ethernet Service Access Node (ESAN) designed for outside plant applications. The E3-48C is designed for rapid deployment of advanced copper-based broadband services and optimized to meet the needs of higher subscriber density. The E3-48C may be AC, DC, or Line powered and supports 48 VDSL2/ADSL2+ ports as well as 48 analog POTs ports.

The Calix E3-48C, E3-48 r2, E3-48C r2 are compact, hardened, high capacity sealed Ethernet Service Access Node (ESAN) designed for outside plant applications. The E3-48C is designed for rapid deployment of advanced copper-based broadband services and optimized to meet the needs of higher subscriber density. The E3-48C may be AC, DC, or Line powered and supports 48 VDSL2/ADSL2+ ports as well as 48 analog POTs ports.

E3-48/C Features

Below is a summary of the key features and benefits of the E3-48/C:

- 48-port VDSL2/ADSL2+ ports with 2-port bonding for high bandwidth needs
- 48 analog POTs ports converted to VoIP for transport on the upstream fiber
- Unit-level vectoring over VDSL2 to maximize the copper bandwidth
- High-capacity network interfaces of 3 x 1/2.5GbE and 2 x 1/10GbE for flexible uplink and downlink configurations
- E3-48 r2, E3-48C r2 units support up to 8 ports of bonded VDSL2 uplink where fiber is not available. The VDSL2 bonded copper uplinks connect to E7 VDSL r2 and E3-48 r2 products.
- E3-48 r2, E3-48C r2 units have 1 vector control interface port to support System Level Vectoring and connection to an E3 VCP
- Interoperability with E-Series and C-Series using ERPS Ethernet protection for main ring or subtended ring topologies
- Daisy chain or RSTP ring (max of 8 nodes) configurations to meet various network topology needs
- Sealed design that meets the most stringent industrial standard on outdoor deployment environments



Calix E3-48/C ESAN

The E3-48/C is designed for fiber to the node installations delivering voice, high speed data, and video services. The E3-48/C unit supports the same voice capabilities of the E7 VDSL2 platform supporting SIP, TDM and H.248 Gateway services.

The E3-48/C supports advanced VDSL2 as well as ADSL2+ fallback technologies with a non-blocking architecture via 3 x 1GE / 2.5GE (SFP) and 2 x 1GE / 10GE (SFP+) ports. The E3-48/C supports one permanent virtual circuit (PVC) or one service mode in ADSL2+ fallback mode. A future software release will add multiple PVC support in ADSL2+ fallback mode. The E3-48/C aligns closely with the performance and configuration characteristics of the E7-2 VDSL2 line card. Supporting 2-port ATM and PTM bonding as well as the full suite of E7 DSL capabilities the unit may be integrated smoothly into network deployments. Additionally, the E3-48/C is capable of supporting VDSL2 vectoring in Release 2.4.

The E3-48/C can be used to deliver high bandwidth residential and business services from a Remote Terminal (RT), fiber-to-the-node (FTTN), fiber-to-the curb (FTTC) or MDU location. It supports environmental flexibility with its sealed enclosure, which has been specially designed for deployment in extreme climate conditions. The E3-48C has been designed to easily fit within existing rights-of-way (ROW) using AC, DC, or Remote power. This approach ensures cost-effective, low-density service deployments from the CO to the access network edge.

Centralized network element management for all E-Series products, including the E3-48C, is supported with the Calix Management System (CMS) release 12.1 and higher. The following management interfaces are also available:

- Command line interface (CLI)
- Embedded web graphical user interface (GUI)
- Simple network management protocol (SNMP v2c)
- Calix Management System (CMS)

E3-48/C Physical Attributes

The Calix E3-48/C has the following physical attributes:

- **Network Interfaces:**
 - 3 x 1GE / 2.5GE small form-factor pluggable (SFP) sockets for 1GE or 2.5GE uplinks and/or subtending downlinks. The units can be daisy chained or ringed with other products supporting standards based RSTP as well as the Calix E-Series and C-Series products. The E3-48C also supports ERPS which seamlessly interconnects with deployed Calix E-Series and C-Series ERPS rings. The E3-48C is simply another node in the ERPS protection domain. In addition, the E3-48C can be aggregated using point-to-point links to the Calix B6, C7, or E7 products.
 - 2 x 1GE / 10GE small form-factor pluggable (SFP+) sockets for 1GE or 10GE uplinks and/or subtending downlinks. The E3-48C supports ERPS which may be used to interconnect with Calix E-Series and C-Series products. The units may also be daisy chained or ringed with other products supporting standards based RSTP as well as the Calix E-Series and C-Series products. In addition, the E3-48C can be aggregated using point-to-point links to the Calix B6, C7, or E7 products.
 - Both the SFP+ ports support the use of the copper direct attach (DA) cables available from Calix.
 - E3-48 r2, E3-48C r2 units support up to 8 ports of bonded VDSL2 uplink where fiber is not available. The VDSL2 bonded copper uplinks connect to E7 VDSL r2 and E3-48 r2 products.
 -
- **Subscriber Interfaces:** 48 x subscriber-facing ports with VDSL2/ADSL2+ and POTS signals
- **Alarm Inputs/Output:** 3 input alarms and 1 output alarm
- **Management Interfaces:**
 - MGT-1 for local craft access to management user interfaces. The connector is 10/100 Base-T Ethernet (RJ-45).
 - MGT-2 can be configured within any VLAN on any Ethernet port interface for an in-band management interface.
 - MGT-4 for local craft management. The connector is RS-232 serial: 38.4 kbps. (Supports the Calix command line interface (CLI) only.)
 - E3-48 r2, E3-48C r2 units have 1 vector control interface port to support System Level Vectoring and connection to an E3 VCP

Power Feeds: AC powered with 110-240VAC, -48vDC, Remote Power +/- 190vDC


Operating Environment: The E3-48/C is environmentally sealed and hardened, supporting an operational temperature range of -40° C to +70° C with humidity of 5-95% (non-condensing).

- **Dimensions:** 17" x 23" x 6" inches



E3-48/C Dimensions

- **Physical Enclosure:** The E3-48C, E3-48 r2, E3-48C r2 are sealed unit enclosures that easily integrate into outside plant deployments. The unit consists of a ruggedized housing with an access compartment for line protection and outside plant cable connections. The access compartment is designed for copper connections for the POTS / DSL output, power, ground, interconnection, and trunk interfaces. This area also includes the 5-pin 50-pair line protection block. Cable access to this area is routed through the bottom plate shown in the diagram below.




E3-48/C Bottom Plate View

The E3-48/C enclosure has one PG36 opening for the subscriber cable, and five PG13 openings for other cables, including fiber patch, power and ground, and E3-48/C Direct Attach (DA) cables. This unit may be deployed in point to point, star, ring, or hub network topology deployments. The E3-48/C unit may serve 3 additional E3-48/C units in nodal configurations by routing the Direct Attach cables through the PG13 openings. The PG form factor is a specific standard. English sized components and connectors are not recommended for use with the PG form factor.

E3-48 r2 and E3-48C r2 Platform Description

Product Overview

The Calix E3-48 r2 products are a family of sealed Ethernet Service Access Node (ESAN) products providing VDSL2 / POTS Combo capabilities and VDSL2 Overlay applications. The E3-48 r2 products introduce new capabilities enabling copper bonded VDSL2 trunk uplinks in addition to fiber uplinks supported today, expanded vectoring capabilities, and higher single pair and bonded pair performance. The E3-48 r2 products are capable of working with an E3 Vector Control Processor (VCP) to optimize subscriber bandwidth by vectoring up to 192 lines. The E3-48 r2 models and the E3-VCP are shown in the diagram below.



Product Description

The Calix E3-48 r2 products have the same E3-48 enclosure, which is optimized for harsh environment OSP installations. The E3-48 products are a compact, hardened, high capacity sealed Ethernet Service Access Node (ESAN) designed for outside plant applications. The E3-48 r2 enables rapid deployment of advanced copper-based broadband services and optimized to meet the needs of higher subscriber density. The E3-48 r2 supports 48 VDSL2/ADSL2+ ports and may be AC, DC, or Line powered. The E3 VCP supports each of these powering options and enables up to 192 vectored VDSL2 ports. The E3 VCP connects to the E3-48r2 products through a dedicated high speed connector that handles the management of the VCP as well as the VDSL2 calculations.

Below is a summary of the key features and benefits of the E3-48 r2 products:

- 48-port VDSL2/ADSL2+ ports with up to 8-port bonding for high bandwidth needs.
- Unit Level and System Level Vectoring over VDSL2 to maximize the copper bandwidth.
- High-capacity vector control interface port for sealed connection to a sealed E3-VCP.
- High-capacity network interfaces of 3 x GE/2.5GE and 2 x GE/10GE for flexible uplink and downlink configurations.
- 8-port bonded, vectored VDSL2 for copper trunk uplinks (available on copper uplink enabled models).
- Universal powering for AC, -48vDC, and Line +/- 190vDC powering
- 8-port bonded copper trunk uplinks may be used with line powered pairs
- Interoperability with E-Series and C-Series using ERPS Ethernet protection for main ring or subtended ring topologies

- Daisy chain or RSTP ring (max of 8 nodes) configurations to meet various network topology needs
- Sealed design that meets the most stringent industrial standard on outdoor deployment environment, IP-67/GR-950.

The E3-48 r2 is designed for fiber to the node (FTTN) installations delivering voice, high speed data, and video services. The E3-48 r2 also supports FTTN extensions and loop shortening by using copper 8-pair bonded trunks to connect to a new node. Shorter loop vectored VDSL2 services can be served rapidly from a copper served, line powered node. Fiber and local powering can be added to this node at a future date. This capability allows for loop shortening and new service turn up capabilities within the existing fiber network. The E3-48 r2 unit supports the same voice capabilities of the E7 platform in Release 2.2 supporting SIP and H.248.

The E3-48 r2 supports advanced VDSL2 as well as ADSL2+ fallback technologies with a non-blocking architecture via 3 x GE/2.5GE (SFP) and 2 x GE/10GE (SFP+) ports. The E3-48 r2 aligns closely with the performance and configuration characteristics of the E7-2 VDSL2 r2 line cards. The E3-48 r2 products support standard 2-port ATM and PTM bonding as well as 8-port VDSL2 PTM bonding. 8-port VDSL2 PTM bonding requires specific multi-pair bonded CPE from Positron or Actellis (at the time of this release). Other multi-pair bonded may be added in the future. The full suite of E-Series DSL capabilities allows the E3-48 r2 products to be integrated smoothly into network deployments.

The E3-48 r2 can be used to deliver high bandwidth residential and business services from a Remote Terminal (RT), fiber-to-the-node (FTTN), fiber-to-the curb (FTTC) or MDU locations. It supports environmental flexibility with its sealed enclosure, which has been specially designed for deployment in extreme climate conditions. The E3-48 r2 has been designed to easily fit within existing rights-of-way (ROW) using AC, DC, or remote power. This approach ensures cost-effective, low-density service deployments from the CO to the access network edge.

Centralized network element management for all E-Series products, including the E3-48 r2, is supported with the Calix Management System (CMS) release 14.1. The following management interfaces are also available:

- Command line interface (CLI)
- Embedded web graphical user interface (GUI)
- Simple network management protocol (SNMP v2c)

Calix Management System (CMS)

E3-48 r2 and E3-48C r2 Chassis Attributes

The Calix E3-48 r2 has the following physical attributes:


- **Network Interfaces**
 - 3 x 1GE / 2.5GE small form-factor pluggable (SFP) sockets for 1GE or 2.5GE uplinks and/or subtending downlinks. The units can be daisy chained or ringed with other products supporting standards based RSTP as well as the Calix E-Series and C-Series products. The E3-48 r2 also supports ERPS and G.8032v2 which seamlessly interconnects with deployed Calix E-Series and C-Series ERPS and G.8032v2 rings.
 - 2 x 1GE / 10GE small form-factor pluggable (SFP+) sockets for 1GE or 10GE uplinks and/or subtending downlinks. The E3-48 r2 supports ERPS and G.8032v2 which may be used to interconnect with Calix E-Series and C-Series products. The units may also be daisy chained or ringed with other products supporting standards based RSTP as well as the Calix E-Series and C-Series products.
 - The SFP and SFP+ ports support the use of the copper direct attach (DA) cables available from Calix. Direct Attach cables are used to interconnect multiple E3-48 r2 units for node locations greater than 48 subscribers. The direct attach cable includes the gland fittings to seal the cable connections to the E3-48 r2 units.
 - The E3-48 r2 and E3-48Cr2 units support up to 8 ports of bonded VDSL2 uplink where fiber is not available. The VDSL2 bonded copper uplinks connect to E7 VDSL2 r2 and E3-48 r2 products.
- **Subscriber Interfaces**
 - 48 x subscriber-facing ports with VDSL2/ADSL2+ and POTS signals
- **Vector Control Interface**
 - 1 vector control interface port to support System Level Vectoring and connection to an E3 VCP
- **Alarm Inputs/Output:** 3 input alarms and 1 output alarm
- **Power Feeds:** AC powered with 110-240VAC, -48vDC, Remote Power +/- 190vDC
- **Operating Environment:** The E3-48 r2 is environmentally sealed and hardened, supporting an operational temperature range of -40 °C to +70 °C with humidity of 5-95% (non-condensing).

E3-48 r2 chassis dimensions


The Calix E3-48 r2 chassis exterior dimensions follow:

- Height: 23.2 inches (58.9 cm)
- Width: 17.0 inches (43.2 cm)
- Depth: 6.0 inches (15.2 cm)
- Weight: 48.0 pounds (21.77 kg)

The sealed enclosure lid can be removed for turn-up and servicing.




The E3-48 r2 enclosure has one PG36 opening for the subscriber cable, four PG13 openings for other cables, including fiber patch, power and ground, and Direct Attach (DA) cables, and one PG16 opening for the Vector Control Interface cable (Note: the PG-16 opening is in the upper right hand corner in the diagram above which is slightly bigger than the PG13 openings). This E3-48 r2 units may be deployed in point to point, star, ring, or hub network topology deployments. The E3-48 r2 unit may serve 3 additional E3-48 r2 units in nodal configurations by routing the Direct Attach cables through the PG13 openings. DA cables are recommended for locally connected E3-48 devices. The right-most DA connection hole is sized with a PG-16 grommet to support the connection of the Vector Control Interface cable. The E3-48r2 units use a QSFP interface to interconnect the E3-48r2 unit to a sealed E3-VCP. Fiber connections are not recommended due to thermal and cost considerations. Multi-mode optics require more power and generate more heat than copper SFP and Direct Attach cables. The PG form factor is a specific standard. English sized components and threaded connectors do not fit into the PG form factor holes. Calix provides a PG to conduit adapter kit that is listed in the ordering section of this guide which allows for more outside plant installation options.




Access Compartment Connections

The access compartment contains all the network ports and port connections.

E3-48C r2 Access Compartment



E3-48 r2 Access Compartment



The following table describes the E3-48 r2 access compartment features.

Interface	Description
Power Terminal	- Remote/line-power option: Power terminals for connecting up to six power pair leads to a remote power supply. Terminal pins with a "+" label are for positive 190 V power; pins with a "-" label are for negative 190 V power. Connect one 5-pin protector rated for remote power (not included with E3-48 R2/E3-48C R2 shipment) for lightning protection for each power source. - Local –48 VDC power option: A and B power terminals for connecting to a DC power supply: A "+", A "-", B "+", B "-" - Local AC-power option: AC IN Line labeled "L," Ground labeled "G", and Neutral labeled "N."
ALM (Alarm)	The DB-9 port is used for external alarm input and output. A three-pair alarm cable ships with E3-48 R2 units or can be purchased as a separate item from Calix.
MGMT-1 and MGMT-4	The MGMT-1 port is used for local management using the E3-48 R2 Web interface. Use an Ethernet cable to connect to a computer. The MGMT-4 port is used for local management using the E3-48 R2 CLI.
VCC	1 vector control interface port to support System Level Vectoring and connection to an E3 VCP.
SFP(1), SFP(2), SFP(3), and SFP+(1), SFP+(2)	Use fiber-optic cables with the SFP ports or Direct Access (DA) cables to connect to an uplink or downlink device, or for daisy chaining with other E3-48 R2 devices. - SFP ports 1, 2, and 3 support a 1-Gbps or 2.5-Gbps SFP module. - SFP+ ports 1 and 2 support a 1-Gbps or 10-Gbps SFP module.
RJ-21 Connectors	E3-48C r2: Two CAT5 RJ-21 connectors. The 25th pair of each connector is not wired. E3-48 r2: Two CAT5 RJ-21 connectors. The 25th pair of each connector is not wired.
5-pin protector sockets	Connect 5-pin protector module sockets (not included with E3-48 R2 shipment) for lightning protection for subscriber ports. E3-48C r2: 48 5-pin protector sockets E3-48 r2: 96 5-pin protector sockets
Cable entry	Cable entry holes sheath the cables, hold them in place, and provide weatherproofing. Note the following: - Calix-supplied cables include above-grade fittings. - For Calix E3-48 R2/E3-48C R2 packages, subscriber and power cables may be pre-installed. Fiber patch and any required Direct Attach cables must be field-installed. - One of the holes is slightly larger than the others (top right) to accommodate the VCP connection.

Related topics

- RJ-21 Pin Assignments


Cable Entry Locations

The following diagram illustrates the recommended cable entry locations.

The Calix E3-48 r2 chassis six openings:

- One PG 36 opening (for subscriber cables).
- Four PG 13 openings (for other cables, including fiber patch, power, and Direct Attach).
- One PG 16 opening in the upper rightmost corner (for the VCC connector only).

E3-48 r2 Chassis, Bottom View



E5-48 and E5-48C Platform Description

The Calix E5-48C (combo – integrated voice and data) and E5-48 (overlay – data with external voice) are hardened, fixed form factor products that offer high bandwidth, advanced broadband services over existing copper outside plant infrastructure in a 1RU form factor.

- The E5-48C Combo chassis features 48 VDSL2/ ADSL2+ Fallback ports with integrated analog POTS ports to support high-bandwidth DSL and VoIP based POTS services.
- The E5-48 Overlay chassis features 48 VDSL2/ ADSL2+ Fallback ports with integrated splitters to support high-bandwidth DSL and analog POTS services.

Similar to the E7-2 chassis, the E5-48 and E5-48C Ethernet Service Access Nodes (ESANs) have physical outside plant connectors that are routed out of the rear of the E5-48 chassis. The E5-48/E5-48C user interface and configuration parameters align with the provisioning, performance, and configuration characteristics of the E7-2 VDSL2 line card and the E3-48C.

E5-48C Combo Unit Overview

The Calix E5-48C 48 Port VDSL2 POTS Combo Unit combines forty-eight VDSL2 / ADSL2+ and POTS subscriber ports, with four Gigabit Ethernet SFP ports, to provide high speed copper services with integrated Ethernet transport. The Calix E5-48C 48 Port VDSL2 POTS Combo Unit is a fixed form factor product that fits into a standard 19” rack to create a compact, high density DSL node, ideal for copper based delivery of IP services across the access network. The Calix E5-48C 48 Port VDSL2 POTS Combo Unit supports a full set of subscriber services and network topology protocols, and may be integrated into a number of topology configurations. The E5-48C utilizes the same VDSL2 and POTS feature sets as the E7-2 VDSL2 and POTS products. The E5-48C supports both SIP and H.248 soft switches and supports Class 5 switch interface through the C7 Trunking Gateway capabilities.



Calix E5-48C Node Front View

Connections to the outside plant are made through two 50 pair RJ-21 connectors located on the rear of the chassis labeled as P1 and P2. The E5-48 Combo unit is powered through connecting the A-B -48vDC power feeds after connecting the dual lug chassis ground. The E5-48 units provide up to 7 alarm inputs and one alarm output on each E5-48 chassis.



Calix E5-48C Node Rear View

Interface Capacity

The Calix E5-48C 48 Port VDSL2 POTS Combo unit is equipped with 48 VDSL2 / ADSL2+ Fallback ports and 48 analog POTS ports for VoIP based services. The E5-48C supports pair bonding in both ADSL2+ and VDSL2 modes. SFP ports on the E5-48C operate at 1 Gbps speed.

Industry Standards Compliance

The Calix E5-48C features 48 ports of VDSL2 combo (voice and data) and it complies with all applicable ADSL / VDSL2 ITU standards ensuring interoperability with a wide range of CPE, including:

- G.992.1 – ADSL / G.dmt
- G.992.2 - G.Lite
- ANSI T1.413
- G.992.3/4 - ADSL2
- G.992.5 - ADSL2+
- G.993.2 - VDSL2
- G.993.5 - VDSL2 Vectoring
- G.994.1 - G.hs
- G.998.2 - VDSL2 Ethernet Pair Bonding
- G.998.1 - ATM-based Multi-Pair Bonding

Voice Connectivity with H.248 and SIP

The E5-48C supports switched voice services from the E5-48C's POTS and DSL ports. A Voice Frequency (VF) splitter integrated in the product terminates the analog POTS signal, and passes it to a voice digital signal processor (DSP) for native VOIP conversion. The DSP performs the call handling and IP packetization required for VOIP implementation. The integration of VoIP on the product uses Session Initiation Protocol (SIP) as specified in RFC-3261 and ITU standard H.248. Supplemental protocol enhancements allow Calix voice capable products to support TDM voice services with VoIP packet streams with legacy Class 5 switches in conjunction with a C7 Trunking Gateway.

The E5-48C provides QOS for voice traffic throughout the Ethernet transport network and uplink interfaces. The E5-48C VDSL2-48C's VoIP traffic is tagged with an 802.1Q VLAN tag and prioritized using 802.1p priority bits. The E5-48C based voice services may be delivered in conjunction with either or both of the following:

- Calix C7 TDM switch interfaces (GR-303, TR-08 Mode I and Mode II) using a Voice over IP Resource card (VIPR) or Ethernet Gateway (EGW) card, as the VOIP to TDM gateway. The Calix C7 VOIP Services Guide provides a complete service description and provisioning instructions.
- Third party softswitches from MetaSwitch and Genband (C15, C20) using SIP VoIP services, and MetaSwitch and Genband (C20) for H.248 VoIP services.

The E-Series H.248 VoIP implementation supports the standard MGCP command structure exchanged between the Call Agent and the Media Gateway. The E-Series supports Media Gateway Controller communication using version 1 and 2 of the H.248 protocol which is the predominant protocol used with deployed Media Gateway Controllers.

The following package identifiers are supported within the command structure packages:

Package Name	Package ID (number)	H.248 Document
Generic	g (0x0001)	H.248.1
Base Root	root (0x0002)	H.248.1
Tone Generation	tonegen (0x0003)	H.248.1
Tone Detection	tonedet (0x0004)	H.248.1
DTMF Detection	dd (0x0006)	H.248.1
Call Progress Tones Generator	cg (0x0007)	H.248.1
Network	nt (0x000b)	H.248.1
RTP	rtp (0x000c)	H.248.1
TDM Circuit	tdmc (0x000d)	H.248.1
Indicator	ind (0x0019)	H.248.3
Basic Call Progress	bcg (0x0023)	Q.1950 Annex A
Expanded Call Progress Tones	xcg (0x0024)	Q.1950 Annex A
Basic Services Tones	srvt (0x0025)	Q.1950 Annex A
Enhanced Alerting	alert (0x003b)	H.248.23
Analog Display Signaling	andisp (0x003c)	H.248.23
Extended Analog Line Supervision	xal (0x0043)	H.248.26
Inactivity Timer	it (0x0045)	H.248.14
Extended DTMF Detection	xdd (0x0052)	H.248.16
Coin Operated Phone Control	coin (0x0095)	H.248.35

Lifeline Voice over IP with SIP

The E5-48C natively packetizes traditional lifeline, analog voice through an integrated SIP user agent (UA) for delivery across the IP / Ethernet network, ultimately terminating in a VOIP softswitch service environment. The SIP feature set of the E5-48C is common to the E7-2 VDSL2-48C SIP capabilities. The E5-48C supports the G.711 voice codec with packetization rates of 10ms or 20ms. Calix has tested SIP VOIP interoperability with various soft switches. Please refer to the Calix E7 Product Planning Guide for a current list.

Lifeline TDM Voice with C7 GR-303 and TR-08 Gateway

The E5-48C has the ability to have VOIP traffic routed to and terminated on a Calix C7 for conversion to TDM, for a GR-303 or TR-08 interface hand off to a traditional Class 5 switch.

POTS based VOIP Services

All 48-subscriber ports on the E5-48C provide broadband DSL access and lifeline POTS service. Analog POTS is carried over the service loop using standard Voice Frequency carriers (300 to 3400 Hertz). A line card integrated voice splitter separates the voice from the data service for each subscriber signal, handing the POTS service to a digital signal processor that digitizes and converts it to native, standards based VOIP streams.

The E5-48C supports different VOIP signaling protocols natively, the first of which used are SIP and a SIP-based Calix proprietary method of supporting TDM based GR-303 and TR-08 uplinks to a Class 5 switch deployed in conjunction with C7.

POTS interfaces of the E5-48C have the following characteristics:

- Loop-Start signaling
- Impedance of $900\ \Omega + 2.16\ \mu\text{f}$, supporting ANSI based loop plants
- Subscriber loop lengths of $1830\ \Omega$, including $430\ \Omega$ telset
- Ringing support for 15 REN / card ring source (5 REN at 1400 Ohms)
- Ringing Voltage of Balanced 86 Volts AC at 20Hz
- Loop current of up to 25mA (during normal operation), and 18mA (during battery backup)
- User configurable Gain (+2dB for A to D, -2dB for D to A)

E5-48 Overlay Unit Overview

The Calix E5-48 48 Port VDSL2 Overlay Unit combines forty-eight VDSL2 / ADSL2+ fallback subscriber ports and corresponding $600\ \Omega$ / $900\ \Omega$ integrated splitters, with four Gigabit Ethernet ports, to provide high speed copper services with integrated Ethernet transport. Integrated splitters extract analog voice from the service loop and terminate it on an existing DLC or analog POTS port on a switch line bay. The E5-48 48 Port VDSL2 Overlay Unit is a compact, high density DSL overlay node, with Ethernet aggregation and transport, providing copper based delivery of IP services and leveraging existing POTS infrastructure across the access network.

The E5-48 48 Port VDSL2 Overlay unit may be used when traditional POTS already exists and is passed through the remote location, originating from the access equipment located in the CO or an existing remote DLC. The DSL overlay is passed through an integrated splitter, allowing the delivery of voice, video, and data to a customer across a common, twisted pair cable.



Calix E5-48 Node Front View

Connections to the outside plant are made through four 50 pair RJ-21 connectors located on the rear of the chassis labeled as P1 through P4. Ports P3 and P4 are used for the POTS line input from the existing POTS infrastructure and ports P1 and P2 are used for POTS / DSL services to the subscribers. The E5-48 Overlay unit is powered through connecting the A-B - 48vDC power feeds after connecting the dual lug chassis ground. The E5-48 units provides up to 7 alarm inputs and one alarm output on each E5-48 chassis.



Calix E5-48 Overlay Node Rear View

Interface Capacity

The Calix E5-48 48 Port VDSL2 Overlay unit is equipped with 48 VDSL2/ ADSL2+ Fallback ports. In addition, the E5-48 supports 48 integrated splitter ports, to support analog POTS services. The E5-48 may be configured to support pair bonding for both ADSL2+ and VDSL2 modes.

Broadband Overlay

The E5-48 48 Port VDSL2 Overlay Unit may be used when traditional POTS already exists and is passed through the remote location, originating from the access equipment located in the CO or an existing remote DLC. The DSL overlay is passed through an integrated splitter, allowing the delivery of voice, video, and data to a customer across a common, twisted pair cable.

Overlay POTS Services with E5-48

All 48-subscriber ports on the VDSL2-48 line card provide broadband DSL access and access to overlay lifeline POTS service. Analog POTS is carried over the service loop using standard Voice Frequency carriers (300 to 3400 Hertz). A line card integrated voice splitter separates the voice from the data service for each subscriber signal, extracting the POTS service, as an analog signal, back out of the line through the POTS overlay connector ports. The analog POTS service can then be handled via a collocated Digital Loop Carrier (DLC), or sent off to a TDM switch analog front end line bay.

POTS interfaces of the VDSL2-48 line card have the following characteristics:

- Impedance of 600 Ω / 900 Ω + 2.16 μ f, supporting ANSI and ETSI based loop plants.
- Loop current of up to 25mA (during normal operation), and 18mA (during battery backup)

Power Consumption and Heat Dissipation

The power consumption and heat dissipation of the E5-48C is 92Watts and 89Watts respectively. The power consumption and heat dissipation of the E5-48 is 72 Watts and 69Watts respectively. In a typical central office (CO) environment, the E7-2 fans run at low speed and draw less power than in a remote terminal (RT) location where temperatures may peak at 65° C.

E5-48/E5-48C Chassis Attributes

The E5-48 and E5-48C chassis have the following physical attributes.

Network Interfaces

4 x 1GE small form-factor pluggable (SFP) sockets for 1GE uplinks and/or subtending downlinks. The units can be daisy chained or ringed with other products supporting standards based RSTP as well as the Calix E-Series and C-Series products. The E5-48/E5-48C also supports ERPS to interconnect with deployed Calix E-Series and C-Series ERPS rings as another node in the ERPS protection domain. In addition, the E5-48/E5-48C can be aggregated using point-to-point links to the Calix B6, C7, or E7 products. The SFP+ ports supports the use of the copper direct attach (DA) cables available from Calix.

Subscriber Interfaces

48 x subscriber-facing ports with VDSL2/ADSL2+ and POTS signals

Management Interfaces

The Calix E5-48 and E5-48C are equipped with two out-of-band 10/100 Ethernet management ports (RJ-45 connectors). The front Ethernet management port is located on the fan module (labeled **MGT-1**). The rear Ethernet management port is located on the rear panel (labeled **MGMT-3**). Use the front Ethernet port for (temporary) local Craft access to the Calix E5-48C. If you require a permanent out-of-band management connection to the Calix E5-48/E5-48C, Calix recommends using the rear Ethernet management port to connect to the LAN.

Alarm Inputs/Output

7 input alarms and 1 output alarm

Power Feeds

- -48 VDC redundant battery feeds (A and B)
- -42.5 VDC to -72 VDC input range

Operating Environment

- Temperature: -40° C to +70° C (-40° F to +158° F)
- Humidity: 10 to 95% (non-condensing)

Fan Assembly

The Calix E5-48/E5-48C chassis requires a E5 Fan Tray Assembly


Management Interfaces

- MGT-1: 10/100 Base-T Ethernet (RJ-45) located on the fan module and typically used for temporary local craft access to management user interfaces
- MGMT-3: 10/100 Base-T Ethernet (RJ-45) on rear of the chassis; typically used for permanent network connection to back office management systems
- MGT-4: RS-232 serial (RJ-11 plug) 38.4 kbps serial interface on the fan module; typically used for local craft management and only supports the Calix command line interface (CLI).
- (MGT-2 can be configured within any VLAN on any Ethernet port interface for an in-band management interface.)

Chassis Exterior Dimensions


- Height: 1.7 inches (4.3 cm)
- Width: 17.5 inches (44.5 cm)

- Depth: 11.45 inches (29.1 cm)




Fan Module Exterior Dimensions

- Height: 1.63 inches (4.1 cm)
- Width: 2.67 inches (6.8 cm)
- Depth: 10.82 inches (27.5 cm)




E7 Converged Network Solutions


With a high-availability architecture and ubiquitous 10GE support on all line cards, the Calix E7 system bridges the gap between traditional IP service access nodes and Ethernet aggregation switches. The E7 delivers the most demanding converged voice, video, and data services over a variety of optical network topologies and technologies.



High-availability Ethernet transport: Providing transport and aggregation for any Ethernet-based access network, the 10GE-4 card is the core element of the E7 Ethernet transport solution. Ring protection protocols and link aggregation are used along with the E7's native dual-card redundancy to provide unparalleled Ethernet transport redundancy. The 10GE-4 card includes four 10GE ports for use as uplink and transport for the local system; or to subtend additional chassis from the primary shelf. The GPON-4 card includes the same lineup of 10GE ports and can be used interchangeably with the 10GE-4 card.



Multiple systems can be linked together using low cost 10GE SFP+ copper cables to create a high-speed 10Gbps transport ring with aggregated, high-density GPON and point-to-point Ethernet subscriber services. Each 10GE-4 card provides twelve GE ports that can be used to provide protected aggregation links to other Calix E7, C7 and E5 platforms.




GPON residential services: The Calix E7 IP services platform is ideal for delivery of GPON and GE FTTP-based services. With the GPON-4 card, network operators can support GPON residential services and dedicated point-to-point GE service delivery in a single platform. The flexibility of the Calix E7 allows service providers to maximize revenues and customer satisfaction while reducing both operational and capital costs.



Point-to-point Ethernet services: Point-to-point Ethernet network design compliments point-to-multipoint GPON access networks. As a dedicated service, point-to-point Ethernet allows service providers to offer individual customers very high-bandwidth, dedicated symmetric services. Point-to-point Ethernet also has the greatest deployment flexibility and OSP simplicity, with the tradeoff of requiring additional fiber management in the CO and OSP.



Mobile backhaul: With integrated network synchronization, hierarchical QoS, and support for T1 services, the E7 provides uncompromised transport of mobile broadband traffic while also supporting triple-play residential and MEF certified business services from a single platform. A powerful collection of classification, policing, and scheduling algorithms let operators manage per-subscriber and per-service traffic flows to maintain priority/delay/loss differentiation within the network.




Metro Ethernet business services: The E7 can deliver Metro Ethernet Forum (MEF) compliant business services. All E7 cards support up to 4,092 Ethernet Virtual Connections (EVCs) and can be configured for E-Line (point-to-point) and E-LAN (multipoint-to-multipoint) services over both GPON and GE networks.

E3-8G Network Solutions

The E3-8G Node-Based Optical Line Terminal (OLT) is optimized to offer cable operators a solution and form factor that eases the transition from an all Hybrid Fiber-Coaxial (HFC) network to PON. The E3-8G supports a variety of deployment scenarios:

- Full HFC to FTTH conversion
- HFC DOCSIS node offload of high bandwidth, data hungry customers
- Business services overlay
- MDU focused deployments

The E3-8G delivers cost-effective symmetrical Gigabit services today, while operators prepare their network for 10 Gigabits per second (Gbps) and multi-wavelength 10 Gbps based PON service in the future. The E3-8G aligns with the provisioning, performance, and configuration characteristics of the E7-2 GPON-8-line card, eight (8) GPON ports with two (2) 10GE XFP ports for aggregation and transport, supporting the full suite of GPON features and functionality for residential and business customers.



Similar to HFC optical nodes, the E3-8G can be installed in a wide variety of outside plant deployment conditions such as strand mount, attached to a pole, within a pedestal, or on a wall. Constructed with a ruggedized aluminum housing, the E3-8G can withstand extreme environmental conditions without dependency on external cooling or shelter. The “lid” of the E3-8G contains the GPON and transport interfaces as well as accommodates the fiber management for pre-connectorized “node” cable. The “tub” of the E3-8G contains the power supply and additional electronic components associated with the unit.

The E3-8G can either leverage an operators existing cable power plant or alternatively can be deployed with a stand-alone power supply. 60/90VAC power is delivered to the E3-8G using a coaxial cable mated to a stinger connector which comes preinstalled from the factory. Stand-alone powering and battery back-up units can be provided through reference-sell equipment from Alpha Power. Leveraging 60/90VAC cable plant power and a node based design, the E3-8G delivers advanced broadband services to residential and business subscribers over GPON.

RFoG overlay using the E3-8G is accomplished in the same manner as the E7-20 and E7-2. However, as the E3-8G is a stand-alone unit, WDM couplers for the ingress of 1550nm downstream and egress of 1610nm upstream, must occur at the following:

- splitter distribution cabinet
- within a splice enclosure
- within a node based EDFA and/or return path transmitter

Special consideration must be given to the return path of the RFoG devices in an E3-8G deployment, if the distance from the headend exceeds the RFoG return loss budgets. Node based return path transmitters, such as those offered by Calix's partner Innotrans, are available to insure return path signal are managed in a manner to ensure long range reach.


E3-8G Management and Provisioning

The E3-8G is provisioned and managed by Calix's CMS system. For cable operators with DOCSIS back office systems, the E3-8G, the E7-2, and E7-20, can be managed with Calix's DOCSIS Mediation Layer (DML) – Open Link Cable. From a DOCSIS back office system, OpenLink Cable translates cable modem configuration files into actionable commands to provision services on the E7 portfolio.

E3-48C, E5-48/C, E3-48/C r2 Network Solutions

The majority of traditional telecommunications services are delivered from the central office (CO) or remote terminal (RT) locations, with local loop lengths ranging between 12kft and 18kft. Based on these loop lengths, ADSL and ADSL2+ services have been used for high-speed Internet service delivery. These solutions have limited data rates and flexibility required to deliver ever increasing broadband services. With the rapid expansion of IPTV, Internet video and locally sourced video on demand (VOD), and ever increasing demand for broadband, service providers will need to shorten or bond copper loops for residential or business customers. VDSL2 will be targeted for loops <5kft while ADSL2+ will be targeted for loops >5kft.

With the exception of green-field deployments or locations requiring OSP rehabilitation due to copper cable age / damage, the use of the existing copper infrastructure maximizes economic advantages by leveraging depreciated OSP facilities, and in many instances, it may be the only practical, cost-effective alternative for delivering 20 Mbps or greater broadband services. Based on common Carrier Serving Area (CSA) deployment practices, RBOCs typically have SAI cross-connects and ILECs have copper bundle taper points, which create an optimal location for VDSL2 FTTN / FTTC service delivery. Generally speaking, this is the point in the access network OSP where transmission cables (ex: 600 pair) transition to multiple, smaller distribution bundles creating an optimal location for VDSL2 FTTN service delivery.



Sample Access Network Outside Plant Carrier Serving Area (CSA)

In a serving area / distribution network, the Calix E3-48C, E5-48/C, E3-48/C r2 will be deployed close to a cross-connect node location, hence the name “fiber to the node” (FTTN). Given that this deployment location positions the node within a few kilofeet of the customer premises, as opposed to 12kft to 18kft for DLC deployment, it will serve a smaller number of subscribers. The E3-48C, E5-48, E5-48C maximizes service delivery, for data / IP video and overlay POTS services, on loops of up to 12Kft in length.

In the case of FTTP architectures, the node is the location of a passive optical network (PON) splitter or active Ethernet optical line terminal (OLT). Therefore, they enable a future migration to PON or point-to-point Ethernet FTTP OLT. This migration path permits service providers to offer competitive triple play services with minimal initial investment, transitioning to a pure fiber based model when the technology, price point, and customer demand all become feasible for the next evolutionary network transformation step.

The Calix E3-48C, E5-48/C, E3-48/C r2 enables facilities-based wireline service providers to deliver up to 100 Mbps bandwidth over the existing copper infrastructure by decreasing the size of carrier serving areas (CSA) thus maximizing broadband performance. The increased subscriber bandwidth enables service flexibility, which is critical to the delivery of advanced IP services, such as IPTV, network VOIP, and high-speed Internet services. The Calix E3-48C, E5-48/C, E3-48/C r2 leverages vectored VDSL2 and ADSL2+ technologies, along with a robust Ethernet switching fabric, to perform Layer 2 aggregation / switching and Layer 3-aware service delivery. Up to five dedicated Ethernet network interfaces can be leveraged to provide high capacity service delivery.



Chapter 2

System Management

Calix E-Series can be managed using powerful, secure, OSMINE-compliant Web-based network management tools. This chapter provides an overview of these client management interfaces, available network management tools, and SNMP support.

Topics Covered

This chapter covers the following topics:

- Web Browser Interface
- Command Line Interface
- Calix Management System
- Security Management
- Fault Management and Event Monitoring
- Link Layer Discovery Protocol (LLDP)
- ERPS Topology View
- SNMP

Client Management Interfaces

The E-Series management interfaces provide the following features:

- Equipment and software configuration
- System turn-up and test
- Service provisioning
- Performance monitoring
- System logging
- Debugging tools
- Software upgrades
- Security (authentication, authorization, and privacy)

Client management interfaces

The E5-series has four client management interfaces:

- **Embedded Web interface:** An embedded browser-based graphical user interface allows you to use a standard Web browser to manage the E-Series devices.
- **Command-Line Interface (CLI):** A Command Line Interface text interface for provisioning, maintaining, and troubleshooting E-Series devices.
- **Calix Management System (CMS):** This graphical user interface is a stand-alone Java-based client for CMS server administration and for provisioning, maintaining, and troubleshooting multiple E-Series (E7 OS) and C7 networks.
- **SNMP:** A standards-based Simple Network Management Protocol interface, supporting gets and traps, as the primary E-Series management interfaces for integration with third-party Operations Support Systems (OSSs).

See *E7-2 Platform Description* (on page [31](#)) for details on the E7-2 management access port.

See *E7-20 Chassis* (on page [71](#)) for details on the E7-20 management access port.


See *E3-8G Node-Based OLT Description* (on page [78](#)) for details on the E3-8G management access port.

See *E3-48C Platform Description* (on page [84](#)) for details on the E3-48C management access port.

See *E5-48/E5-48C Chassis Attributes* (on page [103](#)) for details on the E5-48/E5-48C management access port.

- Allows complete provisioning and service activation
- Enables individual device software upgrades and backups to minimize network disruption of subscriber services
- Allows technicians to quickly provision services
- Comprehensive troubleshooting and debugging tools
- Accessible from any of the E-series Ethernet ports

Embedded web interface design



The embedded web interface is comprised of three functional areas (frames), defined as follows:

- 1 **Navigation Tree:** Displays the physical and logical elements of the system. The E7 chassis is at the highest level, with its fan tray, installed cards with their physical ports, logical sub-interfaces (VLANs, logical Ethernet ports), and installed ONTs. Click an object on the navigation tree to display its attributes in the main window (Work Area).
- 2 **Work Area:** Displays information and attributes about objects selected on the Navigation Tree. View and modify settings in the Work Area, which includes overhead tabs and sub-tabs for displaying more specific functions for the selected category.
- 3 **Alarm Status Area:** Displays a table of all standing alarms, events, and threshold events, and includes a color-coded status monitor.

Embedded web interface controls

The embedded web interface's basic operational controls are performed as follows:

1. Click an object on the navigation tree to display its attributes in the Work Area.
2. Click the overhead tabs and sub-tabs to display menu options for that function.
3. Click in the ID column of listed objects to view and edit parameters, and then click **Table View** to return to the table of objects previously viewed.
4. Click the **Apply** button on overhead tabs, sub-tabs, or in pop-up dialog boxes to apply provisioning changes.
5. Click **Logout** to end a session and log out of the web interface.

Command Line Interface

The Calix E-Series has an embedded Command Line Interface (CLI) for system management access via local or remote TCP/IP connections and local console connections.

The CLI is a text interface for provisioning, maintaining, and troubleshooting E-Series nodes. The CLI is a straightforward command interface. You type commands and their parameters on a single line and the commands execute when you press the Enter or Return key. You can access all features and functions of the E-Series platform through the CLI interface.

For detailed information, see the *Calix E-Series (E7 OS) CLI Reference*.

Command assistance

The E-Series offers command interactive help, command completion and history, as well as the ability to recognize abbreviations. As you type, the following keys can assist you with specifying the command you want:

- The **?** key displays the interactive help that is a list of valid keywords or values, along with brief descriptions. This helps new users learn the system as they use it.
- The **↑** and **↓** arrows allow access and use of previous commands.
- The Ctrl+P key combination accesses the previous command.
- The Ctrl+N key combination accesses the next command.
- The **Tab** key expands a partially typed keyword and completes commands when unique abbreviations are recognized by the system. These features allow experienced users to administer system control with a minimum of typing.
- The **←** and **→** arrows allow editing in the current line.
- The Ctrl+A key combination moves the cursor to the beginning of the line.
- The Ctrl+C key combination cancels commands that have more than one step. For example, prompting for a password.
- The Ctrl+E key combination moves the cursor to the end of the line.
- The Ctrl+B key combination moves the cursor back a space.
- The Ctrl+F key combination moves the cursor forward a space.
- The Ctrl+U key combination erases text from the cursor to the beginning of the line.
- The Ctrl+K key combination erases text from cursor to the end of the line.

Command syntax

The straight-forward command syntax has the following rules:

- The commands are “imperative,” meaning they start with a verb and instruct the system to perform some action.

- For commands that have optional parameters, you can specify them in any order or omit any that are unnecessary.
- When naming a database backup file, only use the following acceptable characters: letters, digits, underscore, and dash.
- The CLI command buffer allows a maximum of 4k characters for a single "paste" operation.
- When restoring a backup database file from a Windows-based server, use a forward slash (/) as a path separation character when working with FTP application operations.
- When referencing an entity name that has spaces, use quotation marks around the entity name. For example, use "Test 123" to enclose the name Test 123.

Note: If you type an invalid string—for example, the name of a command or statement that does not exist—the message “Invalid input detected at '^' marker” appears. The caret (^) indicates where the error is located.

Displaying valid keywords

The CLI provides the ability to display a list of valid keywords at the current level. This section presents some examples of how to evoke lists of particular command keywords and corresponding information.

Example 1:

Type "?" at the prompt and press **Enter** to see a list of the first-level commands available.

Possible commands:

activate	Activate system components.
add	Adds items to a collection.
apply	Apply a template.
clear	Clear system diagnostic information.
commit	Commit to a software version.
create	Create system equipment and service objects.
delete	Deletes system components.
dig	Find the IP address for a hostname.
disable	Disable system equipment and service objects.
download	Download system components.
enable	Enable system equipment and service objects.
exit	Logout from this session.
extract	Extract system data to server.
help	Ask for help.
load	Load a database backup onto system.
logout	Logout from this session.
ping	Ping another host.
quit	Logout from this session.
reboot	Reboot shelf, using the current software version.
remove	Remove items from a collection.
reset	Reset system components.

restart	Restart a system service.
revert	Revert to a software version.
set	Configure system equipment and service objects.
show	Show attributes of equipment and service objects.
snapshot	Captures the current system database.
switch	Switch to the loaded database.
telnet	Start a telnet session to another host.
test	Test system components and services.
tracroute	Find the route to another host.
turn-up	Starts the interactive turn-up tool.
unlink	Unlink system components.
upgrade	Upgrade system components.

Example 2:

Type **clear ?** to see a list of all possible completions for the clear command.

Possible completions:

dsl-template	Clear parameters in a DSL port template.
interface	Clear interface counters.
lldp-neighbor	Clear LLDP neighbor information.
log	Clear system logs.
ont	Clear ONT counters.
ont-port	Clear ONT port information.
pm	Clear performance monitoring statistics.
pots-port	Clear information for POTS ports.
stats	Clear cumulative statistics.
sw-resource-alarm	Evaluate software resource usage, clearing alarm if appropriate.

Example 3:

Type **clear log ?** to see a list of all possible completions for the clear log command.

Possible completions:

alarm	Clear the alarm log.
dbchange	Clear the database change log.
event	Clear the event log.
security	Clear the security log.
tca	Clear the TCA log.

Displaying configuration settings

The CLI also provides information on the current configuration settings and performance monitoring data. This section presents an example of how to evoke lists of particular information.

Example 4:

Type **show interface <interface name>** and press **Enter** to list the current configuration settings for the specific Ethernet port interface.

Note: For E7, you can specify the name of a LAG interface, a card number, or a card/Ethernet port specification. For example, 2/g1. Permitted values for the port type are "g" (for Gigabit Ethernet) or "x" (for 10-Gigabit Ethernet).

Configuration settings shown:

- Name
- Description
- Role
- Current Rate
- # VLANs / Tag Actions
- Status

Example 5:

Type **show ?** to see all of the possible sets of configuration parameters.

Calix Management System

The Calix Management System (CMS), allows operators to manage services while understanding their relationship to the network infrastructure. CMS is a client/server application that provides a graphical user interface for managing multiple Calix nodes. Service-oriented management includes:

- Display of Calix nodes in the network tree
- Integrated service activation
- Global service templates and policies
- Scheduled backups
- Scheduled software upgrades
- Node and port search

Comprehensive network management tools let operators create physical and logical topology maps, engineer traffic flows, and manage network commissioning and software upgrades. Network inventory, alarm surveillance and PM collection are enabled by the E7 element management systems.

- See the *Calix Management System (CMS) Guide* for instructions on using the Calix Management System.
- See the *Calix Management System (CMS) Installation and Upgrade Guide* for Linux or Solaris, depending on your host server operating system, for instructions on how to install or upgrade the CMS server software.



Note: See the CMS release notes for the minimum requirements on client computer hardware and software. (To obtain CMS software, submit an online request through the Calix Resource Center).

CMS Features

Calix Management System (CMS) is a client/server application for managing, monitoring, and configuring multiple Calix Active Ethernet (AE) and GPON ONTs, B6 cards, C7 networks, E3/E5-100 service units, E-Series, E5-30x/E5-520 and F5 and FD Passive Optical Network (GPON) platforms.

For detailed instructions on how to manage Calix network elements, see the configuration, services, and other guides listed in Calix documentation set, accessible by logging into My Calix (www.calix.com/my-calix (<https://www.calix.com/my-calix>)) and browsing the My Calix Documentation Library.

Tasks that you can perform in CMS include the following:

- Manage up to 7,500 network elements and 20,000 AE ONTs using the java-based CMS Desktop application and the browser-based CMS Web application on your Windows or Linux PC or Solaris workstation.
- Support up to 150 concurrent CMS user sessions: Up to 100 with CMS Desktop, 30 with CMS Web, and 20 with TL1 (C7 only), CLI and Embedded Web interface (EWI) cut-through.

Note: XML NBI login sessions count toward the 150-session client session limit for the CMS server.

- Create user groups to customize user privileges by topology (down to the port level) and service type; manage user and interface connection security settings. See Security Management.
- Create customized network groups to organize network elements. See Configuration Management. In this section you can also find instructions for using the following features:
 - Create and manage Calix network elements.
 - Search for and generating custom inventory reports of physical elements (such as networks, nodes, shelves, cards, and ports) and logical elements (such as FTTP, TDM, Ethernet and IP, SONET, video, and interface groups). See Searches and Inventory Reports.
- View aggregated alarms throughout your Calix devices and query historical logs of alarms, events, threshold crossings, network changes, CMS updates, and security events. Optionally modify alarm severity and descriptions, suppress alarms, and utilize the alarm acknowledgment feature. See Fault Management.
- Download and query historical performance monitoring data from C7 networks and E-Series platforms. See C7 and E-Series Performance Monitoring.
- Use a wide variety of C7 business reporting tools. See Business Reports.
- Perform a variety of administrative tasks, such as backup and restore shelf and node data and CMS objects. See Administrative Tasks.
- Troubleshoot CMS alarms and perform maintenance tasks..

View and Manage Calix Equipment

Your user privileges determine which network and device functions and services you can view and configure within CMS. This guide describes CMS features available at the network or node level and above on the Navigation Tree. These include regions, networks, and system and security settings.

CMS supports up to three versions of Calix network element software: the current version and the two previous versions. CMS handles most version translation internally and seamlessly.

For descriptions of and instructions for network element-level provisioning features, see the Calix Documentation Set. A few of the key management features only available in CMS include:

- Global profiles for Calix C7 networks, E-Series service units and platforms, AE ONTs, and global subscriber templates for C7 networks and E3/E5-100 nodes (for details, refer to the platform-specific documentation).
- Service activation for AE ONTs; xDSL and POTS ports and ONTs for C7 networks, E-Series systems, and F5 nodes (for details, refer to the platform-specific documentation).
- A tool for configuring Ethernet services and devices across Calix platforms (for details, see Provisioning Cross-Platform Ethernet Topologies).
- A tool for turning up Ethernet-over-SONET transport and service drops on C7 networks (for details, refer to the *Calix C7 Provisioning Guide*).

Future platform software release support

You can access C7 networks and E-Series system running a future platform release from CMS. For more information, see Future Calix Platform Release Support.

Features not supported

CMS does not support the following:

- Performing software upgrades for a C7 network. Note, however, that you can view the current software versions of all C7 networks in the Network Details screen to track the upgrade progress across the all connected networks.
- Making bulk provisioning changes to multiple nodes or shelves belonging to different C7 networks. You can only change node or shelf parameters for one network at a time.
- Making bulk provisioning changes to platforms other than C7.
- Business reporting on platforms other than C7.

CMS Interfaces

CMS has five client interfaces:

- **CMS Desktop:** A stand-alone, Java-based client for CMS server administration; provisioning, maintaining, and troubleshooting Calix network elements; reporting Performance Monitoring statistics; and viewing online help.
- **CMS Web:** A Web browser-based interface for querying and reporting inventory statistics from Calix network elements; querying and reporting Performance Monitoring from Calix nodes; running historical alarm and event queries; downloading CMS Desktop; and viewing online help.
- **CMS TL1 Interface:** A Transaction Language 1 (TL1) text interface for provisioning, maintaining, and troubleshooting C7 networks.
- **CMS SNMP Northbound Interface:** A unified interface for alarm management integration with back office systems.
- **CMS XML Northbound Interface Applications Processing Interface:** An API that provides back-office system integrators with the ability to activate service requests on C7, E-Series, AE ONTs, and E3/E5-100 units. The interface can also use the subscriber profiles within the CMS database. See the *Calix Management System (CMS) XML Northbound Interface API Guide*.

CMS Connections

Use the following methods to manage Calix equipment connections to the CMS server:

- **B6 networks:** CMS Desktop and Web interfaces. Connections between the CMS server and managed B6 networks consist of IP network links from the LAN to an Ethernet port on each B6 card. For options and details about connecting to the Ethernet ports on a B6 network, see the Calix B6 platform documentation.
- **C7 networks:** CMS Desktop, Web, and TL1. Connections between the CMS server and the managed C7 networks consist of IP network links from the LAN to an Ethernet port on an AMP or ATP card in the C7 network. For options and details about connecting to the Ethernet ports on a C7 shelf, see the *Calix C7 Provisioning Guide*.
- **E3/E5-100 service units:** CMS Desktop and Web interfaces. Connections between the CMS server and E3-12C, E3-48, E5-110, E5-111, E5-120, and E5-121 nodes consist of IP links from the LAN to an Ethernet interface on the node. For options and details about connecting to the E3/E5-100 node, see the appropriate Calix installation documentation.
- **E-Series platforms:** CMS Desktop and Web interfaces. Connections between the CMS server, E-Series, and E5-30x/E5-520 platforms consist of IP links from the LAN to the alarm, management, and Ethernet line interfaces on the node. For details about setting up physical connections, see the Installation Guide for each platform. For details about setting up management interface communication with these E-Series platforms, see the *Turn-Up and Transport Guide*.
- **F5 nodes:** CMS Desktop and Web interfaces. Connections between the CMS server and the F5 consist of IP links from the LAN to an Ethernet port on the E-Management card on the node. For options and details about connecting to F5 nodes, see the *Calix F5 User Guide*.

- **700GE, 700GX and 836GE RSG Active Ethernet (AE) ONTs:** CMS Desktop. Connections between the CMS server and AE ONTs consist of auto-discovered IP links based on SNMP traps sent from the ONTs when they boot up. For options and details about AE ONTs, see the Calix publication, *Managing AE ONTs with CMS*.

TL1 northbound interface

The CMS server provides a TL1 northbound interface for integration with third-party Operations Support System (OSS) software. Once the OSS connects to the northbound interface, it receives all events from every C7 network connected to the CMS server. You can open a TL1 interface using a Telnet or Secure SHell (SSH) security session. For instructions, see Using the CMS TL1 Interface.

A TL1 northbound interface agent is also available (requires advance setup requirements with Calix) for specific Calix E-Series products.

SNMP northbound interface

The CMS server provides a normalized SNMP northbound proxy interface for integration of Calix-supported equipment with back office systems for fault and event management. CMS normalizes the device trap object formats by mapping alarm attributes to an internal format. For more information, see Northbound SNMP for Alarms.

CMS Desktop Views

Permission-based views

All CMS screens and reports are filtered to show only those resources that the user group is authorized to view. Screens, and screen items such as menus and functional areas, are filtered by permissions. Administrators can view the full range of screens, menus, and buttons for their functional areas. Users with more restricted roles see only those items they have permission to view and work with. Items they do not have permission to view or work with are not visible.

Initial work area view: Subscriber search pane


When you open CMS Desktop, the subscriber search pane displays in the CMS work area (the area of the screen to the right of the navigation tree).

For instructions on how to use the Subscriber Search tool, see Performing a Subscriber Search.

To toggle to the CMS administrative and node configuration work area menus, at the top of the screen, click the **Previous Panel** link. (To toggle back to the subscriber search pane, click the **Last Search** link.)

CMS Desktop screen frames

The main screen in iMS and CMS Desktop is divided into four frames:



- 1. Navigation Tree:** Displays a customizable, hierarchical list of network element objects: networks, nodes, shelves, cards, ports, and paths. CMS additionally displays administrative settings at the highest level, followed by a root region and regions and network groups for organizing network elements.
- 2. Work Area:** When the Subscriber Search pane is not displayed (CMS only), this area of the screen displays information about the level selected on the Navigation Tree. At the network or node level and below on the Navigation Tree, the Work Area provides access to view or modify network and node configuration settings. On the Navigation Tree, you can drill down to view details down to the port or path level in the Work Area. In CMS Desktop, to toggle to the configuration menus from the Subscriber Search pane, at the top of the screen, click the **Previous Panel** link; to toggle back to the Subscriber Search pane, click the **Last Search** link.
- 3. Status Monitor:** Displays a count of critical, major, and minor alarms. In addition, you can use the drop-down menu to the right of Status Monitor and select **Resource** to display total memory allocated and in use by the application.
- 4. Alarm Table:** Displays a customizable table of all standing alarms, events, and threshold events.

At the bottom of the screen, the status bar displays the user name on the left side and the date and time (for CMS, the CMS server time) on the right side.

Region and network group views

Regions in CMS are groups of networks and network groups. Network groups are groups of network elements that share something in common. You can create custom regions and network groups based on topology (for example, all C7 networks in a particular state), function (for example, networks that manage video services), or some other meaningful classification.

From the Navigation Tree, access and customize the region and network group views that display in the Work Area:

- **At a region level:** The CMS Navigation Tree contains a root region that cannot be deleted or renamed. Under the root region, you create regions, sub-regions, and network groups for viewing network elements.

On the CMS Region tab, view a topology of the selected region. Other options include viewing region, network group, network, and node details within the region, as well as organizing network elements on a map using GoogleMaps.

- **At a network view level:** Options include viewing network and node details within the network group, organizing network elements on a map using GoogleMaps, and creating customized topology links between network elements to view in the Topology Work Area.

Network view

View network element configuration settings down to the port or path level in the Work Area.

Network filtering options

With the Navigation Tree and the Selection Filter, you can narrow the list of network elements that display in the Work Area:

- Use the Navigation Tree: View for a selected level on the Navigation Tree down to the node, shelf, card, and port levels.
- Use the Selection Filter (C7 networks only): At the bottom of the Navigation Tree pane, use the Selection Filter to display elements in the Work Area using an available filter. For example, display network elements relating to a facility type or interface group.

Security view

At the CMS level, the Work Area provides access to CMS user security and system defaults.

CMS user interface modifications and preferences

These features allow the user to customize table layouts and set basic user parameters that are persisted by a user login.

- To modify table layouts for any table in the CMS desktop, hide and order table columns to display the desired data in the order preferred by the user. This enables the user to focus on data of interest for the specific function being performed.
- To define user preferences for the following CMS aspects, access the dialog box under the CMS File menu:
 - **Alarms** – sets the rate at which CMS desktop is updated to the network
 - **Colors** – sets dialog colors
 - **Dialogs** – sets default behavior for progress dialog completion
 - **Logs** – enables CMS Desktop logging
 - **Ping Settings** – defines keep alive for firewall between client and server
 - **Services** – defines default for showing/hiding services
 - **Tabs** – enables tabs to scroll rather than stack
 - **Tables** – enables/disables column auto sizing

CMS Security Features

CMS supports secure and non-secure northbound (Calix network element to CMS server) and southbound (CMS server to Calix network element) communication. The CMS user provisions the northbound security mode. CMS detects the security setting of any southbound Calix node, and connects to it accordingly. For more information, see Connection Security Modes.

Calix Networks have optional Secure Shell (SSH) security. FiberDrive nodes have no southbound security.

The northbound and southbound interfaces may or may not run in the same mode. For example, the CMS southbound interface could run on an unprotected LAN while the northbound interfaces run on a protected LAN. In this mode, the northbound interfaces could communicate with the protected LAN to OSSs that do not support secure interfaces, while the southbound interfaces can talk securely over an unprotected LAN.

By default, SSH communication is always available from the CMS client to the CMS server, regardless of security mode of the server. However, when you enable the secure mode, CMS disables and drops all insecure (non-SSH) interface connections.

The SSH transport for CMS client communication is mostly transparent for CMS users. CMS Desktop embeds the SSH connection libraries; therefore, it requires no additional client software to run the application in secure mode. CMS TL1 users must use an SSH client instead of a Telnet application.

CMS offers the following features for secure communications:

- **Single login restriction:** CMS enforces a single login restriction option for user groups that restricts a user to one session at a time

- RADIUS for CMS user management
- Southbound SSH security for iMS and TL1 connections to C7 networks
- Northbound security to higher level OSSs:
 - An SSH interface for TL1
 - SSH for the interface between CMS Desktop and the CMS server
 - HTTPS for CMS Web and login authentication

Default CMS Server Ports

The following table shows the default incoming CMS server interface ports:

Port Number	Port Description
21	For FTP server, when used for C7 image management
67	For DHCP server for AE ONTs, when configured on the CMS host server
69	For TFTP server for AE ONTs, when configured on the CMS host server
123/UDP	For Network Time Protocol (NTP) for AE ONTs, when configured on the CMS host server
162/UDP	SNMP trap for all Calix network elements
514/UDP	For Syslog for AE ONTs, when configured on the CMS host server
1098	* Used for RMI Naming service (for F5 and FD device management) when a CMS client communicates with the CMS server through an EJB remote call
1099	JNDI Lookup for F5 and FiberDrive (FD) management
1645	For RADIUS server, authentication for C7 image management
2222/UDP	FD trap receiver for database changes and alarms
2626	* Northbound CMS TL1 for C7 networks
2627	* TCP Port (TL1 in non-echo mode). Not commonly used.
5055	* Socket used by OSMINE TL1 agent (for approved customers)
5056	* Telnet connection to OSMINE TL1 agent (for approved customers)
5432	Postgres listen port (for PostgreSQL only)
8080	* CMS Web
8093	Default JMS server port used for communication between the CMS server process and server-side auxiliary processes.

Port Number	Port Description
8179	Debug Console, used by Calix technical assistance during support calls for server diagnostics
8190	* SSH Server / Secure TL1
8196	CMS Desktop / Secure CMS Desktop
8198	* Cut-Through iMS
8443	* Secure CMS Web / HTTPS
8750	SELT service port
9199	* Cut-Through TL1 / CLI
18080	* Port used by CMS Desktop, CMS Web, and the CMS XML NBI for provisioning of supported Calix network elements
18443	* Secure Web port for E3/E5-100, E-Series, and F5
20002/UDP	FD trap receiver for responses to Synchronous requests
28080	* HTTP Cut Through
28443	* HTTPS Cut Through

* Indicates that the port number is configurable in CMS.

Note: Typically, port 22 is used by SSH service and port 23 is used by Telnet service on a UNIX host. To open a SSH connection, you must have access to port 22. To open a Telnet connection, you must have access to port 23.

Default Outgoing CMS Server Ports

The following are the default outgoing CMS server interface ports:

Port Number	Port Description
21	FTP server connection for CMS scheduled tasks
22	Secure connection to C7 and E7 CLI AE ONT configuration file transfer (via SFTP)
23	TL1 connection to C7 and Telnet connection to other Calix network elements
80	Web connection to C7, E3/E5-100, and E7 nodes (HTTP)

Port Number	Port Description
161/UDP	* SNMP requests to AE ONT, E3, E5, E-Series, and F5 devices
443	Web connection to E3/E5-100, and E7 nodes (HTTPS)
830	* Network Configuration Protocol XML communication with E-Series nodes
1812/UDP	Radius server authentication
20002/UDP	Acknowledgments for FD database changes and alarms
20012/UDP	Synchronous requests to F5
50000	Type-Length Value (TLV) connection to C7

* Indicates that the port number is configurable in CMS.

System Software Upgrades

Upgrade Options

A scheduled CMS software upgrade task includes the following options:

- **Both File Distribution and Reset:** Transfer the new software release and reset each node to operate with the new software.*
- **File Distribution Only:** Transfer the new software release to each node included in the scheduled task without initiating a reset.
- **Reset Only:** Reset each node to start operating with the new software.*

The upgrade options available depend on the platform and deployed GPON ONTs.

For E-Series VDSL2 systems and E3-8G OLT, there is only one upgrade process for a system upgrade.

This includes shelves with VDSL2 cards, shelves with a VCP card, or independent VDSL2 nodes (E5-48, E5-48C, E3-48C, and E3-48C r2). However, a VCP is managed as a standalone unit that, **MUST** be upgraded independent of E7-2 systems. A single VCP image supports both VCP-192 and VCP-384 cards.

E7 GPON System Upgrade Options

There are system upgrade options depending on the deployed ONTs and whether the system upgrade package bundles (includes) the ONT images:

- bundled = 700/836 P-Series ONTs (earlier than R3.1 only)
- unbundled = 844G/854G GigaFamily and T-Series ONTs (also P-Series R3.1 and later)

For ONT images that are "bundled" with system software, they are also available as a separate package to allow for ONT-only firmware upgrades.

For ONT images that are only available "unbundled" with system software, —such as the 844G and 854G GigaCenters— you must use the ONT-only upgrade to load the ONT images onto the E7 GPON OLT so that when you deploy the ONTs and attach them to the GPONs, the appropriate firmware for operation will be pushed down from the E7 to the ONTs.

- For R2.3.50 and later versions of E7 GPON systems, Calix supports the following upgrade options:

- **Merge firmware (Retains ONT images and then upgrades OLT and P-Series):**

Merges the previous “include-bundled” and “retain-existing” actions for handling stored ONT firmware images. This option upgrades the E7 system and P-series with ONT images included in the E7 package. All other ONT images (for example, T-Series, GigaFamily, and GPON interface adapter module) that are not part of the system upgrade package and are currently on the system OLT are retained. Then, perform an ONT-only upgrade to install the newest unbundled ONT images. For an OLT and ONT combination upgrade process, you can preload system software and ONT firmware, and then reset the system.

Note: This is the recommended option for upgrading systems with GigaFamily or T-Series ONTs. However, when a system upgrade is from a version prior to the version that includes the "merge" option (R2.3.60), you must re-load any GigaCenter and T-Series ONT images after the system upgrade.

- **Retain existing ONT firmware (Only upgrades OLT):**

Upgrades the E7 system only, retaining GPON ONT image(s) that are currently on the system. You can then independently upgrade P-Series, T-Series, and GigaFamily ONTs, using the ONT-only upgrade process.

- **Include bundled ONT firmware (Removes ONT images and then upgrades OLT and P-Series):**

Upgrades the E7 system software, removes all existing ONT images (regardless of how they were installed onto the system), and then installs the ONT images (P-Series only) in the E7 package.

- For R3.1 and later E7 GPON systems, Calix supports the above upgrade options and the additional following option:

- **No packages of ONT firmware (Removes ONT images and then upgrades OLT)**

Removes all existing ONT images (regardless of how they were installed onto the system) and upgrades the E7 system software. When you preload new ONT images to the system before using this option, upon system reset the OLT and ONTs restart using the new software and firmware.

Note: Although R3.1 does not bundle ONT images with system software, the upgrade options above can still be used.

ONT Mode	Below R3.1.0		R3.1.0 and Above
	P-Series	Other ONTs	All ONTs
Merge-bundled (default)	Upgrade	Retain	Retain
Include-bundled	Upgrade	Remove	Remove
Retain-existing	Retain	Retain	Retain
No-packages	---	---	Remove

Unbundled ONT Firmware

- **ONT-only Upgrade**

Upgrades all compatible P-Series, T-Series, or GigaCenter ONTs configured in the E7 system, or a subset of those ONTs, without upgrading the E7; allows a single instance of an ONT image type to be upgraded on a PON at a time.

Calix recommends using a CMS scheduled task to automate the ONT-only upgrade process.

- **Preloading ONT Firmware**

For R2.5 and later versions of E7 GPON systems, Calix supports the process of downloading unbundled ONT images prior to resetting and activating the new E7 system software. Then you can invoke a single system reset at a convenient time to start the system running new software and the ONTs running new firmware.

See Preloading ONT Firmware for a GPON System Upgrade.

Calix recommends using a CMS scheduled task to automate the ONT-only upgrade process.

* When resetting the software, you apply one method of resetting the nodes included in the scheduled task:

- **Sequentially:** For ring network topologies, perform a reset on one node at a time.
- **In Parallel:** For linear chain network topologies, perform a reset on all nodes in parallel.

Following the software reset, the new system software is committed as the default software.

Note: For a sequential upgrade, CMS acts as the master during the software reset phase. If the CMS server shuts down before the scheduled task completes, you can re-run the task once CMS is restored (in the scheduled task list, right-click the task and click **Re-Run**). This restarts the software upgrade task from the beginning, including file distribution.

Calix recommends using the CMS “ONT Firmware Upgrade” scheduled task to automate ONT-only upgrade process.

Service-affecting procedure

Calix recommends performing the reset phase of the upgrade during a maintenance window.

- For linear chain topologies, the reset phase is service-affecting for several minutes.
- For non-redundant services on ring topologies, the upgrade reset phase is service-affecting for several minutes.
- For redundant services on ring topologies, the reset phase causes a momentary disruption in service, and the network runs without alternate path protection.



ALERT! Calix strongly recommends scheduling separate software upgrade tasks for ring and linear-chain network topologies.

Do not schedule a Sequential upgrade for a linear-chain network topology, as this results in multiple service outages of several minutes each during the reset phase.

(E7-20 Only) In-Service SCP-10GE to SCP2-10GE Upgrades

E7-20 systems can upgrade from SCP-10GE to SCP2-10GE cards while in service, without disruption of currently operating system services.


Prerequisites

The following prerequisites must be met before performing this procedure:

- The system is already running release 3.0.
- The maximum MTU size is set to ≤ 9390 bytes.

Note: SCP cards support up to 9600-byte MTUs, while SCP2 cards support up to 9390-byte MTUs

Once the software is upgraded, the operator can go through a procedure to upgrade the SCP cards, one at a time, while maintaining system operation.




Compass Application Suite

Compass is an award-winning suite of software applications that enable communications service providers to accelerate their business transformation. Compass simplifies the complex challenges of broadband service planning, activation, and maintenance through applications focused on key strategic functions for all service providers:

- Connectivity, service activation, and device management at the customer premises
- Analysis of broadband traffic, usage patterns, and applications

Compass by Calix equips you with the tools and information needed to optimize the communications experience for your subscribers, lower operating expenses, and increased revenue.

See the Calix Documentation Library for information about Compass.



Security Management

The E-Series utilizes a hardened operating system and Calix aggressively follows the security related patches and guidelines to ensure that the E-Series remains the most secure solution possible. Calix also customizes the E-Series (E7 OS) to ensure that unnecessary services that could be a security risk to the system are not included or enabled. The E-Series protects itself against network based attacks by providing host protection features. (e.g. Slow path packets are rate limited aggressively on a protocol basis, to ensure that the system cannot be overwhelmed.)

Configuring Secure System Login Access

Remote Authentication Dial-Up Service (RADIUS) server is for user login authentication and accounting of administrative access via a centralized user database.

In this model, an E-Series acts as a RADIUS client and during the setup of a new login session, asks the RADIUS server if the provided username / password is a legitimate user account for the system.

Both CLI and web GUI sessions can make use of RADIUS. The E-Series can be configured to determine the priority order for user management: RADIUS / local / both (with a priority order).

RADIUS is a connectionless client/server authentication, authorization, and configuration system that provides secure access to E7 networks.

The E7 as RADIUS client:

1. Passes user information to the RADIUS servers.
2. Acts upon the response.

The RADIUS server:

1. Receives user connection requests.
2. Authenticates the user.
3. Returns all configuration information necessary for the client to deliver service to the user by granting the appropriate access permissions.

All communications between the client and server are encrypted using a shared secret known to both the client and server, but not transmitted with the login.

You can configure multiple external radius servers for redundancy. The RADIUS administrator enters users into the RADIUS server database.

- 4 authentication RADIUS servers and 4 accounting RADIUS servers are supported.

- All of the authentication RADIUS servers are assumed to have the same authentication information, that is, server replication. The system communicates with the “best” server, and then only sends to the next server if it does not get a response. Where “best” is determined by the success rate of getting responses to recent requests.

Note: From the System form, you can specify the order to follow for user authorization: local database only, local database then RADIUS server, RADIUS server only, RADIUS server then local database.

User access levels

The E-Series user accounts that are set up on RADIUS servers must specify the Class attribute for the user to successfully login with the correct access level. If the Class attribute is not specified, the default access level is "read-only."

The following examples show the Class attribute specified on a RADIUS server for a user at each of the three user access levels:

```
john    Cleartext-Password := "john123"
        Service-Type = Administrative-User,
        Class = "read-only"

mary    Cleartext-Password := "mary123"
        Service-Type = Administrative-User,
        Class = "provision"

paul    Cleartext-Password := "paul123"
        Service-Type = Administrative-User,
        Class = "admin"
```

Configuring Secure Network Access

The 802.1x defines a port-based network access authentication mechanism for point-to-point devices on untrusted interfaces attempting to connect to a LAN or WLAN. On an 802.1x enabled port, an untrusted device must be authenticated prior to accessing the network. This 802.1x uses Extensible Authentication Protocol (EAP) protocol that is based on encapsulating EAP packets within the data-link layer frame to provide access control mechanisms for devices interconnected by 802 LANs, known as EAP over LAN (EAPOL).

This section describes how to configure network access authentication with the following components:

- **Supplicant** - untrusted device that is requesting network access and interfacing with an authenticator (E-Series) running EAPOL. An ONT port, xDSL interface, or a bonded XDSL interface can be a supplicant.

- **Authenticator** - a network element that is peering on the data link with the supplicant on a given physical port and relays EAP credentials between the supplicant and authenticating server. An E-Series unit can be an authenticator.
- **Authentication Server** - centralized authentication server (typically a RADIUS server) that validates the identity of the supplicant and informs the authenticator to allow or deny access.
- **Accounting Server** - centralized accounting server (typically a RADIUS server) that receives important information for security and billing purposes. If an accounting server is provisioned, accounting messages (requests, responses, etc.) begin after authentication has occurred.

Note: In the E-Series system, 802.1x accounting packets use the same P-Bit and DSCP values as management traffic, P-Bit = 5 and DSCP = 0x00.

802.1x authentication and accounting are supported by the following E-Series (E7 OS R2.5) systems:

- For xDSL ports (associated interfaces) connecting to an xDSL modem/RG with appropriate 802.1x certificate:
 - E3-48C, E5-48/48C, and E7-2 VDSL2 r1/r2 cards
- For GPON ONT Ethernet ports (in half-bridge mode) connecting to an external supplicant with appropriate 802.1x certificate:
 - E3-8G, E7-2 and E7-20 GPON cards with P-Series 700GE-Family ONTs

The following steps describe how to configure a secure network access:

1. Add an 802.1x authentication server to the network.
2. Add an 802.1x accounting server to the network.
3. Create an 802.1x profile.
4. Reference the 802.1x profile in an Ethernet security profile.
5. Assign the Ethernet security profile when adding an Ethernet service to a subscriber port.

For steps 4 and 5, see the *Calix E-Series xDSL (E7 OS) Applications Guide* and *Calix E-Series (E7 OS) GPON Applications Guide*.

Note: The E-Series implementation of security profiles applies to non-TLAN services only.

You may view the collected 802.1x statistics at the port level and card level. See the *Calix E-Series (E7 OS) Maintenance and Troubleshooting Guide*.

Note: For current E-Series systems, the highest level at which statistics can be retrieved (including 802.1x statistics) is the card level.

Forcing the authentication state on the port

- An action to force an 802.1x supplicant to reauthorize immediately.

From the Work Area at the ONT Ethernet port or xDSL associated interface, click Action > Force 802.1x.

For CLI:

```
force dot1x-dsl-bond <intfc-name> reauthenticate
force dot1x-interface <intfc-name> reauthenticate
force dot1x-ont-port <port-id> reauthenticate
```

- An action to force the supplicant to be unauthorized or authorized until the force attribute is set to none at the port level.

From the Work Area at the ONT Ethernet port or xDSL associated interface, in the force 802.1x list, select the parameter.

For CLI:

```
set ont-port <port-id> none|authorized|unauthorized
set interface <intfc-name> none|authorized|unauthorized
set dsl-bond-interface <intfc-name> none|authorized|unauthorized
```

New Key and Self-Signed Certificate

By default, the E-Series uses SSL to enable https access to the unit to provide a secure network management solution. Each E-Series generates a unique key and self-signed certificate at the beginning of life (including database resets) and whenever the IP address of the management interface changes. The first time a browser connects to an E-Series using https, the unit exchanges an SSL certification with the subscriber computer and the browser asks the user to accept the E-Series SSL certification.

To regenerate a new key and self-signed certificate, use the CLI command **create certificate**.

Note: Whenever a new certificate is generated, current web sessions to the E-Series are impacted.

Fault Management and Event Monitoring

The E-Series reports all events through the web GUI, CLI, and CMS, and maintains local logs of system activity. Up to 500 activity records are stored on the local platform for each log type. CMS is utilized to access older records, effectively enabling historical record archiving and searching capabilities. Notification logs are stored for each of the following log types:

- Alarms
- Alerts
- Database changes
- Security events
- Threshold crossings

All local notification logs can be manually cleared.

Syslog

SYSLOG per RFC3164 is also available for Operations, Administration, Maintenance, and Provisioning (OAM&P) alarms, events, database change, security, and threshold crossing alert (TCA) notifications. Each SYSLOG server record contains a DNS resolvable name or IP address, description, administrative state, and an indication of which OAM&P notification types are to be sent to this server.

The system supports four SYSLOG destinations, with service provider defined SYSLOG “facility” values. The E7-20 uses the Local0 to Local7 facilities for each OAM&P notification and allows each type to use a different SYSLOG facility value. For example, alarms could be Local1, events Local2, security also Local2, with database change and TCA turned off.

Historical Performance monitoring

Historical performance monitoring (PM) information is maintained within the E-Series and can be retrieved for traffic management and network engineering analysis. The E-Series captures various systems and service statistics and maintains this information in 15 minute bins for one day and 24 hour bins for one week. Statistics and performance monitoring data is collected and binned for the E-Series 10GE and GE interfaces.

In addition to the Ethernet traffic PM data, the system supports ONT T1 PM, ONT PWE3 service PM, ONT PON PM, and additional ONT Ethernet PM statistics and error attributes.

Threshold-crossing Alerts

The E-Series also has the ability to issue an alert based on user defined criteria such as a high rate of discarded frames or frame errors within a specified period of time.

Real time diagnostics

The E-Series provides several on-demand tools for real-time performance monitoring and diagnostics. Additionally, the E-Series continuously monitors many physical and system level attributes, including:

- All physical interfaces are monitored for power, temperature, loss of signal (LOS), etc.
- Continuity check messages automatically monitor faults within an ERPS ring

MAC table Display

The E-Series MAC learning table can be displayed to enable debugging of layer 2 learning and forwarding values. If desired, the MAC table can be displayed for a specific physical link or VLAN instead of the entire platform.

DHCP snooping lease table

When DHCP snooping is enabled on a VLAN, the E-Series will track all DHCP activity on that VLAN within the system and maintain a table of DHCP leases granted. The table of granted leases is retrievable via any of the management interfaces and contains the following attributes, which may be searched or filtered by: VLAN ID, IP address, MAC address, ONT Port.

VLAN monitoring

The E-Series can monitor up to 30 VLANs and collect the following items:

- Packets transmitted per VLAN per port
- Bytes transmitted per VLAN per port
- Bandwidth utilized per VLAN per port

Rogue ONT detection and quarantine

When the E7 enables the GPON Rogue ONT Detection and Quarantine feature that was introduced in E7 software version R2.2.90, the E7 detects faulty ONTs on a PON that transmit optical power in upstream timeslots, where no ONT is assigned to transmit (quiet timeslot), and then automatically removes these "rogue" ONTs from operation. This detection and quarantine behavior quickly eliminates the rogue operation of faulty ONTs that are causing upstream optical collisions and affecting services for other ONTs on the PON.

Port mirroring

The E-Series is designed to mirror traffic from any Ethernet, xDSL, or GPON OLT port to a “mirror” port on an Ethernet port, or the centralized E7-20 SCP-10GE card. This function is useful in a number of circumstances including debug scenarios and network troubleshooting. In order to enable port mirroring, a port must first be designated as the destination for the mirrored traffic. Once this designation has been made, selected VLANs on any Ethernet, xDSL, or GPON OLT port can be added to the mirror for observation. A port on the SCP configured as the destination for mirrored traffic should be reserved for this purpose and should not be used for regular network traffic.

Important! Port mirroring should only be enabled on a temporary basis as it causes high CPU utilization.

Link Layer Discovery Protocol (LLDP)

In addition to the Modular Chassis architecture of the E7-2 system support for topology discovery in a stacking ring, the Link Layer Discovery Protocol (LLDP) extends to all Ethernet interfaces on the E-Series.

The Link Layer Discovery Protocol (LLDP) (IEEE 802.1ab) passes PDUs across Ethernet interfaces every 30 seconds in order to announce neighbor and port details and capabilities between systems. LLDP is enabled by default on all GE Access interfaces configured as edge or trunk. ONT Ethernet and VDSL2 Access interfaces can be configured for LLDP, but are disabled by default.

LLDP defines a set of information to be transmitted and received periodically on an Ethernet interface to and from connected devices. There is no requirement that a peer will share the same objects as the sender. With LLDP, each interface maintains a set of MIB objects that it will transmit, and maintains a MIB per peer containing all the objects sent by those peers. This information can be leveraged by the management interfaces to build a “network” topology view and identify all connected access nodes. LLDP allows for any device to know the identity and some information about its immediate peers.

This topology can then be used by CMS to provision and troubleshoot network troubles, and by Compass for network and service planning, as well as network resource optimization.

You can view the LLDP details as follows:

- **Local** displays the details of the LLDP PDU sent by a local port.
- **Neighbor** displays the details of the cached LLDP information received by a port.

This reported information can be leveraged by the management interfaces to build a “network” topology view and identify all connected access nodes. See *Calix Management System (CMS) Guide* for instructions on using the Ethernet Topology Tool.

Calix LLDP PDUs contain the following TLV attributes:

- IEEE defined TLVs
 - Chassis ID – Contains the platform model and serial number of the chassis
 - Chassis ID subtype – Value of 7 indicates that the chassis ID is based on local assignment
 - Port ID – Contains the port name of the sending interface (ie 1/1/x1)
 - Port ID subtype – Value of 7 indicates that the Port ID is based on local assignment
 - Port Description - String indicating the ports description
 - Time To Live – How long the TLV should be cached
 - System Name – The provisioned system name
 - System Desc -- The provisioned System Location

- System Capabilities – MAC bridge
- Enabled Capabilities – MAC bridge
- Management Address - One or more management addresses
- Organizationally Specific TLVs
 - Link Agg Status – Indicates whether or not the port is a member of a LAG
 - Link Agg Port – If part of a LAG, contains the LAG port ID

With LLDP support, the Calix Management System (CMS) can now leverage this information to enhance its topology tool to do the following:

- Use the LLDP capability to automate the discovery of neighbors and autonomously construct the topology view
- Allow a link to be created between objects that reside in different network groups
When such a link is created, you can see a model of the remote network group when you double-click a network group object to open up that network group.
- Leverage CMS to identify networks as you do for C7, AE ONTs, and E7-2 Modular Chassis

Example

- a. The links are discovered for 3 nodes, and information is stored on the respective nodes.
- b. Extract the Database by running an inventory snapshot from CMS Scheduled task for the 3 nodes.
- c. With the inventory snapshot, CMS downloads the LLDP data stored on nodes into its own database and creates the topology links.

Note: If the topology links were already manually created for the devices, and then you perform device software upgrades to LLDP-aware releases, the LLDP-based topology links will co-exist with the manually-created links. Navigate to **Topology Links**, where the Link Type column shows which links are LLDP-Discovered and which links are manually created.

ERPS Topology View

With the ERPS topology view, you can determine the physical topology of Calix systems on an ERPS ring from any single node, and see which ports and systems are connected to each other around the ring. This is especially useful when you may not have access to the CMS to determine topology information. For example, when you are physically at a cabinet location.

Each element in this array provides information as follows:

- Port (or Local Port): An ERPS port on the local card.
- Local Neighbor Port: The other ERPS port on this node.
- Remote Neighbor Port: The ERPS port on the other end of the Local Port's link.

At each node on the ring, you configure whether to collect topology information about the ERPS domain. By default, this parameter is enabled. You can also issue the CLI command **set erps-domain <domain name> topology enabled** to provision the topology monitor for the entire domain. An example output follows:

```
e7-9>show erps-domain traffic topology
```

e7-9 31001000210 (master)			
Port	Status	Neighbor	
1/1/x1	forwarding	e7-20-25	1/21/x1
1/1/x3	blocking	e3-93	1/1/x4

e3-93 1969 (transit)			
Port	Status	Neighbor	
1/1/x3	forwarding	e7-9	1/1/x3
1/1/x4	blocking	e3-93	1/1/x3

e7-20-25 31103000804 (transit)			
Port	Status	Neighbor	
1/21/x1	forwarding	e7-9	1/1/x1
1/22/x1	forwarding	e7-20-25	1/22/x1

See "Creating an ERPS Domain" in the *Calix E-Series (E7 OS) Turn-Up and Transport Guide* for instructions and configuration guidelines on creating an ERPS domain.

In addition to gathering topology understanding, the structured CLI command and output will be used by TAC when they develop CLI scripts for maintenance activities.

Troubleshooting activities can occur from any location on the ring. As such it is necessary for the topology information to be available at each node on the ERPS ring, not just the ERPS Master Node. To that end, the show-erps topology command is one that has been designed to be run from any node on the ERPS ring of interest.

SNMP

The E-Series SNMP interfaces support Get and Trap commands as described in SNMP Overview to provide performance data and fault information, including:

- Host processor and memory utilization
- Ethernet port statistics
- Notifications:
 - Alarms
 - Events
 - Database changes
 - Security events
 - Threshold crossing alerts
- Current alarms and alarm counts

Note: The E-Series interfaces do not support provisioning via the SNMP interface.

Simple Network Management Protocol (SNMP) network management components:

- **SNMP Agent:** Configures the listen port for the E-Series SNMP agent, allowing fault management capabilities.
- **SNMP User:** Allows the management of SNMPv3 users, which are distinct from CLI users.
- **SNMP Community:** Allows the management of SNMPv2c communities. The E-Series can accept SNMP requests from each community provisioned in the system.
- **SNMP Trap Destinations:** Allows management of SNMP trap destinations, which are hosts specified to receive alarm and event notifications generated by the E-Series system.

Configuration guidelines

- The following limits exist for the SNMP request side:
 - 10 SNMPv2c community strings
 - 10 SNMPv3 users
- The following limits exist for the SNMP notification side:
 - 5 trap destinations that are either SNMPv2c or SNMPv3
 - SNMPv3 trap destinations must reference an SNMPv3 user on the request side
 - SNMPv2c trap destination community strings are independent of the request side community strings

Configuration task for an SNMPv2c management interface:

Create an SNMP community.

Configuration task for an SNMPv3 management interface:

Create an SNMPv3 user account.

Note: If the network is configured with CMS, the CMS server listens on an SNMP trap for any E-Series autonomous messages including alarms, events, security events, threshold events, and database-change events.

See the *Calix E-Series SNMP MIBs Reference*. Calix provides northbound interfaces for management integration to an Operations Support System (OSS) using Simple Network Management Protocol (SNMP). You can use an OSS to manage SNMP-based Calix network elements without significant interface development. This document provides enterprise MIB definitions for the Calix E-Series products.

SNMPv3 User Accounts

Simple-Network Management Protocol (SNMP) users are distinct from System (Web browser/CLI) users. This is because the SNMP subsystem utilizes its own authentication mechanism, in accordance with the security standards defined by SNMPv3.

SNMPv3 user access level

You can create SNMPv3 user accounts that allow the following levels of privileges:

No security privileges:

- Requests are only processed if they are unencrypted.
- Neither of the following permissions are provided:
 - Authenticate incoming packets
 - Generate packets authenticated by the receiving management entity

Authentication - no encryption privileges:

- Incoming messages are only processed if they can be authenticated.
- Outgoing messages can be authenticated by the receiving management system.

Authentication - encryption privileges:

- Incoming messages are only processed if they are encrypted and can be authenticated.
- Outgoing messages are encrypted and can be authenticated by the receiving management system.

Supported E-Series MIBS

This topic lists the supported E-Series enterprise and standard MIBs, and provides specific information regarding the implementation of the standard IF-MIB.

Enterprise MIBs

The following table lists the E-Series enterprise MIBs.

File Name ¹	Description
CALIX-PRODUCT-MIB	Defines the administrative assignments which are used to uniquely identify physical components of the Calix products.
CALIX-SMI	Defines the structure of management information for Calix Networks.
E7-Calix-MIB	Contains the objects required to display information for the system.
E7-Fault-MIB	Contains E-Series current alarm table and alarm counts. Note: Calix E-Series products support SNMP alarm traps for Valere rectifiers; alarm traps for other DC power sources are not supported.
E7-Notifications-MIB	Contains E-Series notifications.
E7-TC	Contains E-Series support definitions.

¹ Although the MIB file names include E7 only, these enterprise MIBs also address the complete set of E-Series (E7 OS) products.

Standard MIBs

The following table lists standard MIBs implemented in the E-Series products.

File Name	Description
SNMPv2 (mib-2)	Contains system group information.
SNMP-FRAMEWORK-MIB	Contains information for snmpEngine group.
SNMP-MPD-MIB	Contains snmpMPDStats information.
SNMP-USER-BASED-SM-MIB	Contains usmStats and usmUserTable information.
HOST-RESOURCES-MIB	Contains hrMemorySize, hrStorageTable, and hrProcessorTable information.

File Name	Description
IF-MIB (E7-2 only)	The ifTable and ifXTable. The following numbering scheme is used: • 101- 112 - SFP 1GigE ports, shelf 1, card 1, 10GE-4 and GE-12 cards • 101- 108 - SFP 1GigE ports, shelf 1, card 1, GPON-4 cards • 101- 124 - SFP 1GigE ports, shelf 1, card 1, GE-24 cards • 201- 212 - SFP 1GigE ports, shelf 1, card 2, 10GE-4 and GE-12 cards • 201- 208 - SFP 1GigE ports, shelf 1, card 2, GPON-4 cards • 201- 224 - SFP 1GigE ports, shelf 1, card 2, GE-24 cards • 301- 312 - SFP 1GigE ports, shelf 2, card 1, 10GE-4 and GE-12 cards • 301- 308 - SFP 1GigE ports, shelf 2, card 1, GPON-4 cards • 301- 324 - SFP 1GigE ports, shelf 2, card 1, GE-24 cards • ... • 2001- 2012 - SFP 1GigE ports, shelf 10, card 2, 10GE-4 and GE-12 cards • 2001- 2008 - SFP 1GigE ports, shelf 10, card 2, GPON-4 cards • 2001- 2024 - SFP 1GigE ports, shelf 10, card 2, GE-24 cards • 10101-10102 - XFP 10GigE ports, shelf 1, card 1, 10GE-4 and GPON-4 cards • 10201-10202 - XFP 10GigE ports, shelf 1, card 2, 10GE-4 and GPON-4 cards • 10103-10104 - SFP+ 10GigE ports, shelf 1, card 1 • 10203-10204 - SFP+ 10GigE ports, shelf 1, card 2 • ... • 11901-11902 - XFP 10GigE ports, shelf 10, card 1, 10GE-4 and GPON-4 cards • 12001-12002 - XFP 10GigE ports, shelf 10, card 2, 10GE-4 and GPON-4 cards • 11903-11904 - SFP+ 10GigE ports, shelf 10, card 1 • 12003-12004 - SFP+ 10GigE ports, shelf 10, card 2 • 20101-20148 - DSL ports, shelf 1, card 1, VDSL-48C and VDSL-48 cards • 20201-20248 - DSL ports, shelf 1, card 2, VDSL-48C and VDSL-48 cards • ... • 21901-21948 - DSL ports, shelf 10, card 1, VDSL-48C and VDSL-48 cards • 22001-22048 - DSL ports, shelf 10, card 2, VDSL-48C and VDSL-48 cards • ... • 30101-30104 - SFP Pon ports, shelf 1, card 1, GPON-4 cards • 30101-30108 - SFP Pon ports, shelf 1, card 1, GPON-8 cards • 30201-30204 - SFP Pon ports, shelf 1, card 2, GPON-4 cards • 30201-30208 - SFP Pon ports, shelf 1, card 2, GPON-8 cards • ... • 31901-31904 - SFP Pon ports, shelf 10, card 1, GPON-4 cards • 31901-31908 - SFP Pon ports, shelf 10, card 1, GPON-8 cards • 32001-32004 - SFP Pon ports, shelf 10, card 2, GPON-4 cards • 32001-32008 - SFP Pon ports, shelf 10, card 2, GPON-8 cards • ... • 100001-755360 - ONT FE, GE, and HPNA Ethernet ports

File Name	Description
IF-MIB (E7-20 only)	The ifTable and ifXTable. The following numbering scheme is used: • 101- 124 - SFP 1GigE ports, card 1, GE-24x cards • 201- 224 - SFP 1GigE ports, card 2, GE-24x cards • 301- 324 - SFP 1GigE ports, card 3, GE-24x cards • ... • 2101- 2104 - SFP 1GigE ports, scp a, SCP-10GE cards • 2201- 2204 - SFP 1GigE ports, scp b, SCP-10GE cards • 12101-12101 - SFP+ 10GigE ports, scp a, SCP-10GE cards • 12102-12102 - XFP 10GigE ports, scp a, SCP-10GE cards • 12201-12201 - SFP+ 10GigE ports, scp b, SCP-10GE cards • 12202-12202 - XFP 10GigE ports, scp b, SCP-10GE cards • 30101-30104 - SFP Pon ports, card 1, GPON-4 cards • 30201-30204 - SFP Pon ports, card 2, GPON-4 cards • ... • 100001-755360 - ONT FE, GE, and HPNA Ethernet ports
IF-MIB (E3-48C only)	The ifTable and ifXTable The following numbering scheme is used: • 101-112 - SFP 1GigE ports • 10103-10104 - SFP+ 10GigE ports • 20101-20148 - DSL ports
IF-MIB (E5-48C and E5-48 only)	• 101- 104 - SFP 1GigE ports • 20101-20148 - DSL ports

IF-MIB attributes: ifDesc and ifName

For Release R2.1 and lower, both the IF-MIB ifDesc and ifName attributes contain the interface description for all interface ports, and display the same value. For example:

```
IF-MIB::ifDescr.2101 = STRING: Ethport A/g1
IF-MIB::ifName.2101 = STRING: Ethport A/g1
```

```
IF-MIB::ifName.100001 = STRING: Ont 1/g1
IF-MIB::ifDescr.100001 = STRING: Ont 1/g1
```

For Release 2.2 and higher:

- The IF-MIB ifDesc attribute contains a value derived from the user provisionable "subscriber ID" field for all interface ports. If there is no subscriber ID configured, the ifDesc attribute will be blank.
- The ifName attribute contains the interface description.

Note: For additional information on enterprise MIBs, see the *Calix E7/E5-48/E5-48C/E3-48C SNMP MIBs Reference* manual, or download the MIB modules as described in [Downloading MIBs from the Web](#).

Loading the MIBs

Loading a MIB module provides a way for the NMS to learn about new MIB objects, including MIB names, OIDs, and the object type.

The loading order is important when you are loading MIBs. Many MIBs use definitions that are defined in other MIBs. These definitions are listed in the IMPORTS clause near the top of the MIB, and may need to be loaded first. If the load order is incorrect, the MIB loader may report that the imported MIBs are undefined.

Note: You can open a MIB file using WordPad.

This topic describes how to load the MIBs in the correct order.

To load the MIBS in the correct order

1. Load the MIBs in the following order:

Note: A given product may not support all listed MIBs.

- a. SNMPv2-SMI
- b. SNMPv2-TC
- c. SNMPv2-MIB
- d. RFC1213-MIB
- e. IF-MIB
- f. CALIX-SMI
- g. CALIX-PRODUCT-MIB

Note: If you are loading the v1 versions, *V1* displays in the file name.

2. In the Load TRAP-TYPE/NOTIFICATION-TYPE macro? confirmation dialog box, select **OK**.

The Load/Unload MIBs:SNMP dialog box displays.

Chapter 3

Network Topologies

This chapter provides engineering and planning guidelines for E-Series network topologies, including network transport options, node interconnections, and supported uplinks/downlinks.

Topics Covered

This chapter covers the following topics:

- Configurations and Deployments
- Ethernet Topology Protocols
- Ethernet Transport and Aggregation
- Operation Domains

E7 Configurations and Deployments

The E-Series supports a variety of product and network configurations:

- Deployments from central office (CO) and remote terminal (RT) locations
- 10GE Transport and aggregation using point-to-point and ring topologies
- Protected uplink interfaces
- GPON FT*TX (E7 only)
- Point-to-Point Ethernet FT*TX (E7 only)
- VDSL2 (E7-2 and E3-48C)

These core building blocks can construct many diverse business-specific network configurations:

- Mobile broadband transport
- Multipoint MEF-compliant transparent LAN service
- Aggregation of provider owned GE access devices
- Residential and business GPON FT*TP (E7 only)
- Residential and business point-to-point Ethernet FT*TP (E7 only)
- Residential VDSL2 services


The Ethernet transport and aggregation configurations are common to the E3-48C and E7 products. The two systems can be interchanged without restriction in applications that are not specific to the two-card E7 configuration.

The E7 system supports both high-density Central Office (CO) deployments as well as challenging remote terminal (RT) cabinet configurations.

Central Office (CO) FTTP configurations

The majority of GPON and point-to-point Ethernet FT*TP applications are from the central office. GPON has a reach of 20 km (approx. 60 Kft) with standard OIM optics, 40 km with an Extended Reach OIM; far longer than the 18 Kft reach of a typical copper loop. Ethernet single fiber optics support ranges up to 60km. As a result, a larger percentage of end users can be served directly from the central office with a fiber PON than with home-run copper.

Central office configurations typically require multiple E7 systems linked by a 10GE copper SFP+ cable to create a single Ethernet bridging and multicast domain, as shown in the following diagram. The Ethernet Ring Protection Switching (ERPS) ring links the multiple E7 systems, while RSTP and/or Link Aggregation are on the uplink to the network router(s).



GPON and Ethernet Deployment from a Central Office

Voice traffic can be switched to a co-located or remote Calix C7 VoIP TDM gateway with a GR-303 switch interface, or sent to a third-party SIP softswitch.

When deploying GPON, multiple E7 systems can be linked together with a Calix BITS Timing Cable to create a traceable network reference clock for the GPON and T1 services. A single redundant BITS connection in the CO provides an inter-CO network timing reference that can enable differential timing for T1 pseudowire clock recovery.


Remote Terminal (RT) FTTP configurations

The E7 is well suited to meet the requirements for RT deployments:

- Environmentally hardened, redundant architecture
- 1RU design is small yet can expand with an extra slot for very low first install cost
- Additional chassis are added with subscriber growth yielding a near linear cost curve
- Mix and match line cards in a common chassis meet the needs of a diverse remote access network
- Line cards are managed as a single chassis for operational efficiency
- Subscribers are aggregated and network resources are efficiently shared across protected trunk facilities
- Resilient, hot-swappable fan tray requires less power

- Integrated aggregation and transport reduces the number of systems that need to fit in the RT
- The E7 can be easily integrated into most existing outdoor cabinets to extend the reach of the FTTX service network, cost effectively

Remote terminal configurations typically involve multiple cabinets (ODC-100 to ODC-3000) connected by a ring (or star) to the central office. An E7 in the central office terminates the ring and provides the network uplinks, as shown in the following diagram.



GPON Deployment from a Remote Terminal

Transport between E7 systems is typically in a 10GE ring topology, but other topologies such as point-to-point, linear chain, or point-to-multipoint star may be used.

A remotely deployed E7 is able to support both IPTV and overlay RF video. Many Calix ODC cabinets accommodate EDFAs for video overlay.

E3-8G Configurations and Deployments

The E3-8G is intended for the MSO access space, where customers will leverage their existing cable plant. In most cases, RF is expected to be used with the E3-8G (not a requirement).

As the E3-8G is based on E7-2 and GPON-8 card, the deployment modes and north bound configurations are consistent across the platform. delivered and


In regard to the south bound connections to CPE, the E3-8G supports 8 GPON interfaces however due to space constraints in the E3-8G does not support Active Ethernet connections.

RFoG overlay using the E3-8G is accomplished in the same manner as the E7-20 and E7-2. However, as the E3-8G is a stand-alone unit, WDM couplers for the ingress of 1550nm downstream and egress of 1610nm upstream, must occur at one of the following:


- The splitter distribution cabinet
- Within a splice enclosure
- Within a node based EDFA and/or return path transmitter

Special consideration must be given to the return path of the RFoG devices in an E3-8G deployment if the distance from the headend will exceed the RFoG return loss budgets. Node based return path transmitters, such as those offered by Calix's partner Innotrans, are available to insure return path signal are managed in a manner to insure long range reach.


Use Case 1: voice and video services may remain in the RF network while HSI is over the GPON network.




Use Case 2: video services may remain in the RF network while HSI and voice are over the GPON network.





Use Case 3: RF video and RF return may be used through the ONT/GPON network.



RFoG Overlay with < 20km total network distance




E3-8G with RFoG > 20km without return transmitter**E3-8G with RFoG > 20km with return transmitter**

- Using RSTP, E3-48 r2 can interoperate with the C-Series and E-Series products as a member of the RSTP ring or as the RSTP root bridge. Generally, when adding an E3-48 r2 to an existing RSTP protected domain the E3-48 r2 would be a RSTP node. The E3-48 r2 also interoperates with the Calix B-Series products. In these typologies the E3-48 r2 will be configured to act as the RSTP root bridge for the protected nodes. In this topology the B6 is configured to transparently pass the control VLAN traffic back to the E3-48 r2.

The E3-48 r2 supports utilizing multiple links within a LAG group. All interfaces within one LAG group must be at the same transmission rate. A future release will support the ability to provide LAG across an RSTP ring, enabling higher bandwidth rings.

RSTP Rings

The E3-48C, E3-48/C r2 also support IEEE 802.1W Rapid Spanning Tree Protocol (RSTP) which can be used to provide a protected network connection in either ring or point-to-point topologies. Up to eight units may be interconnected in a RSTP ring.




In the case of point-to-point topologies, the E3-48C, E3-48/C r2 supports IEEE 802.3ad Link Aggregation protocol (LAG). Link aggregation provides a network-protected point-to-point and has the benefit of enabling multiple 1GE / 2.5GE / 10GE interfaces to the E3-48C ESAN.

- Using ERPS, E3-48C, E3-48/C r2 can be integrated into C-Series and E-Series deployments.
- Using RSTP, E3-48C, E3-48/C r2 can interoperate with the C-Series and E-Series products as a member of the RSTP ring or as the RSTP root bridge. Generally when adding these nodes to an existing RSTP protected domain, the unit would be a RSTP node. The E3-48C, E3-48/C r2 also interoperate with the Calix B-Series products. In these typologies, the nodes will be configured to act as the RSTP root bridge for the protected nodes. In this topology, the B6 is configured to transparently pass the control VLAN traffic back to the E3-48C, E3-48/C r2.

Daisy Chain

The E3-48C, E3-48/C r2 also support the network topology of a daisy chain. Up to eight units may be interconnected using 1GE / 2.5GE / 10GE interface(s) for uplinks, with the others used for downlink or shelf connectivity. The diagram below shows an E7 providing aggregation of nodes supporting a daisy chain of 4 units. The head end of the daisy chain can also be the E3-48C, E3-48/C r2 itself or any other Calix product.



E3-48C, E3-48/C r2 Linear Daisy Chaining

Ethernet Topology Protocols

The E-Series supports multiple network topology protocols, each with its benefits and reasons for use within the access network. This section describes these protocols and the capabilities of the E-Series platform with respect to each protocol.

Rapid Spanning Tree Protocol (RSTP)

The E-Series supports the IEEE 802.1w Rapid Spanning Tree Protocol (RSTP), which represents the evolution of the IEEE 802.1D Spanning Tree Protocol (STP). This protocol delivers improved network convergence time for point-to-point and small-scale ring protection as follows:

- Eliminates loops in a redundant-link network by selectively disabling links.
- Monitors for failure of active links and reactivates redundant links in order to restore the network to full connectivity


The E-Series also provides the ability for node protection between a pair of E-Series systems, using RSTP.

Note: If a device connected to the E-Series only supports STP, the E-Series will fall back to STP, ensuring seamless protocol interoperability.

RSTP allows a network design to include redundant network uplink / downlinks and provides a backup path to an Ethernet device if one of the active links fails. This is achieved without having to be concerned with network loops or manually enabling the backup links.

If allowed, bridge loops would allow traffic with the same MAC address to be seen on multiple ports, causing a networked device not to know where to send traffic and leading to a network broadcast storm. Broadcast storms can occur when a large volume of broadcast and multicast packets are forwarded in an endless loop between Ethernet switches if multiple paths exist. Broadcast storms consume CPU resources and can take a network or connected device down. Based on the design of the E-Series, bridge loops are NOT allowed to occur.

When two E-Series ports are configured for RSTP, one port is designated as blocking and the other designated as forwarding (active). For the E7, this behavior is true whether the ports are on the same E7 card or on different cards within a single shelf. This designation happens automatically as part of the RSTP protocol and does not require special provisioning. When connecting devices to the E-Series and using RSTP, the devices should have both of its ports connected to like ports on the E-Series. For instance, two GE or two 10GE ports should be used, not one GE and one 10GE port.



Industry Standard Platform / Port Protection Enabled with RSTP

The advantage of RSTP as an E-Series network redundancy protocol over link aggregation is the ability to protect rings of subtended devices, such as an E-Series aggregating multiple E5-100 units in an Calix ODC remote terminal.

Note: The E-Series RSTP does not interoperate with PVSTP (Cisco proprietary STP).

The max-age, hello-time, and fwd-delay values have a small range of valid values. The parameters are interdependent, which further constrains their values. Their relationship is as follows:

$$\text{max-age} \geq (2 * \text{fwd-delay} - 1) \text{ and } \text{max-age} \leq (2 * \text{hello-time} + 1)$$

$$\text{hello-time} \leq (\text{max-age} - 2) / 2$$

$$\text{fwd-delay} \geq (\text{max-age} + 2) / 2$$

Facility and equipment protection using RSTP

Individual Ethernet port interfaces with either a Trunk or Edge role can participate in Rapid Spanning Tree Protocol (RSTP).

- Facility protection can be configured by using two ports on the same node or E7 card.
- Equipment protection can be configured by using one port on two cards in the same E7 shelf. That is, one port on each E7 card.
- For Modular Chassis systems, RSTP is only supported on MCC shelf interfaces.


- For E7-20 systems, RSTP is only supported on SCP cards.
- For cross-card LAG interfaces, RSTP is not supported.

Note: Node protection can be configured by using two nodes on an ERPS ring.


To view the Ethernet port interfaces that are actively participating in RSTP, click the node in the Navigation Tree, and then click the **RSTP** > **Interfaces** tabs.

RSTP Node Protection

The need often arises to have diverse fiber routes from a remote network element back to the aggregation locations. This simple RSTP node protection configuration is shown in the following diagram.



The following figure shows RSTP node protection pair as the recommended method to construct the connection of two 10GE rings. See *Diverse Fiber Route and RSTP Node Protection* (on page [191](#)).



RSTP Planning Guidelines

Individual Ethernet port interfaces with either a Trunk or Edge role can participate in Rapid Spanning Tree Protocol (RSTP).

- Facility protection can be configured by using two ports on the same node or E7 card.
- Equipment protection can be configured by using one port on two cards in the same E7 shelf. That is, one port on each E7 card.
- For Modular Chassis systems, RSTP is only supported on MCC shelf interfaces.

- For E7-20 systems, RSTP is only supported on SCP cards.
- For cross-card LAG interfaces, RSTP is not supported.

Note: Node protection can be configured by using two nodes on an ERPS ring. See "Configuring RSTP Settings" in the *Calix E-Series (E7 OS) Turn-Up and Transport Guide* for the RSTP parameters to apply to the nodes.

- All network connections (ERPS, RSTP, LAG, DHCP servers, multicast routers, IGMP-enabled video servers, network-facing routers) are made on Trunk or Edge interfaces, only.
- When connecting devices to the E-Series using RSTP, connect both ports on the devices to like ports on the node. For example, two GE or two 10GE ports should be used, NOT one GE and one 10GE port.
- RSTP is enabled by default on all port interfaces.
- RSTP can be used in uplink, downlink, transport ring, and node protection pair configurations.
- The E-Series supports port-level RSTP. Therefore, ensure the far-end device is configured similarly (port-level RSTP) and not using VLAN-level RSTP.
- A device that is downlinked or uplinked to an ERPS ring must link either to a single node on the ring or to a node protection pair on the ring.
- BPDU Guard and RSTP cannot be enabled on an interface, simultaneously.
- For Modular Chassis (MC) systems, LAG and RSTP interfaces are only supported on the Modular Chassis Controller (MCC) shelf.
- For E7-20 systems, LAG and RSTP interfaces are only supported on the SCP cards.
- For an RSTP network, Calix recommends setting the following parameters as shown:
 - The VLAN IGMP Mode = snoop-suppress or proxy
 - The IGMP profile Router Learning Mode = static-dynamic
 - The IGMP profile Router Solicit On Topology Change = Y (enabled)
 - The interface NOT be designated as a static router port through the VLAN membership.
- RSTP node protection:
 - Only supported on the E7 platforms.
 - Subtended devices running RSTP connect to two different nodes.
 - One node is designated as the "Primary" and another node is designated as the "Secondary"
 - The Primary and Secondary nodes must be members of the same ERPS ring.

Ethernet Ring Protection Switching (ERPS)

The E-Series allows for provisioning with the recently approved ITU G.8032 Ethernet Ring Protection Switching (ERPS), a ring protection protocol. The ITU ERPS specification is an enhanced version of the Ethernet Automatic Protection Switching (EAPS) protocol defined in the IETF in RFC 3619. Other vendors have implemented similar types of ring recovery protocols, based on RFC 3619 or proprietary ring recovery protocols, which are not compatible with the Calix C-Series or E-Series products. Calix chose this implementation due to our standards-based focus, and emphasis in reaching full ITU-T compliance with a future software upgrade, pending industry / vendor adoption of the ITU G.8032 ERPS protocol. The initial implementation of G.8032 protocol on the E-Series (called Enhanced EAPS) is closely aligned with the ITU specification.

To create an ERPS domain, the following node aspects are designated:


- The role of the node
- A primary port and secondary port on the master node
- A first and second port on each transit node
- A control VLAN
- Service VLANs added to each node
- ERPS domain added to each VLAN membership

Number of Rings Supported

- The E7-2 supports 6 ERPS rings on a single shelf system or on the Modular Chassis Controller shelf.
- The E7-20 supports 2 ERPS rings on the SCP cards.
- The E3-48/C, E3-48/C r2, and E3-8G support 3 "chain of ERPS rings"

ERPS ring VLAN configuration

When an ERPS ring is provisioned, the VLANs to carry traffic are provisioned in the ERPS domain on each node of the ring. Multiple VLANs can be configured across the same physical ring for further separation of traffic. As shown in the following diagram, two logical VLANs are configured for voice and video traffic. The diagram also illustrates the control VLAN, which is provisioned as VLAN 2 in this example. All nodes on the ring send their traffic in the direction of the master node's primary port with the exception of the Control traffic on the ring.




ERPS Ring VLAN Configuration

ERPS Ring Normal Operation

With ERPS, one node in the ring is designated as the ring master and all other nodes are considered transit nodes. The master node is responsible for protecting the Ethernet ring from network loops as well as coordinating the flushing of bridging tables.

During normal operation, the following process occurs:

- The Master node forwards user traffic on its primary interface, and blocks/discards user traffic on its secondary interface.
- The Transit node, adjacent to the secondary interface of the master node, also blocks user traffic.
- The remaining Transit nodes forward traffic on both ring ports.
- The control VLAN is not blocked on the Master node secondary port, as it must be received by the Master node to verify the health of the ring. However, control VLAN packets received by the node master on the secondary ring port are never forwarded.



ERPS Ring Topology—Normal Operational Model


ERPS Ring Failure Detection

The ERPS protocol allows the following methods of detecting a ring failure, that are listed in the order of occurrence:

- a. Loss of Signal (LOS) detection specified in G.8032, occurs within one millisecond of a hardware-based fault detection, declaring a link failure message. The LOS mechanism detects the failures that account for nearly all ring faults, such as a fiber cut, transceiver failure, and node reset during software upgrade.
- b. Link Continuity Check Messages (LCCM) detection which is optional in G.8032, occurs with the loss of 6 consecutive messages (LCCM interval = 200 milliseconds), declaring a link down message. This method uses a point-to-point protocol that is link-local between each node on the ERPS ring. That is, the protocol generates Link Continuity Check (LLC) messages that are sent on the Control VLAN between two Ethernet ports of two connected E7's, where they are trapped at the CPU and are not forwarded to other ports. The LCCM mechanism detects infrequent failures, such as a port that is up but not able to pass frames because of a bad fiber connection, a CPU that is not able to process protocols, and a transponder-based WDM system not protecting links.

- c. ERPS Ring Heartbeat Ring-Level Status Messages detection occurs with the loss of 3 consecutive messages (message interval = 1-10 seconds), resulting in an alarm condition with no topology changes, unless other nodes on the ring report Node Failure of adjacent nodes. In that case, a Ring Fault is declared and traffic is forwarded on both the primary and secondary ERPS interfaces. This method uses health packets that are sent around the ring on the Control VLAN to determine if a ring failure has occurred.

The diagram below illustrates the change in traffic flow when a network link failure occurs.



ERPS Ring Topology—Network Link Failure Mode

Loss of Signal or Link Down Condition

The following process occurs when there is a LOS or Link Down condition on an ERPS ring:

1. A Transit node detects a LOS/Link Down condition on any of its ERPS interfaces, and then the following happens:
 - It immediately flushes its forwarding database and sends a “link down” control frame on the control VLAN
 - Other Transit nodes along the path intercept and forward the link down control frame
 - All Transit nodes flush their local forwarding database upon receipt of the link down control frame
2. The Master node receives the link-down control frame, and then does the following:
 - Flushes its bridge table
 - Instructs the Transit nodes to flush their bridge tables
 - Unblocks its secondary ERPS interface
3. All nodes start re-learning MAC addresses.
4. The Master node starts sending recovery PDUs out its primary ERPS interface on the control VLAN to determine when the ring has recovered.

Ring Fail Message Condition

If the Master node receives a node fail message from a ring node adjacent to a failed node, it declares a ring fault and begins forwarding traffic on both the primary and secondary ERPS interfaces, and then does the following:

1. Flushes its bridging table for the ring-based VLANs, except the control VLAN
2. Begins forwarding traffic on both the primary and secondary interfaces for the ring-based VLANs
3. Sends a LINK_DOWN PDU out both the primary and secondary ports, instructing the transit nodes to flush their local bridging tables, and then begins learning the new topology
4. Sends out Recovery messages to determine when the ring is restored

Recovery Sequence

After receiving six consecutive recovery PDUs (message interval = 1-10 seconds, default of 1 second) on its secondary ERPS interface, the Master node declares the ring has recovered, and then it initiates a switchback, as follows:


1. The Master node again flushes its bridge table and instructs the transit nodes to flush their bridge tables.
2. The Master node starts blocking user traffic on its secondary interface.
3. The Master node sends health PDUs out its primary ERPS interface at the designated interval (message interval = 1-10 seconds, default of 5 seconds).

The total time for a ring to revert after a recovery is declared:

$$6 \text{ RECOVERY tries} * \text{RECOVERY message frequency}$$

ERPS ring device protection

The E-Series is extremely flexible and provides the ability to interconnect with multiple protection protocols to multiple devices at any given node in the ring. As shown in the following diagram, device protection can be accomplished using the redundant card/unit architecture, while network protection is accomplished using ERPS, RSTP, or link aggregation.



Supported Link and Node Protection Topologies

ERPS Planning Guidelines

- Each ERPS domain must have the following:
 - A unique domain name assigned to all nodes in the domain.
 - A unique control VLAN ID designated on each ring node.
 - One node designated as the master node and the remaining nodes designated as transit nodes.
- A primary and secondary port must be designated on the master node.
- A first and second port must be designated on each transit node.
- A single node can be the Master Node for all, some, or none of the ERPS domains on the system.
- Each Ethernet interface can only have one ERPS domain associated with it. Traffic between rings is passed over the backplane between the two E7-2 cards or the E7-20 SCP cards.
- For the ERPS Control VLAN, use a VLAN ID that is NOT used for any services on the network.
- The ERPS ring Control VLAN is limited to passing the ERPS control PDU's.
- All node ports in an ERPS ring must be the same speed (all GE, all 10GE, or all 2.5GE).
- The Ethernet port interface role for ERPS ports is limited to "trunk."
- Ethernet interfaces used for ERPS links cannot be set to "Untrusted."
- Before adding an interface to an ERPS ring, the associated port must have the attributes set as follows:
 - Duplex = **full**
 - Flow Control = **none**

If the interface is removed from the ERPS domain, these attributes can again be modified.

- If you delete an ERPS domain from a node without first disabling the ports configured with the ERPS domain, a forwarding loop may potentially be introduced to the system. Calix recommends disabling the Ethernet ports containing the ERPS domain before deleting the ERPS domain from the system.
- To change any attributes on an ERPS domain, it must be disabled.
- The ERPS feature is compatible with all other Calix products running ERPS and can be mixed within the ring.
- If an E7-2 shelf with a single card is configured in an ERPS ring, the designated ERPS ports for that node will reside on the single card. However, if a second card is subsequently installed in the E7-2 shelf, one ERPS link for the node must be relocated to the second card and reconfigured such that each card has a designated ERPS port.
- An ERPS ring can support up to 32 E-Series units, or up to 16 E7 units (for E7-2 MC nodes with VDSL2-based MCCs). Each E7 card consists of one unit. An E7 chassis containing two cards counts as two units.

Note: The number of units does not include devices or nodes subtended from the ring, including the MCE shelves of an E7-2 modular chassis (MC) node.


- A “chain of ERPS rings” should be limited to no more than 3 interconnected rings.
- Only modify the reserved VLAN of the system during a maintenance window, as the action is service affecting to the ERPS ring.
- All VLANs must be created on each node of the ERPS ring and be associated with the ERPS domain membership. (Do NOT associate VLANs to underlying ERPS port interfaces.)
- VLANs can be members of multiple ERPS rings. Traffic switched between rings and network designs must be reviewed to ensure forwarding loops are not introduced.
- The node periodically verifies that the ERPS control VLAN is still present in the forwarding table of the switch component. If the VLAN is not present in the hardware, an alarm is raised.
- A Health message is transmitted on the Master Node primary interface at an interval of 1-10 seconds (user defined). If the Master Node does not receive 3 consecutive Health messages on the secondary interface, it will generate an alarm, but does not automatically start forward traffic on the normally blocked secondary port.
- A Recovery message is sent from the Master Node at an interval of 1-10 seconds (user defined) after a fault is detected on the ring to determine when the ring has recovered. The Master Node must receive 6 consecutive Recovery messages to declare the ring as recovered, initiating a switchback.
- The total time for a ring to revert when the ring has recovered:
6 RECOVERY tries * RECOVERY message frequency
- All nodes in an ERPS ring must have the same IGMP settings on the video VLAN for traffic to flow on the ERPS ring—either all with snooping (snoop-suppress, proxy) or all without snooping (flood).

- For nodes in an ERPS ring with no directly attached routers, set the IGMP profile parameter Router Learning Mode to Static-and-Dynamic and have no statically defined router ports set. This ensures that only the ring is “learned” and allowed to be the multicast router source for the system. This configuration prevents any subtended devices that might incorrectly flood a general query into the network (typically during a topology change) from being learned as the querying router.
- For nodes with a direct connection to a multicast router that does not respond to Multicast Router Solicitations, configure statically-defined router ports to identify the uplink. See Creating VLAN Members.
- When a node is configured for RSTP Node Protection, the RSTP domain of the two nodes must be associated with one of the ERPS domains.
- If an ERPS ring is configured with a VLAN that uses DHCP snooping and IPv6/LDRA, then the DHCP option 82 must be configured the same on all nodes in the ERPS ring.

G.8032v2 Ethernet Ring Protection Switching

The E7 allows Service Providers to implement network reliability in several ways, including LAG, RSTP, and Ethernet Ring Protection Switching (ERPS). Release 2.4 introduces an additional option based on ITU-T G.8032v2 (ITU-T G.8032/Y.1344 02/2012). The implementation of G.8032 is designed for multi-vendor interoperability, and supports rapid switch times comparable with SONET protection switch times. G.8032 is supported by the Calix E7, E3-8G, E3-48C, E5-48, E5-48C, E5-520 and E5-308 products. This allows support for unified access (both residential and business services) on a single ring. The E7 supports multiple rings, providing Service Providers the flexibility to isolate a group of customers or services into separate logical or physical rings to improve reliability and network operations.

Note: Release 2.4 does not support G.8032 over LAG interfaces, however, this support is planned for a future release.



G.8032v2 Functional Overview

The fundamentals of the G.8032 ring protection switching protocol/architecture are as follows:

- The principle of loop avoidance
- The utilization of learning, forwarding, and Filtering Database (FDB) control (learn and flush) to direct frames correctly around the ring path


The overarching fundamental concept of this ring protection protocol is that once a ring port has been blocked, it may be unblocked only if it is known that there remains at least one other blocked ring port somewhere else on the Ethernet ring.

The nodes in a G.8032v2 ring are identified as follows:

- **RPL owner node**—the node with the designated Ring Protection Link (RPL) port that is responsible for blocking Ethernet data frames at one end of the RPL to prevent loops from occurring.
- **RPL neighbor node**—the node adjacent to the RPL owner node where the connecting port is optionally designated as the RPL Neighbor port that is responsible for blocking its end of the RPL.

Note: All other nodes in the ring are identified as RPL Mode "none".

In the normal resolved protocol state, at the RPL Owner and RPL Neighbor ports the Ethernet data frames are blocked at the RPL port for purposes of establishing the default flow of frames. It is best practice for all ports on a ring to run at the same line rate. However, there is no restriction on allowing the end user to configure asymmetrical line rate ports.



As part of configuring ring instances, a control VLAN is created on each node of the ring. This control VLAN is used to pass ring APS (R-APS) packets between all the nodes on the ring. The R-APS channel carries the protocol frames used to manage the protocol/blocking state of the ring.

As a part of the node configuration when VLANs are created, a VLAN can be associated with a configured ring instance. This is needed for nodes to identify which VLANs to block on the ring ports, and to understand which VLAN forwarding tables must be flushed in the event of a failure in the ring.

When configured as the RPL Owner node in a ring, the RPL Owner node blocks the forwarding of data VLANs on the RPL port when the ring is in the normal state. The control/R-APS VLAN is not blocked on the RPL port, however R-APS packets received on the RPL port are never forwarded.

G.8032 Protection Operation

The RPL Owner node may declare a Signal Fail (SF) and initiate a ring switch event for the following reasons.

- A fiber cut
- Failure to receive consecutive timely or valid CCMs from MEPs
- Receipt of a R-APS PDU, indicating Signal Fail (SF) message from any node on the ring


In the event of a detected or reported SF the RPL Owner node forwards the R-APS with appropriate Flush indication out both ring ports, unblocks the RPL port and flushes the local forwarding database for VLANs who are members of the ring.

Non-RPL Owner nodes may report SF for the following reasons:

- Locally attached fiber failure (i.e. a fiber cut)
- Locally reported Loss of Continuity

In the event of a detected failure, the node generates an R-APS out its non-failed ring port with SF condition set. Upon reception of the SF R-APS PDU at the RPL Owner node, the RPL Owner node will initiate a ring switch.

During the recovery phase (no local SF at the node(s) that had reported the failure), when the failed link gets restored the nodes adjacent to the restored link send R-APS No Request (NR) PDUs. Upon receiving the NR indication, if the ring is configured for revertive switching the RPL owner will start the revert hold off timer (configurable for 1 to 12 minutes). Assuming the NR condition remains in effect; when the timer expires the ring will revert. The RPL Owner node blocks the RPL port and then sends an R-APS No Request/RPL Blocked (NR-RB) PDU on both ring ports. This will cause all other nodes other than RPL Owner on the ring to unblock all blocked ports and initiate a flush as dictated by the Flush logic.



The G.8032 protocol maintains a state machine to identify the state of the ring. Ring nodes may be in one of the following states:

- **Disabled (init)** – Ring protocol is administratively down
- **Idle** – Normal Operation with no ring port/link/node faults detected on the ring. Physical topology has all nodes connected in the ring and the logical topology has all nodes connect without a loop. The RPL is blocking.
- **Protection** – Ring port/link/node faults are detected by nodes adjacent to the failure. The nodes adjacent to the failure block the failed ring port and report the failure using R-APS SF PDUs. The RPL owner has switched the ring to its protected state.
- **Manual Switch (MS)** – The operator has tried to assert a Manual switch at a ring port. The ring will be switched if SF/FS is not present elsewhere on the ring.
- **Force Switch (FS)** – The operator has asserted a Forced switch at a ring port. The ring will switch. If SF is present, a segmented ring may result.
- **Pending** – Transient state – reported during transitions between the states listed above (typically during a timeout period of one of the protocol timers where the RPL owner node is in the process of reverting the ring to the Idle state).

G.8032 Administrative Commands

G.8032 provides a set of administrative commands that induce and manage ring switching. The operator may assert a Manual Switch (MS) or a Forced Switch (FS) condition at any node. They may further assert the CLEAR command to clear the condition at the node.

The functions behave as follows:

- **Forced switch:** Forces a topology change and a change to the ring port where the command is issued, even if there is an existing Signal Fail condition elsewhere in the ring.
 - Allows an operator to block a particular ring port
 - Effective even if there is an existing Signal Fail condition
 - Multiple FS commands for ring supported
 - May be used to allow immediate maintenance operations
- **Manual switch:** Causes a topology change and a change to the ring port where the command is issued, only if the change does not result in a path failure elsewhere on the ring.
 - Allows an operator to block a particular ring port
 - Effective only if there are no existing Forced Switch or Signal Fail conditions
 - Overridden by new Forced Switch or Signal Fail conditions
 - Multiple Manual Switch commands cancel all Manual Switch commands
- **Clear:** Removes an existing forced or manual switch command on the ring port. The Clear command is also used at the RPL owner to clear a non-revertive mode condition and move the block to the RPL.


G.8032 Revertive Support

G.8032 has the option to operate in either revertive or non-revertive mode. Configure the revertive/non-revertive mode at the RPL owner node. Note that as the RPL owner node controls the revert behavior, it does not matter what the revert configuration is at other nodes. When in revertive mode, the RPL port will be the link that is normally blocking in the absence of any failures. In the face of a clear from failure, all other ports are non-blocking after the revert timeout. This is the default operational mode.

When in non-revertive mode, the last failed link will be blocked. Conceptually the blocked ring port follows the last failed link on the ring. Revertive mode offers more predictability as the operator can control which ring port is the normally blocked port.

G.8032 Supported Topologies

The E7 supports up to six G.8032 rings per node. Each G.8032 ring must have its own Ethernet ring ports. An example of this topology is shown in the following diagram:



Note: The E7 does not support sub-tended G.8032 rings in the R2.4 release.

Configuration process

1. Set the Ethernet ports and interfaces.
2. Configure the G.8032 ring.
3. Configure the RPL modes for the Owner and Neighbor ring ports.
4. Add the traffic VLANs to the G.8032 ring.
5. (Optional) Configure CCMs.
6. Enable the Ethernet Ports.

Configuration guidelines


Follow these guidelines when configuring an Ethernet transport ring with G.8032v2:

- One logical ring instance is supported per physical ring.
- One G.8032v2 physical ring supports up to 32 E-Series units, or up to 16 E7 units (for E7-2 MC nodes with VDSL2-based MCCs). Each E7 card consists of one unit. An E7 chassis containing two cards counts as two units.
- Some devices do not send to the multicast address as calculated by using the Ring ID but send to the G.8032 Version 1 default MAC of 01-19-A7-00-00-01. Calix strictly adheres to the standard and will send the R-APS packets for each Ring ID to the multicast address 01-19-A7-00-00-[Ring ID]. If the addresses do not match, the G.8032 protocol will fail as messages are not properly interpreted by RPL owner and neighbor for each Ring ID. The destination address may be configurable on the 3rd party equipment as a workaround. Calix is **not** configurable.
- You can map a ring instance to any pair of GE or 10GE Ethernet ports in the ring.
- You can map a ring instance to LAG interfaces where the members of each LAG interface are on the same line card.

- You can mix GE, 2.5G and 10GE line rate in a ring to allow for migrating from a GE ring to a 10GE ring.
- Configure the revertive/non-revertive mode at the RPL owner node. Note that as the RPL owner node controls the revert behavior, it does not matter what the revert configuration is at other nodes.
- Disable/Enable ring on nodes without owner/neighbor will not form into loop. (All nodes' ring ports are unblocked)
- Adding SOAM (CFM) to an existing ring is service affecting; the following procedures cover how to configure SOAM as part of the G.8032v2 ring configuration process.

Fault monitoring

An E-Series G.8032v2 ring supports Ethernet OAM Continuity Check Messages (CCMs) for monitoring link connectivity by configuring maintenance group end points (MEPs) on the individual ring ports. At a specified interval, MEPs send CCM protocol data unit (PDU) messages to check continuity on the ring. The following graphic shows an *example* of a simple four node main 10GE ring configured for CCM. Notice for each side of the link between adjacent nodes on the ring, the same MEG, MEP, and Remote MEP values are assigned.



Note: Configuring Continuity Check Messages (CCM) on an existing ring (rather than during the initial ring configuration) is service-affecting. See "Configuring Ethernet OAM" in the *Calix E-Series Maintenance and Troubleshooting Guide* for complete instructions on configuring Ethernet OAM.

Configuration process

Configuring the Ethernet OAM functionality and initiating various actions and data retrievals from the Ethernet network consists of the following process:

1. Create a Maintenance Entity Group (MEG).
2. Add a Maintenance Endpoint (MEPs) to the MEG.
3. Since the MEG auto-discovery parameter is not supported on G.8032v2 rings, add Remote Maintenance Endpoints to the MEG to avoid alarms caused by detecting unknown MEPs.

The continuity checks begin immediately after the MEP is created, if the Continuity Check parameter is enabled.

Configuration guidelines

- G.8032v2 ring port health can be provisioned to rely on Continuity Check Message (CCM) operation.
- MEG name and local and remote MEPs must be assigned on a ring port in order to enable CCM.
- The MEG must be in Y.1731 format, the VLAN must be the same as the ring control VLAN, and the MEG level should be 0, 1, or 2.
- Per shelf, the MEG name on a G.8032v2 ring port must be different than the MEG name for other ring port's.
- The MEG used for G.8032v2 ring port objects must have the following:
 - One and only one local MEP
 - One and only one remote MEP

Note: If you attempt to add more than one local or remote MEP, a "CCM LOC" error occurs.
- A MEP is provisioned against a G.8032v2 ring port object.
- MEP direction must be down and the continuity check must be enabled.
- A MEP and a Remote MEP are required.
- Per port, the local MEP and remote MEP must be different.
- For each side of the link between adjacent nodes on the ring, use the same MEG, MEP, and Remote MEP values.
- The parameters that can be configured on local MEP are "Admin State (enable/disable)" and "CCM LTM Priority (0-7)."

CLI Example: `set meg <megname> mep g8032-ring-interface <ring/1> admin-state <enable/disable> ccm-ltm-priority <0-7>`

Note: When the Admin State is disabled, the G.8032v2 ring does not consider CCM as configured.

- Only Ethernet OAM Continuity Check Messages (CCMs) are supported on G.8032v2 ring ports. Link trace and Loopback are not supported.
- The following MEG attribute is not supported for G.8032v2 rings:
 - Auto discovery (must always be set to disabled)
- The following MEP attributes are not supported for G.8032v2 rings:
 - Port/IP Host
 - Frame Measurement Profile
 - Delay Loss MAC
 - Delay Measurement
 - Loss Measurement

G.8032 Interoperability

Although G.8032v2 is an ITU standard designed to promote interoperability, different vendor implementations can result in changes in behavior. For Release 2.4, the following third-party equipment has been tested for interoperability: equipment has been tested for interoperability:

- Alcatel-Lucent ALU 7750
- Cisco ASR 903
- Juniper MX80

Link Aggregation Group (LAG)

The system uses IEEE 802.3ad/802.1AX Link Aggregation (LAG) to bond multiple GE or 10GE or 2.5GE ports (not mixed) into a single link aggregation group with a single logical Ethernet interface. The ports that comprise a LAG can be on the same line card or on two different line cards in the E7 same shelf, but, must be of the same module type. For example, either all XFP or all SFP+.

Link aggregation provides two principle values to the network operator:

- Point to point link protection / redundancy between network elements
- Bandwidth expansion of the logical interface beyond the capacity of a single link

You can configure the following LAG configurations:

- The **active LAG** has all ports in the LAG active and the combined bandwidth of the ports is available to carry the traffic on the logical link of the LAG. When a failure occurs, the available bandwidth is reduced by the amount of bandwidth carried over the failed port.

For an E7 cross-card LAG, this configuration is known as **active/active**.

- The **active/standby LAG** has multiple ports active and one or more port in standby, positioned to come online when an active port fails.

For an E7 cross-card LAG, a failure threshold defines how many active ports in the group can fail before the group is taken out of service, and the system switches the LAG operation to the ports on the standby card.

Example Single-Card LAGs


Two different single-card LAG configurations are shown below:

- The active LAG configured for the E-Series has all ports in the LAG active and the combined bandwidth of the ports is available to carry the traffic on the logical link of the LAG. When a failure occurs, the available bandwidth is reduced by the amount of bandwidth carried over the failed port.
- The active/standby LAG configured for the router has multiple ports active and one port in standby, positioned to come online when an active port fails.

Example parameter settings:

- RSTP Enabled = disabled
- LAG cross-card = N (No) or disabled
- LACP Min Ports = 1
- LACP Max Ports = 2**

Note: **Ports added to the LAG that exceed the LACP Max Ports parameter are designated as standby ports.



Example Cross-Card LAGs


The E7 cross-card protection LAGs shown below have half of the LAG ports on one card and the other half of the LAG ports on another card, in the same shelf. You can designate the links on either card as active or standby.

Note: If the number of active ports falls below the Min Ports value, then the system switches the LAG operation to the ports on the standby card. When using cross-card LAG, the LAG on the switch connected to the E7 must have RSTP disabled.

Note: Keep the LACP role set to the default value of **active**, unless there is a very clear need. Specifically, the topology where at least one side of the LAG is cross-card will not operate when either side of the LAG is set to LACP role = passive.

Example parameter settings:

- RSTP Enabled = disabled
- Access router LAG:
 - Role = Trunk
 - LACP Cross Card = active-active
 - LACP Min Ports = 1
 - LACP Max Ports = 4
 - Card 1 Ethernet ports LACP Priority = 128
 - Card 2 Ethernet ports LACP Priority = 100
- C7 LAG:
 - LACP Cross Card = active-standby
 - LACP Min Ports = 1
 - LACP Max Ports = 1
 - Card 1 Ethernet port LACP Priority = 100
 - Card 2 Ethernet port LACP Priority = 128




LAG per card and RSTP

As a network configuration, RSTP can be enabled on two LAG interfaces to create a protected connection to routers. As shown below, a separate LAG is configured on each card and the LAG interfaces have an aggregated bandwidth capacity of 2Gbps, while RSTP provides E7 cross-card equipment protection. All of the links in the LAG must terminate on the same card for RSTP to be supported.

Example parameter settings:

- RSTP Enabled = enabled
- LAG cross-card = N (No) or disabled
- Access router LAG:
 - LACP Min Ports = 2
 - LACP Max Ports = 2



Configuration guidelines

- **LAG interfaces:**
 - The ports in a LAG must be either all GE or all 10GE.
 - For GE ports:
 - ♦ up to 8 GE ports per LAG
 - ♦ up to 6 LAGs per shelf
 - For 10GE ports:
 - ♦ up to 4 10GE ports per LAG
 - ♦ up to 2 LAGs per shelf
 - For 1GE LAG groups, the connectors should be of the same type: SFP or SFP+, not both.
 - All VLAN associations must be removed from a LAG before it can be deleted.
 - Before a port can be assigned to a LAG interface, the port's default associated interface must be disabled. When a port is assigned to a LAG, the LAG interface provisioning (for example, VLAN membership) applies to the port. See *Configuring an Ethernet Port* for details on how to add ports to the LAG interface.
 - For Modular Chassis (MC) systems, LAG and RSTP interfaces are only supported on the Modular Chassis Controller (MCC) shelf.
 - For E7-20 systems, LAG and RSTP interfaces are only supported on the SCP cards.
 - A LAG can include ports on the same line card, or on two different line cards within the same shelf, known as cross-card protection.
 - LACP System Priority is used between two systems connected by the LAG to determine which system should be controlling the LAG. The lower value takes priority. Typically, the upstream side of the LAG is configured for the LAG master (lower value).
 - The port priorities on each side of the LAG should be set to the same values.
 - Keep the interface LACP Role set to the default value of **active**, unless there is a very clear need. Specifically, the topology where at least one side of the LAG is cross-card will not operate when either side of the LAG is set to LACP role = passive.
 - Keep the Ethernet port parameter LACP Timeout set to the default value of **short**, unless there is a very clear need. Specifically, the topology where at least one side of the LAG is cross-card will not operate when either side of the LAG is set to LACP Timeout = long.
 - Set the following parameters as shown to achieve the best re-convergence time:
 - ♦ The VLAN IGMP Mode = proxy
 - ♦ The IGMP profile Router Learning Mode = static
 - ♦ The IGMP profile Router Solicit On Topology Change = N (disabled)
 - ♦ The LAG interface is designated as a static router port through the VLAN membership.

- For a single-card LAG, the ports added to the LAG that exceed the LACP Max Ports parameter value are designated as standby ports and come online when a LAG active port fails. (This is the equivalent of active/standby LAG on a single card.)
- The Link Aggregation Control Protocol (LACP) assigns a LAG the same MAC address of the associated Ethernet link with the highest port priority. When two or more ports have the same priority, the MAC address of the port with the lowest port number is used. Therefore, if ports 3, 5, and 7 are configured for a LAG, then the LAG MAC address is the MAC address of port 3. The system MAC address identifies the system in the LACP control messages, and the interface MAC address identifies the Link Aggregation Group.
- RSTP must be disabled in order to configure a cross-card LAG.
- For Active/Active cross-card protection LAGS, the interface role parameter must be set to Trunk.
- For Active/Standby cross-card protection LAGs:
 - ♦ An equal number of links should be configured for the active card and the standby card (1-4).
 - ♦ The ports on only one of the cards will be active while the ports on the other card are standby.
 - ♦ When provisioning active/standby LAG, the port priority should NOT be provisioned to 0.
 - ♦ For both ends of a LAG, the ports on the active card must all have the same LACP Priority value for the Ethernet port parameter, and the LAG ports on the standby card must all have the same priority value. Yet, the priority value on the active card ports must be lower than the value set for the ports on the standby card, giving the priority to the active card ports.
 - ♦ A failure threshold defines how many active ports in the group can fail before the group is taken out of service, and the system switches the LAG operation to the ports on the standby card. This threshold is indicated by the LACP Min Ports parameters.
 - ♦ LACP Min Ports = LACP Max Ports (When one active port fails, the system switches the LAG operation to the ports on the standby card.)

Ethernet Transport and Aggregation

The E-Series provides flexible, high-capacity Ethernet transport and network aggregation options to service providers. Also, Metro Ethernet Forum (MEF) compliant services (MEF 9 and 14) are supported for business service delivery.

Integrated High-Capacity GE Aggregation

The E7 is built on a core Layer 2/3 switch capable of full-duplex, wire speed forwarding at all frame sizes and traffic types across all interfaces. This capacity makes the E7 ideal for aggregation and transport of IP/Ethernet services across the access network.

Ethernet Ring Topologies

Ring topology networks are often used to ensure carrier-grade service delivery because rings enable dynamically protected network service delivery to multiple locations.

Ethernet Ring Protection Switching (ERPS)

The E-Series ring-based protection protocol for Ethernet networks is aligned with the ITU G.8032 Ethernet Ring Protection Switching (ERPS) standard. The ERPS ring can operate across a single GE interface or 10GE interface, depending on network requirements. The E-Series system supports a single ERPS ring per interface. The primary application for multiple ERPS rings is to enable the E-Series to interconnect ERPS rings. See *Ethernet Ring Protection Switching (ERPS)* (on page [167](#)) for more information.

Rapid Spanning Tree Protocol (RSTP) ring


The IEEE 802.1w Rapid Spanning Tree Protocol (RSTP) is supported on most switches and routers to prevent network loops and reroute traffic in traditional Ethernet enterprise networks. RSTP can also provide service restoration when two or more links share common end points in a point-to-point or ring topology. When utilized for small-scale ring topologies or point-to-point configurations, RSTP converges and provides service restoration within several hundred milliseconds, depending on the number of Ethernet switches in the network. RSTP is therefore not recommended for large-scale ring topologies. See *Rapid Spanning Tree Protocol (RSTP)* (on page [163](#)) for more information.

Multiple-ring configuration

The E7 with redundant cards is ideally suited to support subtended rings from the main transport ring. Each ring is terminated on a separate line card in a single E7 shelf and the E7 backplane becomes the alternate protection path between rings. ERPS is used for the main transport ring, typically a 10GE ring, and multiple small RSTP rings can be subtended, as shown in the following diagram.

The E-Series is an ideal platform to build high speed Ethernet aggregation networks in the access network. In the figure below, a 10GE ERPS aggregation ring is built to link five E-Series locations. The E-Series at right and bottom of each figure have subtended GE ERPS rings, while the E-Series on the left are configured for RSTP Node Protection and used to “dual-home” connect a 1GE RSTP ring to the 10GE aggregation ring.


E7 Multiple Rings




Note: A chain of ERPS rings should be limited to no more than three (3) interconnecting.

The number of ERPS rings supported on each E7 system is as follows:

- The E7-2 supports 6 ERPS rings on a single shelf system or on the Modular Chassis Controller shelf.




- The E7-20 supports 2 ERPS rings on the SCP cards.



- The E3-48/C, E3-48/C r2, E5-48, and E5-48C supports 3 "chain of ERPS rings".

Connecting transport rings


The following figure shows the recommended method to construct the connection of two 10GE rings with either E7-2s. RSTP Node Protection is configured on the two E-Series that link the two 10GE rings and one of the rings uses RSTP for ring-based protection. Only a single link exists between the two E-Series nodes interconnecting the rings. See *Diverse Fiber Route and RSTP Node Protection* (on page [191](#)).



2.5GE Subtended Ring Topologies with E3-48/C r2

Standard E7 topology rules apply to E3-48/C r2 units. Max 32 units (hops) per ring. Deployment scenarios include a 10GE ring where each E3 is a member of the ring as well as a subtended 2.5Gbps nodes in a FTTN deployment.

The SFP and SFP+ ports support the use of the copper direct attach (DA) cables available from Calix. Direct Attach cables are used to interconnect E3-48 r2 nodes in a 2.5GE ring located in a Cooled cabinet.



SFP Usage Constraints with E3-48 r2 Port Combinations	
Port Combinations	Support Description
2-10GE SFP+ 2-1GE SFP Fiber	Requires the use of Direct Attach (DA) cables on the SFP ports (1Gbps).

Diverse Fiber Route and RSTP Node Protection


Two E7-2 systems can be provisioned to provide route-diverse network redundancy using RSTP domains in the following ways:

- Aggregate Ethernet devices
- Create an RSTP protected uplink from multiple locations or E-Series shelves
- Create isolated RSTP domains that are connected by the redundant ERPS / Enhanced EAPS transport ring

RSTP domain isolation performs the following functions:

- Localizes topology awareness
- Prevents traffic from transit across the transport ring when not required
- Reduces failover times by reducing the number of RSTP convergence points


In order to provide protection of subtended devices that have a fiber connection to two locations on an ERPS ring, E7s can be configured for RSTP Node Protection, as shown below.



Ethernet Aggregation with Diverse Fiber Routes


When RSTP Node Protection is configured, the two E-Series nodes become part of the same RSTP domain. The RSTP logic within the node protected pair recognizes the ERPS link between the two E7 nodes, and will never to block this link when converging the RSTP topology.

Node protection used for Uplink




Node protection flexibility

With E-Series node protection, the ERPS ring and protection switches are managed independently of the RSTP domain. As a result, topology changes or failures within the RSTP domain are not propagated to an ERPS ring, and vice versa. Also, RSTP topology changes in one protected node pair are not propagated to other protected nodes. Multiple node protection RSTP domains are depicted in the following diagram.



Multiple RSTP Node Protection Domains for Uplinks / Aggregation

Two E-Series systems forming a protection domain are typically located adjacent on the ERPS ring, but they are not required to be physically co-located or immediately adjacent. The following diagram illustrates the flexibility of RSTP node protection and how it can be applied within the transport ring. The network design below is supported for adjacent and non-adjacent Node Protection Pairs.




Logical RSTP Node Protection Domains within an ERPS / Enhanced EAPS Ring

The E7 RSTP node protection features are identical where the units can be combined to form a protected node configuration.

Node protection pair to connect transport rings

The following diagram shows the recommended method to construct the connection of the two 10GE rings. RSTP Node protection is configured on the two E7s that link the two 10GE rings and one of the rings uses RSTP for ring-based protection. only a single link exists between the two E7 nodes interconnecting the rings.



RSTP Node Protection planning guidelines:

- A node protection pair consists of two E7-2 nodes present on the ERPS ring, that are configured for a protected node hand-off from the ERPS ring to the protected trunk or edge ports, for a maximum of 16 locations (32 E7s) in an ERPS ring where each node on the ERPS ring can be utilized for uplink and/or downlink interfaces.
- A E7 node may only belong to one protection pair.

- Each node protection pair must be given a unique ID.
- When E7s are configured with node protection pairs, they can still use the ERPS ring for traffic that is destined for another node on the ring. However, the ERPS ring must ensure the ring integrity.
- When an E7 is configured for RSTP Node Protection, the RSTP domain of the two E7 nodes must be associated with one of the ERPS domains.
- Multiple node protection pairs can be defined within an ERPS ring. This allows the ERPS ring to have multiple hand-off points where trunk link or edge link ports are protected.
- RSTP links on the E7-2 can be to any port on any card, however, both of the device ports should be connected to like ports on the E-Series. For example, either two GE or two 10GE ports should be used.
- For each E7 RSTP link, add the interface to the desired traffic VLAN membership(s), or use the default Native VLAN.
- Typically, two E7 nodes forming a protection pair will be located in the same physical location. Yet, they are not required to be located immediately adjacent to each other on the ERPS ring.
- When configuring a protection domain with a pair of E7 systems, one E7 is provisioned as the primary RSTP node the other is provisioned as the secondary node. The link on the primary node carries traffic and the corresponding link to the secondary node is blocked, if the link to the primary node is operational. If a link fails on the primary node, traffic flows on the secondary link. When the primary link has been restored, the traffic automatically reverts back to the normal operational mode.
- There is no correlation between the ERPS ring master node and the assignment of the protection pair primary and secondary nodes. The ERPS ring master node can be one of the protection pair nodes or can be another node on the ring.
- When an E7 node is provisioned as part of a protection pair (primary or secondary), then all of the node trunk links or edge links provisioned with node protection enabled will become part of the node protection pair.
- Node protection pairs automatically enable node protection on the ERPS ring ports.
- For E7, cross-card LAGs cannot be used as a protected interface in a node protected pair.
- Node protection pairs automatically assign special path costs for the trunk link, edge link, and ERPS ring ports that participate in the node protection.


Note: RSTP Node Protection is not supported for E7-20s or E7-2 Modular Chassis.

Ethernet Point-to-Point Aggregation

Integrated high-capacity aggregation: The E7 is built on a core Layer 2 and Layer 3 switch capable of full-duplex, wire speed forwarding at all frame sizes and traffic types across all interfaces. This capacity makes the E7 ideal for aggregation and transport of IP/Ethernet services across the access network.

Hub and spoke network


The E7 ESAP can serve as the fundamental building block for transport and aggregation networks regardless of size or location. In the central office, the E7 acts as the central node that serves remote E7s over 10GE or multiple GE links using standards including RSTP and Link Aggregation.



Point-to-point star configuration

At the most basic level, the E-Series can aggregate Ethernet services from the Calix C7, E5, and F5 platforms into a single network uplink. This traditional “star” configuration can be deployed within the same E7-2 system used for GPON services with the GPON-4 card’s integrated GE and 10GE interfaces.

The E-Series supports protected connectivity for ports configured as uplinks that are connected to any standards-based Ethernet switches or routers using the same RSTP and/or link aggregation protocols. For E7, RSTP protected interfaces must be on different cards within the E7 shelf to provide equipment protection.



Redundant Calix E-Series Point-to-Point Ethernet Aggregation


Point-to-point Ethernet services on every E7 card

The E7-2 GPON-4 and 10GE-4 cards and the E7-20 GE-24x card include a large number of GE ports for aggregating other Ethernet equipment or for delivering point-to-point business and residential services alongside the primary E7 applications. Deployment of customer-based Ethernet services is accomplished with Calix 700GX ONT's or provider-owned switches or routers.



E7 Point-to-Point Ethernet Service Delivery

The high port count of the E7-2 GE-24 card and the E7-20 GE-24x card make them ideal for high density residential and business services, and the cards' EXA Powered core also allow them to provide Ethernet aggregation to subtended network elements. Both hub-and-spoke and ring-based connections to subtended devices are supported.




Point-to-point or Active Ethernet (AE) is a complimentary network design to point-to-multipoint GPON access networks. As a dedicated service, point-to-point Ethernet allows network operators to provide individual customers with very high-bandwidth, dedicated symmetric services. It also has the greatest deployment flexibility and OSP simplicity, with the tradeoff of requiring additional fiber management in the CO and OSP.

The E7-2 GE-24 card includes two 10GE ports supporting industry standard SFP+ modules for use as uplink and transport for the local E7 system, or to subtend an E7 chassis with GE-24 cards from a primary E7 system used for aggregation and transport of IP services. Multiple E7-2 systems can be linked together in a Modular Chassis system, using low cost 10GE SFP+ copper cables, resulting in a high-density point-to-point Active Ethernet configuration of 480 AE ONT subscribers in 10RU space.


Point-to-Point Ethernet Services on E3-48C

The Calix E3-48C easily integrates into Calix B-Series, C-Series, and E-Series networks via protected or point to point connections to deliver short loop high speed service deployments. The E3-48C supports ERPS and directly connects to Calix E-Series and C-Series networks using either GE or 10GE connections. As shown in the diagram below, the Calix E3-48C products have been deployed in Fiber-to-the-Node (FTTN) locations via 10GE delivering high speed services to subscriber locations.



Ethernet Aggregation Topology


The E-Series provides the ability to aggregate standards-based Ethernet traffic from the full range of Calix broadband solutions, Ethernet switches, and routers. Uplink and downlink aggregation can be provided through one or more GE or 10GE port. The protocols that the E-Series supports for aggregation include EAPS and single link or redundant links running RSTP or link aggregation.



E-Series Ethernet Aggregation

The above diagram depicts the E7 deployed in a remote location aggregating C7, E5-100, and E5-300 broadband access platforms. In this configuration, the E7 is aggregating the Ethernet traffic from the remote devices and switching it onto the transport ring for delivery to the network uplink. This diagram shows a redundant two-card E7 system providing aggregation redundancy and network uplink protection.

The following diagram shows the physical GE connections that enable link protection (RSTP) from the E5-100 Ethernet service platform to the redundant E-Series system.




Ethernet Aggregation with RSTP Downlink Protection

Protected Uplink with Virtual Router Redundancy (VRRP)

The E-Series can be used with IETF RFT 3768 Virtual Router Redundancy Protocol (VRRP) to increase the availability of the default gateway. Service providers use VRRP to provide router backup, should one of the routers fail. The E-Series does not play a role in this protocol and only responds to the routers' ARP and gratuitous ARP messages to determine if a router protection switch has occurred. RSTP is not run on the links between the E-Series and the routers because the routers are preventing loops by controlling the advertising of IP and MAC addresses into the E-Series Layer 2 network.

The redundant E-Series connections to the VRRP-enabled routers can be from from two different E-Series systems or from the same E7 shelf on separate line cards within the Layer 2 broadcast domain. This latter configuration, with two E-Series systems on the same 10GE ring but at different locations, is shown in the following diagram.



E-Series Network Connected to Routers Running Virtual Router Redundancy Protocol (VRRP)

Transparent LAN (TLAN)


The E-Series can be arranged in a point-to-point Ethernet configuration that delivers Transparent LAN (TLAN) services (TLS) when symmetrical bandwidth is required. You could also adapt this arrangement for a multi-point drop configuration and restrict the bandwidth on ingress traffic using a policy map. The following diagram shows how the E-Series can be used to transparently trunk business traffic across a network to another location, typically a remote office or secondary business location.

The traffic the E7 receives from the business can have any of the following VLAN configurations:

- Tagged
- Untagged
- Double-tagged

The E7 performs an add outer tag (S-tag) action to all incoming traffic from the business to simplify the local switching of TLAN traffic from any designated TLAN interface to any other TLAN interface. This accommodates any RSTP BPDUs as well.

In the example diagram below, the incoming business traffic is untagged, single-tagged, or double-tagged traffic. The E7 adds an outer tag (S-tag) of 50 that is a designated TLAN interface. The business traffic is transported to the other TLAN interface where the E7 removes the outer tag and delivers the traffic.



Note: For instructions on how to configure a TLAN service on ONT Ethernet ports, see "Creating Data Service on an ONT Ethernet Port" in the *Calix E-Series GPON Applications Guide*.

Configuration guidelines

The following conditions must apply for the VLAN and Ethernet interfaces that are used to configure the transparent LAN service:


- For TLAN to be enabled, the following must be/remain disabled on the VLAN: DHCP Snoop, MAC Forced Forwarding, IP Source Verify, and IGMP Mode must be/remain Flood (no snooping).
- In Ethernet security profiles that are applied to ONT ports, the L2CP Filter parameter must be set to all-tunnel.
- Interfaces at each end of the TLAN:
 - Split-horizon forwarding is disabled.
 - Role is designated as an edge link.
 - RSTP Enabled is set to tunneled.
 - LACP Tunnel is set to Yes.
 - If subscriber endpoints on the TLS have a maximum transmission unit (MTU) of 1500, the MTU for the Ethernet interfaces on the E7 should be increased to at least 1504 to accommodate the additional TLAN tag that is placed on the traffic within the E7.
- A Link Aggregation Group (LAG) configured across the entire TLAN (end-to-end):
 - Provision a separate VLAN for each link of the LAG.
 - Configure the link aggregation timers on the external link aggregation equipment to lower values so remote failures are detected faster.

- A link aggregation group terminated on each end of the TLAN:
 - Create Link Aggregation interfaces on the E7s.
 - LACP Tunnel is set to No so remote failures will produce an immediate response, such as link switching.

Operation Domains

Operation domains within the system

- The Modular Chassis Controller cards, and E7-20 SCP Controller cards:
 - Primarily participate in the domain that uplinks to the network, connects to transport rings, and aggregates other Calix access nodes.
 - Secondly participate in the domain that allows subscribers to access services.
- The Modular Chassis Expansion cards, and E7-20 Line Cards only participate in the domain that allows subscribers to access services.



Interface role configuration guidelines

	Trunk	Edge	Access
E7-2	X	X	X
E7-2 MCC	X	X	X
E7-2 MCE			X
E7-20 SCP	X	X	
E7-20 GE-24			X
ONT Ethernet Port			X
Tag Actions		X	X
Native VLAN	X	X	
Networking Protocol (RSTP, ERPS, LAG)	X	X	

Interface roles

- **Trunk:** A port connecting to other equipment belonging to the service provider or to another service domain with consistent VLAN tagging levels. These ports may also be referred to as Network ports or Provider ports in industry standards. These ports support outer VLAN tag plus MAC switching.

Examples of trunk ports are 10G ERPS transport ports and GE uplinks. Trunk ports can be configured for link aggregation, RSTP, or ERPS. To properly process ingress double tags, the GE network interface (uplink) must be configured as a Trunk role.

- **Edge:** A port facing customer equipment or facing reduced functionality devices, alternative administrative domains, or managed CPE. Generally, this E-Series Ethernet port interface is where all classification is first performed on ingress traffic (if customer facing). The E-Series Ethernet interface is also expected to add, replace, or remove one or more VLAN tags on edge traffic. ERPs, RSTP, and LAG networking protocols are supported.

Examples of an edge port would include GE ports to managed CPE, a GE/10GE port to external equipment which may use different tagging levels, or GPON ports.

- **Access:** A port facing untrusted customer equipment or other devices serving subscribers. Generally, this port interface is where individual subscriber services are defined and enforced (bandwidth limits, security, multicast profiles). Networking protocols are not supported.

Examples of an access port would be point-to-point connections to subscribers or other devices serving subscribers.



Chapter 4

Quality of Service

The Quality of Service (QoS) plan for a network defines how services are prioritized relative to each other. This chapter provides engineering and planning information regarding E-Series QoS support.

Topics Covered

This chapter covers the following topics:

- Quality of Service Overview
- Traffic Classification and Queuing
- Ethernet Traffic Management
- GPON Traffic Management
- xDSL Traffic Management

Quality of Service Overview

Quality of Service (QoS) refers to the ability of a network to provide better service by prioritizing network traffic. QoS allows you to manage traffic flows to maintain priority / delay / loss differentiation within the E-Series network, using the following methods.

- **Classification** of packets to determine its service level
- **Policing** of ingress traffic to ensure guaranteed traffic bounds can be met
- **Queuing and scheduling** which prioritizes and shapes output traffic (scheduling) across multiple traffic classes (queuing)

Ingress traffic

- Bandwidth limits on a per class-map
- Classifies traffic packets of interest based on P-bit, VLAN ID, DSCP, MAC OUI (not supported for T-Series ONTs)
- Marks the traffic with a new P-bit value to properly queue the traffic

Egress traffic

- Bandwidth shapes on a per class of service (P-bit)
- Ensures that all queued traffic with the highest priority is delivered to the network first
- Sets the maximum allowable processing rate (rate limiting) at which packets are forwarded
- Ensures the bursting nature of the traffic is reduced to the value set (traffic shaping)

The E-Series uses the following traffic-control objects that support the Quality of Service (QoS) plan for a network, as described below.

Ingress QoS Features (for traffic entering an access or edge port)

QoS Feature	E-Series Implementation
Class of Service Mapping By predefined matching criteria, ingress frames are mapped to a class of service (CIR or PIR) and forwarded to specific VLANs.	The ingress traffic can be classified based on the P-bit, VLAN, and DSCP (IP Precedence) values. Additionally, the ONT and VDSL2 card can classify traffic based on the MAC OUI values of the ingress traffic., and frames are forwarded to the specified VLAN (via VLAN-IFs). Each traffic class can be treated differently to differentiate the service implied for the subscriber, such as rate limiting, shaping (burst accommodation), or prioritizing relative to other traffic classes. • Ethernet Ports use Policy maps which are lists of QoS-related actions to perform on packets that match certain criteria. The policy map is associated with an E7 Ethernet port interface. The matching criteria are specified by an associated classification map that lists rules to identify packets: match P-bit, outer tag, inner tag, specific DSCP value. The QoS actions that can be taken on matched packets: --Specify P-bit (priority)-- The P-bit value defined in the policy map must be consistent with the class of service type, as defined in the Ethernet port Class of Service (CoS) table. --Specify rate limit – ensures ingress traffic does not exceed a specified bit rate. --Specify maximum burst size – ensures the bursting nature of traffic is reduced to the specified value. • GPON Ethernet and xDSL services use Service Tag Actions when a service is provisioned at an ONT Ethernet port or an xDSL port, a service-tag action (and associated match list) specify the classification and marking of packets from the subscriber port into the service VLAN.
Rate Limiting Ingress frames are rate-limited as required.	Ingress frames bound for a specific VLAN are rate-limited by the following: • Ethernet ports use Policy Maps that indicate an action that specifies rate limit for the packets that match certain criteria specified in an associated Class Map – ensures ingress traffic does not exceed a specified bit rate. • GPON Ethernet and xDSL services use Ethernet Bandwidth Profiles to specify the upstream and downstream bandwidth rates to apply to individual Ethernet services on ONT Ethernet ports and xDSL ports. • ONT ports can use a Dynamic Bandwidth Allocation (DBA) feature that dynamically changes upstream GPON bandwidth assignments at the port level based on traffic class and demand.
P-Bit Add/Overwrite The P-bit values of incoming frames are maintained or overwritten with new values	The P-bit values can be honored, changed to a predefined value, or mapped from the DSCP values of incoming IP packets by the following: • Ethernet ports use Policy Maps to specify a P-bit value on a matched P-bit value, or use a Tag Action that allows you to add, change, or add and change VLAN tags on packets arriving on a specified GE or 10GE interface. You can use a set of matching criteria to select which packets to transform. Or, packets can be transformed unconditionally. • GPON and xDSL services use Service Tag Actions when a service is provisioned at an ONT Ethernet port or an xDSL port, a service-tag action (and associated match list) specify the classification and marking of packets from the subscriber port into the service VLAN.

QoS Feature	E-Series Implementation
VLAN Mapping¹ VLANs are mapped from the CVID space to the SVID space.	Ethernet ports use Tag Actions that allow you to add, change, or add and change VLAN tags on packets arriving on a specified GE or 10GE interface. GPON and xDSL services use Service Tag Actions when a service is provisioned at an ONT Ethernet port or an xDSL port, a service-tag action (and associated match list) specify the classification and marking of packets from the subscriber port into the service VLAN.

1: The priority field of ingress frames is also mapped/promoted to the service provider space. See P-bit Add/Overwrite.

Egress QoS Features (for data traffic leaving an access or edge port)

QoS Feature	E7 Implementation
Class of Service Mapping Downstream-bound frames of the same traffic class (CIR or PIR) are on the same VLAN, and this same mapping is maintained toward subscriber access ports.	• Ethernet and xDSL ports are associated with a Class of Service (CoS) profile so traffic is sent at a particular rate by controlling the queue of packets. At egress, the E7 maps 8 priority CoS queues (1–8) to the P-bit values (0-7) of the traffic. You can configure each of the CoS queues with a maximum and minimum bandwidth rate. • ONT ports are associated with a system-level PON Class of Service (CoS) for scheduling upstream traffic.
Traffic Prioritization When traffic flows from service VLANs to a subscriber edge port, guaranteed frames (CIR) always take priority over non-guaranteed frames (EIR). Among frames of the same traffic class (CIR or EIR) on the same VLAN ¹ , higher-priority queues are always emptied before lower-priority queues.	• Ethernet Ports are associated with a Ethernet Port Class of Service (CoS) profile so traffic is sent at a particular rate by controlling the queue of packets. At egress, the E7 maps 8 priority CoS queues (1–8) to the P-bit values (0-7) of the traffic. You can configure each of the CoS queues with a maximum and minimum bandwidth rate. • GPON and xDSL services use Service Tag Actions when a service is provisioned at an ONT Ethernet port or an xDSL port, a service-tag action (and associated match list) specify the classification and marking of packets from the subscriber port into the service VLAN. • ONT ports are associated with a system-level PON Class of Service (CoS) for scheduling upstream traffic.
Rate Limiting (Shaping) Egress frames are rate-limited as required.	• Ethernet Ports uses a per-port basis (as part of the Ethernet Port CoS Profile) ability to shape the aggregate traffic that includes a maximum rate and a maximum burst size that determines the level to which traffic is allowed to deviate from the configured rate shaper. The E7 auto-calculates the burst size if a value is not specified. • GPON Ports and xDSL ports use an Ethernet Bandwidth Profile to set the peak information rate and the peak burst size for downstream traffic.
VLAN Mapping² VLANs are mapped from the SVID space to the CVID space.	• Ethernet ports can have tag actions applied that are typically described as being applied to ingress traffic on Ethernet port interface, the reverse of the tag action is applied to packets on the egress direction of the interface. For example, if an Add-Tag action is applied to an interface, that same tag is stripped from packets leaving the interface. • GPON services and xDSL use Service Tag Actions when a service is provisioned at an ONT Ethernet port or an xDSL port, a service-tag action (and associated match list) these objects are typically described as being applied to traffic from the subscriber port into the service VLAN. The reverse of the tag action is applied to packets on the egress direction of the service port. For example, if an Add-Tag action is specified in service provisioned on a subscriber port, that same tag is stripped from packets leaving the subscriber port.

Class map

Defines a set of criteria (class rules) for classifying ingress traffic on GE and 10GE ports, such as VoIP traffic. Also defines the name for the class. It is applied to a policy map.

If the "Match Type" is set to "all," a traffic packet matches the class map criteria if it matches all of the rules in the map. In this case, you can add multiple rules to a class map as long as the rules do not conflict. In practice, this limits you to adding one match-tag rule, one match-2-tags rule, and one match-p-bit rule or match dscp rule.

If the "Match Type" is set to "any," you can add up to the limit of 100 class rules.

Policy map

Defines the traffic management policy (policy rules) for a set of class names. It is applied to an E7 Ethernet port interface.

The QoS actions that can be taken when packets match the selection criteria are as follows:

- **Out P-bit** - marks the traffic by specifying a P-bit value to properly prioritize the traffic. Upon egress, the frame is still priority tagged with the original P-bit value. The upstream traffic on a service VLAN must have the same P-bit value as is set for downstream traffic. The P-bit value defined in the policy map must be consistent with the class of service type, as defined in the Ethernet port Class of Service (CoS) table.

P-bit value that specifies the VLAN priority value. The default setting of pbit-none does not alter the existing pbit. The upstream traffic on a service VLAN must have the same P-bit value as is set for downstream traffic. The P-bit value must also be consistent with the class of service type as defined in the Ethernet port CoS table.
- **Rate Limit** - ensures ingress traffic does not exceed a specified bit rate. Maximum allowable processing rate in Mb/s at which packets entering an E7 Ethernet port interface are forwarded.
- **Maximum Burst Size** - ensures that the bursting nature of the traffic is reduced to the specified value. Set the maximum burst size to a value greater than default when using jumbo frames (for example 9000 kbyte) at higher rates. Typically, set the burst to the amount of traffic sent in twice the round trip time. Also consider whether the traffic is shaped before reaching the E7; the edge device that shapes traffic greatly reduces the buffer requirements in the infrastructure device. Also see Creating Ethernet Port Class of Service for information on maximum rate and burst size on a per-port basis.

Maximum allowable burst size in kilobytes or use "m" suffix for megabytes. You must set a non-zero value for the operation to complete. Typically, set the burst to the amount of traffic sent in twice the round trip time.

Note: If you create multiple policy rules that specify an Out P-bit, only the policy rule with the lowest sequence number ID is carried out.

Match list

Defines a set of criteria (matching rules) for classifying traffic on an ONT Ethernet port or an xDSL port. The match list defines how the ONT or VDSL2 card classifies subscriber traffic into a service VLAN. It is applied to a service tag action that is then associated with an Ethernet subscriber service.

A match list can contain both “tagged” and “untagged” match rules, up to 12 tagged rules and up to 16 untagged rules.

- **Untagged match rules** can match on a portion of the source MAC address as indicated by the Source MAC and Source MAC Mask attributes of a video Set-Top box. If no Source Mac value is available, as in data traffic, you can use the system default match list "all-untagged," to match all untagged traffic.
- **Tagged match rules** can match on any combination of the outer tag VLAN-ID and P-bit values. The Tag Protocol Identifier (TPID) can also be specified at the E-Series egress Ethernet port interface, but defaults to 0x8100.

Note: If a Residential Gateway (RG) is installed in the subscriber's premises, the RG classifies traffic into VLANs. Then for each service type, create a match list that contains a tagged rule, matching the service VLAN.

Class of Service (CoS) Profile

Defines a CoS profile to associate with a port so traffic is sent at a particular rate by controlling the priority queuing of packets. At egress, the E7 maps 8 priority CoS queues (1–8) to the P-bit values (0-7) of the traffic, with the highest priority queue matching P-bit 7 and the lowest priority queue matching P-bit 0. You can configure each of the CoS queues with a maximum and minimum bandwidth rate.

On a per-port basis, the E7 also supports the ability to shape the aggregate traffic that includes a maximum rate and a maximum burst size that determines the level to which traffic is allowed to deviate from the configured rate shaper. The E7 auto-calculates the burst size if a value is not specified.

Auto-calculated burst size = Number of bits that can be transmitted in 80ms at the configured rate limit

Example

- Rate limit = 100 Mbps
- Calculated burst size applied = 8,000 Kbits

Calix recommends the port burst size parameter reflect double the Round Trip Delay (RTD) of the network, with 40 ms used as the default RTD. It is not possible to have a burst size less than the MTU of the interface.

CoS application for Link Aggregation Groups (LAG)

Rate limiters and shapers on a LAG are applied as follows:

- The maximum and minimum bandwidth rates per CoS are applied to the LAG interface on ingress. The resulting rate limit is distributed over all links in the LAG.
- The port shaping rate and burst size for traffic shaping are applied per-Ethernet port independent of the other links in the LAG interface.

Note: The upstream traffic on a service VLAN must have the same priority value as is set for downstream traffic. The priority value must also be consistent with the class of service type as defined in the Ethernet port CoS table.

Traffic Classification and Queuing

Traffic classification categorizes network traffic according to various parameters into a number of traffic classes. Each traffic class can be treated differently to differentiate the service implied for the subscriber, such as rate limiting, shaping (burst accommodation), or prioritizing relative to other traffic classes.

The E-Series supports the ability to classify ingress traffic based on the P-bit, VLAN, and DSCP (IP Precedence) values. Additionally, the P-Series and GigaCenter ONTs and VDSL2 card can classify traffic based on the MAC OUI values of the ingress traffic. The E-Series then marks the traffic with a new P-bit value, if appropriate.

Using the newly marked P-bit values, the egress queues on the port then deliver the traffic to the network based on priority (with minimum CoS bandwidth guarantees if desired).

Most networks use the following commonly-defined Per-Hop Behaviors (PHB):

- **Default PHB(Per hop behavior)**—which is typically best-effort traffic
Any traffic that does not meet the requirements of any of the other defined classes is placed in the default PHB. The recommended DSCP for the default PHB is '000000'(0).
- **Expedited Forwarding (EF) PHB**—dedicated to low-loss, low-latency traffic
This class is suitable for voice, video and other real-time services. The recommended DSCP for expedited forwarding is '101110' (46).
- **Assured Forwarding (AF) PHB**—gives assurance of delivery under prescribed conditions
Traffic that exceeds the subscription rate faces a higher probability of being dropped if congestion occurs. The AF behavior group defines four separate AF classes. Within each class, packets are given a drop precedence (high, medium or low). The combination of classes and drop precedence yields twelve separate DSCP encodings from AF11 through AF43.
 - If congestion occurs between classes, the traffic in the higher class is given priority.
 - If congestion occurs within a class, the packets with the higher drop precedence are discarded first.
- **Class Selector PHBs**—maintains backward compatibility with the IP Precedence field
If a packet is received from a non-DiffServ aware router that used IP precedence markings, the DiffServ router can still understand the encoding as a Class Selector codepoint.

Queuing Overview

Packet priority may be enforced by queuing, which processes higher-priority packets before lower-priority packets.

Traffic queuing manages transmission of frames waiting to be forwarded over an interface. Queuing processes the frames based on priority, which is determined by the priority bits (P-bits) assigned to the frames.

Queuing on line card Ethernet ports and VDSL2 subscriber ports

The E-Series maps 8 priority CoS queues (1–8) to the P-bit values (0-7) of the traffic, where 0 is the lowest priority and 7 is the highest priority. Each CoS queue is assigned a scheduling priority. The E-Series maps directly from P-bit value to CoS queue to scheduling priority, as shown in the table below:

P-Bit	CoS Queue	Scheduled Priority	
0	0	0	(lowest)
1	1	1	
2	2	2	
3	3	3	
4	4	4	
5	5	5	
6	6	6	
7	7	7	(highest)

Ethernet frames transmit from each CoS queue in a strict priority fashion. Strict priority queuing ensures that all traffic (in the queues) with the highest priority is delivered to the network first, followed by the lower priority traffic. For example, a frame with priority 5 transmits from the CoS queue before a frame with priority 0.

Queuing on ONT Ethernet ports

See *Queuing Methods on ONT Ethernet Ports* (on page [264](#)).

Layer 3 Priority Mapping

The most commonly used method for marking traffic for quality of service treatment, outside of the internal mechanisms used in Service Provider networks, is Layer 3 use of the DiffServ Code Point (DSCP) field in the IP header. This mechanism derives from previously used Type of Service (TOS) bits, later IP Precedence field, and uses some pre-defined 6-bits on the packet for classification purposes.

Because CPE devices may exclusively support Layer 3 QoS classification and traffic conditioning mechanisms, and the E7-2 works in the Layer 2 Ethernet domain, the E-Series maps DSCP classified traffic into the Layer 2 format using IEEE 802.1p (P-Bit). Calix recommends that all received subscriber traffic be tagged into a P-Bit setting to avoid unnecessary traffic congestion or traffic receiving lower than desired priority.

The E-Series supports mapping Differentiated Services Code Point (DSCP) or IP Precedence priority bit values, used for packet classification on DiffServ networks (defined in RFC 2474 and RFC 2475), to IEEE 802.1p priority-bit values to classify traffic priority.

- **Traffic coming from the subscriber to xDSL Access ports** can reference the use of layer 3 priority values from the service tag action associated with the service, and then from the xDSL port associated interface, indicate whether to use the system-default profile named "access" or a custom profile to map DSCP or IP precedence values to P-bit values.

Note: You can either apply IP-Precedence or DSCP mapping tables, but not both. The upstream traffic on a service VLAN must have the same priority value as is set for the downstream traffic.

- **Traffic coming from the subscriber to GPON Ethernet ports** can reference the use of layer 3 priority values from the service tag action associated with the service, and then indicate whether to use the system-default profile named "access" or a custom profile to map DSCP values to P-bit values.
- **Traffic coming from the network to Ethernet Trunk and Edge ports** can use class map, class rules, and policy maps to match the ingress Layer-3 value of incoming frames and assign a corresponding P-bit (priority) value.

Some service providers' networks include core routers or aggregation switches that are not setup to use P-bits for traffic management downstream. Instead DSCP is used for this purpose. In these cases all traffic received by the access node has a P-Bit value of 0, and the access node needs to set the P-bit at the network interface, or traffic internal to the access node will be discarded at random (service unaware).

DSCP and IP Precedence

Differentiated services or DiffServ is a computer networking architecture that specifies a simple, scalable and coarse-grained mechanism for classifying and managing network traffic and providing quality of service (QoS) on modern IP networks. DiffServ can, for example, be used to provide low-latency to critical network traffic such as voice or streaming media while providing simple best-effort service to non-critical services such as web traffic or file transfers.

DiffServ uses the 6-bit Differentiated Services Code Point (DSCP) field in the IP header for packet classification purposes. DSCP replaces the outdated type of service (TOS) field.

Prior to its deprecation, the Type of Service field was defined as follows.

0	1	2	3	4	5	6	7
Precedence			Type of Service				

The last bit of Type of Service (bit 7) was defined as "Must Be Zero."

DSCP traffic classes

The DSCP traffic classes are described in RFC 2474, RFC 2597, RFC 3246, the P-Bit construct in IEEE 802.1p. The table below shows the mapping between Traffic Classes and P-bits:

IP DSCP Name	IP DSCP Abbreviation	IEEE 802.1p
CSx	Class Selector (where x is 0 through 7) as defined in RFC 2474	Mapped to P0 through P7 respectively. For example: CS0 mapped to P0.
BE	Best Effort (also default)	P0
AFxy	"x" refers to Traffic Class, and "y" refers to Drop Precedence (RFC 2597). Where x is 1,2,3,4 (higher is better), and y is 1, 2, 3 (lower is better)	
AF1y	Assured Forwarding	P2
AF2y	Assured Forwarding	P3
AF3y	Assured Forwarding	P4
AF4y	Assured Forwarding	P5
EF	Expedited Forwarding (RFC 3246)	P6

The following table shows the various traffic types and the acronyms used for them:

User Priority	Acronym	Traffic Type
1	BK	Background
2	-	Spare
0 (Default)	BE	Best Effort
3	EE	Excellent Effort
4	CL	Controlled Load
5	VI	Video, < 100 ms delay
6	VO	Voice, <10 ms delay
7	NC	Network Control

For more information, see the document *Calix Engineering & Planning Guide: L2 Ethernet Access Networks*, available on the Calix Resource Center.

Assured Forwarding (AF) PHB group

The IETF defines the Assured Forwarding behavior in RFC 2597 and RFC 3260. Assured forwarding allows the operator to provide assurance of delivery as long as the traffic does not exceed some subscribed rate. Traffic that exceeds the subscription rate faces a higher probability of being dropped if congestion occurs.

The AF behavior group defines four separate AF classes with Class 4 having the highest priority. Within each class, packets are given a drop precedence (high, medium or low). The combination of classes and drop precedence yields twelve separate DSCP encodings from AF11 through AF43 (see table). RFC 2597 recommends the following DSCP codepoints for the four AF classes.

Drop Probability	Class 1	Class 2	Class 3	Class 4
Low	AF11 (DSCP 10)	AF21 (DSCP 18)	AF31 (DSCP 26)	AF41 (DSCP 34)
Medium	AF12 (DSCP 12)	AF22 (DSCP 20)	AF32 (DSCP 28)	AF42 (DSCP 36)
High	AF13 (DSCP 13)	AF23 (DSCP 22)	AF33 (DSCP 30)	AF43 (DSCP 38)

The AFxy, format uses x as the class number and y as the drop precedence.

Some measure of priority and proportional fairness is defined between traffic in different classes. Should congestion occur between classes, the traffic in the higher class is given priority.

Expedited Forwarding (EF) is defined in RFC 2598 and falls under the DSCP Class 5 (CS5). EF has the characteristics of low delay, low loss and low jitter. These characteristics are suitable for voice, video and other realtime services. EF traffic is often given strict priority queuing above all other traffic classes. The RFC recommended DSCP value for EF corresponds to a DSCP value of 46.

Best Effort (BE) is indicated by a DiffServ value of 0.

802.1P to DiffServ DSCP values

Below is a pre-defined mapping table of 802.1P to DiffServ DSCP values based on the recommendations of the Internet Engineering Task Force (IETF).

Service	IEEE 802.1p Priority	DiffServ Code Point
Network Control	7	CS7
Voice, NSP, CES	6	EF
Network Management	5	CS5
Video	4	CS3
Data3	3	AF3-3
Data2	2	AF21-3 & CS-2
Data1	1	AF11-3
Best Effort	0	0

The service provider network administrator is responsible for ensuring that no application will be introduced into the network that may contend for bandwidth at a given priority, unless the application was explicitly designed into the QOS scheme of the network. For instance, accepting traffic from a business network or savvy residential user without remarking the DSCPs on ingress may introduce unanticipated congestion at a critical priority level. This is both a network design and security issue.

About Creating a DSCP Map for Ingress Traffic

You can create and modify an DSCP Map profile that is applied to xDSL port interfaces to enable and/or specify a mapping of the DSCP values to IEEE 802.1p priority bits.

To designate the DSCP profile values as the ingress traffic access list, do the following:

Proprietary Information: Not for use or disclosure except by written agreement with Calix.

1. In the service tag action that is going to be applied to the data or video service, select **Map a Layer-3 priority** from the P-Bit Source list.
2. On the xDSL port associated interface where you are going to be provisioning the data or video service, select the name of the DSCP profile to use from the DSCP/IP Precedence Profile list.

Note: The upstream traffic on a service VLAN must have the same priority value as is set for downstream traffic.

Note: If IP traffic on an Access Interface has a DSCP setting other than the values included in the assigned DSCP map, the E7 system interprets this condition as a mismatch and uses a "Default" P-bit value from the assigned DSCP Map profile. For example, if you were using the system-default DSCP Map "Access" and the incoming frames had a DSCP value of 001 001, that value is not included in the assigned DSCP map, so the E7 would use the default P-bit value of 0. However, you can create a DSCP Map with a Default P-bit value of your choice, and then assign the DSCP Map to xDSL port interfaces.

The system-default DSCP profile "Access," has values as follows:

DSCP Value	Meaning	Decimal	Drop Probability	802.1p Priority-Bit
000 000	BE-CS0	0	n/a	0
001 000	CS1	8		1
001 010	AF11	10	Low	1
001 100	AF12	12	Medium	1
001 110	AF13	14	High	1
010 000	CS2	16		2
010 010	AF21	18	Low	2
010 100	AF22	20	Medium	2
010 110	AF23	22	High	2
011 000	CS3	24		3
011 010	AF31	26	Low	3
011 100	AF32	28	Medium	3
011 110	AF33	30	High	3
100 000	CS4	32		4
100 010	AF41	34	Low	4
100 100	AF42	36	Medium	4
100 110	AF43	38	High	4
101 000	CS5	40		5
101 110	EF	46		5
110 000	CS6	48		0
111 000	CS7	56		0
	Default			0

Best Effort (BE)

Assured Forwarding (AF) gives assurance of delivery under prescribed conditions

Expedited Forwarding (EF) dedicated to low-loss, low-latency traffic

Default = P-bit assigned for a DSCP value that is not included in the DSCP Map

About Creating an IP Precedence Map

You can create and modify an IP precedence map that is applied to xDSL port interfaces to enable and/or specify a mapping of the IP precedence values to IEEE 802.1p priority bits.

To designate the IP precedence profile values as the egress traffic access list, do the following:

1. In the service tag action that is going to be applied to the data or video service, select **Map a Layer-3 priority** from the P-Bit Source list.
2. On the xDSL port associated interface where you are going to be provisioning the data or video service, select the name of the IP precedence profile to use from the DSCP/IP Precedence Profile list.

The system-default IP precedence profile "access," has values as follows:

IP Precedence Value	Meaning	Decima 1	802.1p Priority-Bit
000 xxx	Routine/Best	0	0
001 xxx	Effort	1	1
010 xxx	Priority	2	2
011 xxx	Immediate	3	3
100 xxx	Flash/Voice/Video	4	4
101 xxx	Flash Override	5	5
110 xxx	Critical/Voice RTP	6	0
111 xxx	Internet Network	7	0

Note: Calix recommends NOT assigning P-bit value 7 to any DSCP value. The upstream traffic on a service VLAN must have the same priority value as is set for downstream traffic.

About Creating a Policy Map for L3 Priority Mapping

Ingress traffic on Trunk or Edge Ethernet ports can use class map, class rules, and policy maps to match the ingress Layer-3 value of incoming frames and assign a corresponding P-bit (priority) value. This assumes that the upstream routers or aggregation switch are not P-bit aware and do not manage traffic based on P-bit marking.

For detailed instructions on how to create a policy map for L3 priority mappings, see the *Calix E-Series (E7 OS) xDSL Applications Guide*.

Note: The upstream traffic on a service VLAN must have the same priority value as is set for downstream traffic.

The following table shows the different classes of traffic as they correlate to AF ranges. Classes 1 to 4 are referred to as Assured Forwarding PHB (Per Hop Behavior). AF is divided into four independently forwarded AF classes and is defined in RFC 2597. Within each AF class there are three different levels of drop precedence. In case of network congestion, a packet with a higher drop precedence (drop probability) value would be discarded first.

RFC 2597 recommends the following DSCP codepoints for the four AF classes. The AF_{xy} format uses *x* as the class number and *y* as the drop precedence.

Drop Probability	Class 1	Class 2	Class 3	Class 4
Low	AF11 (DSCP 10)	AF21 (DSCP 18)	AF31 (DSCP 26)	AF41 (DSCP 34)
Medium	AF12 (DSCP 12)	AF22 (DSCP 20)	AF32 (DSCP 28)	AF42 (DSCP 36)
High	AF13 (DSCP 13)	AF23 (DSCP 22)	AF33 (DSCP 30)	AF43 (DSCP 38)

The AF_{xy} format uses *x* as the class number and *y* as the drop precedence.

Expedited Forwarding (EF) is defined in RFC 2598 and falls under the DSCP Class 5 (CS5). The RFC recommended DSCP value for EF corresponds to a DSCP value of 46.

Best Effort (BE) is indicated by a DiffServ value of 0.

VLAN Tag Actions and Traffic Classification

The E-Series provides standards-based IEEE 802.1Q VLAN tagging. VLAN tagging was developed as a means to allow multiple networks to transparently traverse the same physical network. In addition, 802.1Q can be used to create transparent, virtual LANs (VLAN) and separate traffic, ensuring end user traffic segregation and network security. When VLANs are provisioned as a unique customer or port identifier, VLAN C-tags (customer tag) are utilized to create a VLAN per customer / port association.

The VLAN tagged frames are identified as having a tag by utilizing the Ethertype field, which is part of a standard Ethernet frame and is part of the Ethernet packet overhead. The Ethertype field is set by default for VLAN tagging (0x8100). This value is a software provisionable value in the E-Series since some switches and routers require specific values that don't follow the standard recommendation for Ethertype designations.

E-Series also supports and Q-in-Q VLAN stacking, based on IEEE 802.1ad Bridged VLANs - 802.1Q Amendment 4, Provider Bridges, which describes the ability to add multiple tags onto an Ethernet frame. Multiple VLAN tags allow for support of multiple provisioning models as described in Broadband Forum's TR-101 standard ("Migration to Ethernet-Based DSL Aggregation"). The ability to add multiple VLAN tags (called VLAN stacking or Q-in-Q) enables the following:

- Expands the addressable VLAN space from 4094 VLANs to well over 16 million VLANs
- Allows logical separation and trunking of VLANs through a network by using a VLAN tag to group a larger range of VLAN tags together

The most common way to use VLAN stacking is by inserting two tags on the traffic. These tags are often referred to as the inner, Customer tag or C-tag, and the outer, Service Provider tag or S-tag. As stated previously, the C-tag is used to uniquely identify a customer, and typically is used on a per port basis. The S-tag, is used to logically group C-tags together, typically is used by the ISP but can be used for any logical grouping (such as DSLAM ID, line card ID, identifying an RT of remote CO, etc).

Ethernet Traffic Management

In order to ensure proper Ethernet network management, the <e-type2> has been designed to support rate limiting and traffic shaping. The primary application for this functionality is to ensure service-level agreements and provide control over the amount of bandwidth utilized for service delivery, especially for TLS type services. This functionality allows the <e-type2> to limit the rate at which a subscriber can send traffic to the network and prevents a negative impact on the network from one high bandwidth user. This also enables the service provider to offer tiered services and gain additional revenue from customers requiring more bandwidth.

Ethernet Traffic-Control Objects


The E-Series have the ability to classify ingress traffic based on the P-bit or VLAN values, and for the E7, MAC OUI values.

Depending on the provisioned classifications, an E-Series marks the incoming traffic with a new P-bit value to properly queue the traffic at the egress port. The E-Series supports a queuing structure on the port for egress traffic that defines 8 priority queues: 7 specifies the highest priority queue where the traffic is delivered to the network first, and 0 specifies the lowest priority queue. The Quality of Service (QoS) scheme allows service providers to determine the priority of traffic within the network and properly traffic manage the entire Ethernet network. See "Creating Ethernet Port Class of Service" in E-Series system user guide for details.

Note: P-bit values are defined by IEEE 802.1p. This specification is used for Layer 2 Ethernet switches to prioritize traffic.

The E-Series uses two traffic-control objects for Ethernet traffic management:

- **Class map:** defines a set of criteria for classifying traffic into a class, such as VoIP traffic. Also defines the name for the class.
- **Policy map:** defines the level of bandwidth or queue priority to apply to traffic that matches the criteria specified in a referenced class map. This map is applied to an Ethernet port interface.



The E7 also uses the following traffic-control object for Ethernet traffic management:

- **Class of Service (CoS):** Defines a CoS profile to associate with an Ethernet port or GPON port so traffic is sent at a particular rate by controlling the priority queuing of packets.

Traffic management and rate limiting

To ensure proper Ethernet network management, the E7 supports rate limiting and traffic shaping. The primary application for this functionality is to ensure SLA service level agreements and provide control over the amount of bandwidth utilized for service delivery.

The E-Series QoS features provide the following:

- **Ingress traffic**
 - Bandwidth limits on a per class map / policy map
 - Classifies traffic packets of interest based on P-bit or VLAN ID
 - Marks the traffic with a new P-bit value to properly queue the traffic
- **Egress traffic**
 - Bandwidth shapes on a per Class of Service (CoS)
 - Ensures that all queued traffic with the highest priority is delivered to the network first
 - Sets the maximum allowable processing rate at which packets are forwarded
 - Ensures the bursting nature of the traffic is reduced to the value set (traffic shaping)

Class Map and Policy Map

Ingress rate limiting is sometimes called traffic policing because it ensures ingress traffic does not exceed a specified bit rate, keeping a subscriber from exceeding their data rate contract. Rate limiting applies to traffic entering an E-Series network from a downstream device.

If FTP is being used to determine the actual throughput of an interface, the results can be distorted in certain scenarios if the acknowledge (ACK) sent to indicate receipt of a packet is dropped. This causes a retransmission of the original packet even though it may have been successfully delivered. The results can then be distorted by the retransmission of traffic, making the throughput look substantially lower than the actual capability of the interface.

Ethernet Port Class of Service Profile

Egress traffic shaping ensures that traffic leaving the E7 network does not exceed a particular bandwidth requirement by controlling the queuing of packets. This functionality reduces the bursting nature of traffic, enabling a downstream device to process the traffic without overflowing traffic queues and dropping packets. If a queue does overflow, the traffic is dropped at the tail of the queue.

The E-Series supports the ability to shape traffic on a per CoS basis with up to 8 CoS queues supported.

Egress rate shaping is supported on a per-port basis, and the E-Series also supports the ability to shape the aggregate traffic that includes a maximum rate and a maximum burst size that determines the level to which traffic is allowed to deviate from the configured rate shaper. The E-Series auto-calculates the burst size if a value is not specified.

Auto-calculated burst size = Number of bits that can be transmitted in 80ms at the configured rate limit

Example

- Rate limit = 100 Mbps
- Calculated burst size applied = 8,000 Kbits

Calix recommends the port burst size parameter reflect double the Round Trip Delay (RTD) of the network, with 40 ms used as the default RTD. It is not possible to have a burst size less than the MTU of the interface.

CoS application for Link Aggregation Groups (LAG)

Rate limiters and shapers on a LAG are applied as follows:

- The maximum and minimum bandwidth rates per CoS are applied to the LAG interface on ingress. The resulting rate limit is distributed over all links in the LAG.
- The port shaping rate and burst size for traffic shaping are applied per-Ethernet port independent of the other links in the LAG interface.

Traffic Rate Limiters

The possible 1500 individual ingress traffic rate limiters (policing) per card classify services by S-tag, C-tag, or P-bit and can be individually policed on ingress interfaces. Typically, service bandwidth policing and traffic shaping is performed at the provider edge port (e.g. ONT Ethernet port) to provide the most scalable network design without impact to high speed aggregation points in the network. However, some applications may require intermediate points in the network to provide additional service rate policing. Network examples might include:


- Mid-network policing of multi-point Transparent LAN services (Metro Ethernet Forum E-LAN)
- Downstream rate limiting of specific service (e.g. Internet) sharing the same unicast subscriber VLAN as other traffic (e.g. VOD)

Ethernet Traffic Data Path

This topic describes the path and process that data packets undergo during an E-Series ingress and egress traffic flow.


Overall data path and logical interface roles

Depending on the direction (ingress or egress) of the traffic through the E-Series and the role assigned to the logical interface, the following data path processes exist for the E-Series.




- Trunk interfaces can forward traffic to/from all interface types
- Edge interfaces:
 - Can forward traffic to/from trunk interfaces
 - Can forward traffic to/from another edge interface if Split Horizon Forwarding is disabled

Data path summary



- 10GE / GE line rate forwarding with Quality of Service (QoS)
- 64G non-blocking Layer 2 switch (64-9000 byte frames)
- 32k MAC addresses
- IEEE 802.1Q VLANs and 802.3ad Provider Bridging (Q-in-Q VLAN stacking)
- IEEE 802.1p service prioritization
- Integrated QoS and traffic management functions
- Integrated IGMP v2 snooping
- Layer 3 awareness (DHCP Relay, L3 ACL, etc.)

Ingress interface summary



- GE ports or 10GE ports
- IEEE 802.3ad link aggregation
- Ports are associated with:
 - Ethernet Interfaces
 - Link Aggregation Group Interfaces
- Interface roles:
 - Trunk (generally network facing, ERPS), Edge (generally edge facing)
 - Split Horizon Forwarding prevents forwarding of frames between Edge interfaces (can be disabled as required to support TLS and other trusted provider network applications)

Ingress tag-actions



- Tag actions are supported on Edge interfaces only
- Two types of Tag Actions:
 - **Add Tag**
 - Adds an outer tag based on matching VLAN and/or P-bit
 - **Change Tag**
 - Changes outer tag based on matching VLAN
- At Egress on an Edge interface, opposite tag action is applied
- If multiple tag actions are specified, most-specific criteria takes precedent
 - Total 768 tag actions per system
 - “Match any” rules are a single tag action

Ingress Quality of Service (QoS) classifying




- Ingress Quality of Service (QoS) has two components:
 - Class Map – identifies (**classifies**) frames
 - Policy Map – applies policy to traffic (policing)
- QoS can be applied to a VLAN service and/or an interface based on rules
- Classification rules match using:
 - Outer Tag
 - Inner Tag
 - P-bit
- Class Map can match any or all rules

Ingress Quality of Service (QoS) policing




- Ingress QoS has two components:
 - Class Map – Identifies (classifies) frames
 - Policy Map – Applies policy to traffic (**policing**)
- Applies policy (policing) to frames identified through a Class Map
 - Set or change P-bit
 - Rate limit (default is “no rate limit”)
- Policy Map gets applied to an Ethernet port interface

Data path switching




- 64G non-blocking switch (64-9000 byte frames)
- Unicast 10GE / GE line rate forwarding
- Multicasting at line rate: any-to-any interface without restriction
- 32k MAC address capacity
- IEEE 802.3ad link aggregation
 - No additional latency introduced (hardware-based flow distribution)

Egress Class of Service (CoS) queuing



- Class of Service (CoS) based queuing and scheduling per port
- CoS queues
 - Strict priority scheduling with minimum guaranteed bandwidth per class
 - Eight CoS queues based on P-bit (P-bit 7 = highest priority)
 - Rate shaping max rate – 1 Mbps increments
 - Minimum bandwidth for queue – 1 Mbps increments


Egress inverse tag-actions



- Tag actions are supported on Edge interfaces only
 - Ingress tag actions get reversed on egress

- Frames forwarded from a Edge interface to an ERPS or Trunk interface do not have second tag-action opportunity
- Frames switched from one Edge interface to another can have tag actions applied at ingress and egress
 - Typical application would be TLS (MEF EVP-LAN)
 - Split Horizon forwarding must be disabled on both ports
 - Bridging occurs within the common VLAN

Egress interface



- GE ports or 10GE ports
- IEEE 802.3ad link aggregation
- Interfaces are associated with a port or Link Aggregation Group
- Interface roles:
 - Trunk (generally network facing), Edge (generally edge facing)
 - Split Horizon is enabled on Edge interfaces by default (frames cannot switch between Edge interfaces)

Class Map and Class Rules

A class map with class rules specifies certain selection criteria, such as P-bit or VLAN values, to classify traffic from the network and arriving at the E7. If packets match the criteria listed in the class map, an associated policy map states the Quality of Service (QoS) actions to perform on those identified packets. The class map can specify whether packets must match any or all rules in order to be selected.

Class map rule for untagged traffic destined for the Native VLAN

The E7 Ethernet port Trunk and Edge interfaces are always associated with at least one VLAN, through the native VLAN attribute (VLAN 1 by default). The native VLAN ID is specified on the interface and is used for untagged user traffic on that interface.

Note: Interfaces with the Access role do not support a native VLAN. To forward untagged traffic on E7 Ethernet ports with an Access interface, an add-tag action must be applied to untagged frames, assigning the traffic to a designated VLAN.

To create a class rule for matching untagged traffic, do the following:

Proprietary Information: Not for use or disclosure except by written agreement with Calix.

1. Specify the native VLAN ID when defining the Ethernet Interface attributes.
2. Set the following parameters when creating a class rule:
 - Match pbit = pbit-any
 - Match Outer = native VLAN ID
 - Match Inner = Any

For example, if you specify the native VLAN ID as 4 for the Ethernet interface, the class rule for matching untagged traffic would appear as follows:

3. Apply the class rule to a class map, and then in turn, associate the class map to a policy map that is associated to the Ethernet interface where you specified the Native VLAN ID.

Configuration guidelines

- If the "Match Type" is set to "all," a traffic packet matches the class map criteria if it matches **all** of the rules in the map. In this case, you can add multiple rules to a class map as long as the rules do not conflict. In practice, this limits you to adding one match-tag rule, one match-2-tags rule, and one match-p-bit rule or match dscp rule.
- If the "Match Type" is set to "any," you can add up to the limit of 100 class rules.

Class map and class rule parameters

You can provision the following parameters for class maps and class rules:

Parameter	Description	Valid Options
Name*	Name of classification map.	Any existing class map
Map Type	Match type.	any, all (default)
Index*	Index of rule in map. This is a numeric index value that uniquely identifies this object within the system. Index values start with 1.	1-100
Match pbit	P-bit value (range 0-7) that specifies the VLAN priority value to match. Alternately, if "pbit-any" is used or the parameter is not specified, the P-bit value is not considered.	pbit-0 to pbit-7, pbit-any (default)

Parameter	Description	Valid Options
Match Outer	Outer VLAN ID to match. VLANs can be specified by name or by numeric VLAN ID (range 1-4093). Selecting "Any" or "Ignore" indicates a match for both tagged and untagged traffic.	Ignore (default) Enter Value
Match Inner	Inner VLAN ID to match. VLANs can be specified by name or by numeric VLAN ID (range 1-4093).	Ignore (default), any, untagged, Enter Value
Match DSCP	DSCP value to match.	Not used (default), Enter Value: DSCP value 0-63, or: be cs0, cs1, af11, af12, af13, cs2, af21, af22, af23, cs3, af31, af32, af33, cs4, af41, af42, af43, cs5, ef, cs6, cs7.

*Required field

Policy Map and Policies

A policy map with policies specifies what Quality of Service (QoS) actions to perform on ingress traffic packets when the packets match selection criteria. The selection criteria are listed in a class map that is associated with the policy map. In turn, the policy map is associated with an E-Series Ethernet port interface.

See *Example: Application of Policy Map and Tag Action* (on page [232](#)) for an example of an E-Series Ethernet interface tag action and an interface QoS policy being applied to an Ethernet interface.

The QoS actions that can be taken when packets match the selection criteria are as follows:

- **Out P-bit** - marks the traffic by specifying a P-bit value to properly prioritize the traffic. Upon egress, the frame is still priority tagged with the original P-bit value. The upstream traffic on a service VLAN must have the same P-bit value as is set for downstream traffic. The P-bit value defined in the policy map must be consistent with the class of service type, as defined in the Ethernet port Class of Service (CoS) table.
- **Rate Limit** - ensures ingress traffic does not exceed a specified bit rate.
- **Maximum Burst Size** - ensures that the bursting nature of the traffic is reduced to the specified value. Set the maximum burst size to a value greater than default when using jumbo frames (for example 9000 kbyte) at higher rates. Typically, set the burst to the amount of traffic sent in twice the round trip time. Also consider whether the traffic is shaped before reaching the E7; the edge device that shapes traffic greatly reduces the buffer requirements in the infrastructure device. Also see *Creating Ethernet Port Class of Service* for information on maximum rate and burst size on a per-port basis.

Configuration Considerations

- The ports have a single-rate traffic shaping function, differing from the MEF scheme and terms. For example, in MEF terms, Committed Burst Size (CBS) is a component of a two-rate traffic shaping system. Where you have:
 - CIR (Committed Information Rate)
 - CBS (Committed Burst Size) for the CIR buffer
 - EIR (Excess Information Rate)
 - EBS (Excess Burst Size) for the EIR buffer

However, with the E-Series single-rate configuration settings, you have:

- Maximum Rate
- Burst Size
- Consider the provisioned Maximum Rate to be the equivalent of the provisioned Committed Information Rate (CIR) rate in a two-rate system when the Excess Information Rate (EIR) value = zero. That is, $CIR + (EIR = 0) = PIR = Max$.
- The policies are processed in a specified sequence (lowest sequence number first), so when a match occurs, no more policies are processed for that packet. This process is important when you have overlapping match criteria. For example: If you create multiple policy rules that specify an Out P-bit, only the policy rule with the lowest sequence number ID is carried out.
- Ethernet port interfaces have the following aspects:
 - For Trunk interfaces, any policy map assignment is allowed.
 - For Edge and Access interfaces, if a policy map contains a two-tag classification, the edge or access link must have a tag-action that adds the outer tag being matched by the class rule.

Before starting

Ensure the class map that you want to associate to the policy map is already created and includes class rules.

Policy map and policy parameters

You can provision the following parameters for policy maps and policies:

Parameter	Description	Valid Options
Name*	Name of policy map.	text string
Seq Num*	Numeric value that specifies the sequence for processing the policies in the policy map. The lowest sequence number is processed first. If you create multiple policy rules that specify an Out P-bit, only the policy rule with the lowest sequence number ID is carried out.	1-1500

Parameter	Description	Valid Options
Class Map*	Name of class map to associate to the policy action.	Any existing class map
Out pbit	P-bit value that specifies the VLAN priority value. The default setting of pbit-none does not alter the existing pbit. The upstream traffic on a service VLAN must have the same P-bit value as is set for downstream traffic. The P-bit value must also be consistent with the class of service type as defined in the Ethernet port CoS table.	pbit-0 to pbit-7, pbit-none (default) Default = pbit-none
Rate Limit	Whether to set an maximum allowable ingress rate limit. this is the allowable processing rate in Mb/s at which packets entering an E7 Ethernet port interface are forwarded. If you select "Set a Rate Limit" you must also enter a limit value.	none (default) Set a Rate Limit (64kb/s to 10000mb/s)
Max Burst Size	Maximum allowable burst size in kilobytes or use "m" suffix for megabytes. You must set a non-zero value for the operation to complete. Typically, set the burst to the amount of traffic sent in twice the round trip time.	none (default), Set a Burst Size (4-16000 kilobytes)

*Required fields

Example: Application of Policy Map and Tag Action

When both an Ethernet interface tag action and an interface QoS policy are applied to an Ethernet interface, the order of operations is as follows:

1. The VLAN tag action occurs.

The P-bit from any incoming tag is automatically copied into the P-bit of the added tag.

2. The policy map is applied.

Policies are processed in a specified sequence (lowest sequence number first), so when a match occurs, no more policies are processed for that packet. This process is important when you have overlapping match criteria. For example: If you create multiple policy rules that specify an Out P-bit, only the policy rule with the lowest sequence number ID is carried out.

Example:

A business purchased a MetroE service and will be passing QoS from end to end. The Service Provider has agreed to give higher priority to the Enterprise voice traffic and default lower priority to the rest of the TLS traffic (but still higher priority than Internet).

Customer L3 switch sends traffic as follows:

- Voice tagged with VLAN 4000 and P-bit 6
 - Data tagged with VLAN 4000 and P-bit 0
 - Traffic ingresses an Ethernet port (no ONT involved)
1. A VLAN Tag action adds VLAN 100 and copies P-bits 0 and 6 into the new tag.

2. A class map criteria is applied:

- Class Rule A = match on P-bit 6, ignore VID (or match 4000).
- Class Rule B = match all (ignore VID, ignore P-bit)

3. A policy map is applied:

- Match on Rule A: set P-bit to 5 (SP scheme for TLS voice), rate limit set to 2 Mbps (negotiated)

Voice packets are tagged with inner VID 4000/P-bit 6 and outer VID 100/P-bit 5.

- Match on Rule B: set P-bit to 2 (SP scheme for TLS data), rate limit to 100 Mbps (negotiated)

Data packets are tagged with an inner VID 4000/P-bit 0 and outer VID 100/P-bit 2.

4. Egress CoS shaping limits the aggregate downstream rate on the port to 100 Mbps, or each queue to 2 Mbps/100 Mbps, as appropriate.

Ethernet Port Class of Service Profile

This topic describes how to create a Class of Service (CoS) profile to associate with an Ethernet port so traffic is sent at a particular rate by controlling the queue of packets. At egress, the E-Series maps 8 priority CoS queues (0–7) to the P-bit values of the traffic, with the highest priority queue matching P-bit 7 and the lowest priority queue matching P-bit 0. You can configure each of the CoS queues with a maximum and minimum bandwidth rate.

On a per-port basis, the E7 also supports the ability to shape the aggregate traffic that includes a maximum rate and a maximum burst size that determines the level to which traffic is allowed to deviate from the configured rate shaper. The E7 auto-calculates the burst size if a value is not specified.

Auto-calculated burst size = Number of bits that can be transmitted in 80ms at the configured rate limit

Example

- Rate limit = 100 Mbps
- Calculated burst size applied = 8,000 Kbits

Calix recommends the port burst size parameter reflect double the Round Trip Delay (RTD) of the network, with 40 ms used as the default RTD. It is not possible to have a burst size less than the MTU of the interface.

CoS application for Link Aggregation Groups (LAG)

Rate limiters and shapers on a LAG are applied as follows:

- The maximum and minimum bandwidth rates per CoS are applied to the LAG interface on ingress. The resulting rate limit is distributed over all links in the LAG.
- The port shaping rate and burst size for traffic shaping are applied per-Ethernet port independent of the other links in the LAG interface.

See "Traffic Management" and "Data Path" in the *Calix E-Series Engineering and Planning Guide*.

Note: The upstream traffic on a service VLAN must have the same priority value as is set for downstream traffic. The priority value must also be consistent with the class of service type as defined in the Ethernet port CoS table.

Class of service parameters

You can provision the following parameters for class of service:

Parameter	Description	Valid Options
Name*	Name of COS queue.	text string
Queue 1 Rate	Shaping rate for queue 1, specified in Mbps. The allowed range is 1–10000.	1–10000 Mbps
Queue 1 Min Bandwidth	Minimum bandwidth for queue 1, specified in Mbps. The allowed range is 1–10000.	1–10000 Mbps
Queue 2 Rate	Shaping rate for queue 2. Shaping rates are specified in Mbps.	1–10000 Mbps
Queue 2 Min Bandwidth	Minimum bandwidth for queue 2, specified in Mbps. The allowed range is 1–10000.	1–10000 Mbps
Queue 3 Rate	Shaping rate for queue 3, specified in Mbps. The allowed range is 1–10000.	1–10000 Mbps
Queue 3 Min Bandwidth	Minimum bandwidth for queue 3, specified in Mbps. The allowed range is 1–10000.	1–10000 Mbps
Queue 4 Rate	Shaping rate for queue 4, specified in Mbps. The allowed range is 1–10000.	1–10000 Mbps
Queue 4 Min Bandwidth	Minimum bandwidth for queue 4, specified in Mbps. The allowed range is 1–10000.	1–10000 Mbps
Queue 5 Rate	Shaping rate for queue 5, specified in Mbps. The allowed range is 1–10000.	1–10000 Mbps
Queue 5 Min Bandwidth	Minimum bandwidth for queue 5, specified in Mbps. The allowed range is 1–10000.	1–10000 Mbps
Queue 6 Rate	Shaping rate for queue 6, specified in Mbps. The allowed range is 1–10000.	1–10000 Mbps
Queue 6 Min Bandwidth	Minimum bandwidth for queue 6, specified in Mbps. The allowed range is 1–10000.	1–10000 Mbps
Queue 7 Rate	Shaping rate for queue 7, specified in Mbps. The allowed range is 1–10000.	1–10000 Mbps
Queue 7 Min Bandwidth	Minimum bandwidth for queue 7, specified in Mbps. The allowed range is 1–10000.	1–10000 Mbps
Queue 8 Rate	Shaping rate for queue 8, specified in Mbps. The allowed range is 1–10000.	1–10000 Mbps
Queue 8 Min Bandwidth	Minimum bandwidth for queue 8, specified in Mbps. The allowed range is 1–10000.	1–10000 Mbps

Parameter	Description	Valid Options
Port Shaping Rate	Aggregate shaping rate for the port, specified in Mbps. The allowed range is 1–10000. Alternately, select none to indicate that shaping should not be done.	none (or keyword "unshaped") Enter Value: 1–10000 Mbps or select a value from the list
Port Burst Size	Aggregate burst size for port, specified in Kbits. The allowed range is 1 to 128000. Alternately, select auto to automatically calculate the burst size.	auto (or keyword "auto") Enter Value: 1–128000 Kbits or select a value from the list

* Required field

VLAN Tag Actions

A tag-action allows you to add or change, or add and change VLAN tags on packets arriving on a specified Edge or Access link GE or 10GE interface. You can use a set of matching criteria to select which packets to transform. Or, packets can be transformed unconditionally.

See *Example: Application of Policy Map and Tag Action* (on page [232](#)) for an example of an Ethernet interface tag action and an interface QoS policy being applied to an Ethernet interface.

While tag actions are typically described as being applied to packets received (ingress traffic) on an Ethernet port interface, the reverse of the tag action is applied to packets transmitted (egress direction) from the interface. For example, if an Add-Tag action is applied to an interface, that same tag is stripped from packets leaving the interface.

The following VLAN tag actions can be assigned to an edge or access link interface within the networking domain:

- **Add Tag** - The "add-tag" action adds an outer tag either on all packets on an interface, or only on packets that match a specified VLAN ID or P-bit.
- **Add 2 Tags** - The "add-2-tag" action adds an outer tag and an inner tag with the specified VLAN IDs on packets of untagged traffic.
- **Change Tag** - The "change-tag" action changes the outermost tag to the specified VLAN ID on packets that match a specified VLAN ID.
- **Add and Change Tag** - The "add-and-ch" action changes the tag and then adds an outer tag on packets that either match a specified VLAN ID or are priority-tagged frames.

Note: Tag-actions can only be applied to edge or access link interfaces. To create tag actions for an ONT Ethernet or xDSL port subscriber service, you must use Service Tag Actions.

Configuration guidelines

- Tag actions can only be assigned to edge or access link interfaces within the networking domain and are defined with respect to ingress traffic.

- When you create a tag-action, the interface automatically becomes a member of the VLAN specified in the tag-action.
- When you create an add-tag action, the P-bit of the incoming tag is automatically copied into the P-bit of the added tag.
- You can configure only one of the following attributes on a given VLAN on a given interface:
 - VLAN membership
 - Tag-action (add, add-2-tags, change, or add-and-change)
 - Native VLAN
- To forward untagged traffic, an access interface must have an add-tag action, add-and-change, or add-2-tags action applied to untagged frames, assigning the traffic to a VLAN.
- An edge or access link interface can be associated with a particular VLAN either through a VLAN membership or through a tag-action, but not through both.
- Interfaces can be associated with additional VLANs through memberships or tag-actions.
- An interface role cannot be modified when it is assigned a tag-action.
- The E-Series can perform the following tag actions on a VLAN that has IGMP enabled:
 - Add-tag action to untagged Ethernet frames
 - Change-tag action to change the VLAN ID at network administrative boundaries
- If no matching criterion is assigned to a tag action, then the tag action is performed on all packets entering the specified interface, except for the packet traffic on VLANs that have an associated membership with the interface. When the match rule is set to ignore and p-bit any that is the same as no matching criterion assigned. You can see this if you create a tag action and leave those fields blank. It will be set to ignore.
- For double-tagged packets, the outer VLAN ID must match a provisioned VLAN. The inner tag does not have to match a provisioned VLAN ID.
- A priority tagged frame has a new tag added or changed-to with the specified tag value. Upon egress, the frame still carries the original priority bits onto the newly added or change-to VLAN tag.
- If an Ethernet interface has both a VLAN tag action and a policy map applied, the tag action function is applied to the interface before the policy map function.

Example add-tag-action configurations

Create an action to add a tag (VLAN 400) to all traffic on the associated interface.

Create

Create VLAN Tag Action

Associated Interface * 1-1-GE-1

Action Performed * add-tag

Matching Tag

Matching pbit

S-VLAN (Outer Tag) * 400

CREATE CANCEL

Create an action to add a second tag (VLAN 400) to all VLAN 401 traffic on the associated interface.

Create

Create VLAN Tag Action

Associated Interface * 1-1-GE-1

Action Performed * add-tag

Matching Tag 401

Matching pbit

S-VLAN (Outer Tag) * 400

CREATE CANCEL

Create an action to add a tag (VLAN 400) to all traffic based on the P-bit for the associated interface.

Create

Create VLAN Tag Action

Associated Interface * 1-1-GE-1

Action Performed * add-tag

Matching Tag ignore

Matching pbit pbit-6

S-VLAN (Outer Tag) * 400

CREATE CANCEL

Create an action to add a tag (VLAN 400) to all traffic based on the P-bit and outer tag on the associated interface.

Create

Create VLAN Tag Action

Associated Interface * 1-1-GE-1

Action Performed * add-tag

Matching Tag 401

Matching pbit pbit-6

S-VLAN (Outer Tag) * 400

CREATE CANCEL

Example application configuration for tag-action versus VLAN membership:

For incoming traffic where the voice and video services are single-tagged and the per-port-data VLANs are double-tagged:

- VLAN-per-port for data traffic
- Single voice service VLAN
- Single video service VLAN

Configure the following:

1. Create a voice VLAN and a video VLAN and add the interface to each of the VLAN memberships.
2. Create a no-match add-tag action to the interface for all remaining VLANs to add an outer data service tag.

Before starting

Before starting the VLAN tag action process, check that the following conditions are met:

- The VLANs must already exist in order to create tag actions.
- The interface that is the target for the tag action must have the edge or access role assigned.

Parameters

You can provision the following parameters for VLAN tag actions to an Ethernet interface:

Parameter	Description	Valid Options
Matching Tag	Specifies the match criteria. If you select Enter Value, then you must also specify the VLAN ID value to match. VLANs can be specified by name, or by numeric VLAN ID. (VLAN IDs 1002-1005 are reserved for E7 operation.)	ignore, any, untagged, priority, Enter Value (then enter a VLAN ID 2-4094)
Matching pbit	P-bit value that specifies the VLAN priority value to match.	pbit-0 to pbit-7, pbit-any, or leave blank
Action Performed*	Specifies whether to add, change, or add and change tags to incoming packets.	add-tag add-2-tags ch-tag add-and-ch
S-VLAN (Outer Tag)*	New VLAN ID. (VLAN IDs 1002-1005 are reserved for E7 operation.)	2-4093
C-VLAN (Inner Tag)	New VLAN ID. (VLAN IDs 1002-1005 are reserved for E7 operation.) Note: This parameter only applies if the Action Performed selection is "add-2-tags" or add-and-change.	ignore, any, untagged, Enter Value (then enter a VLAN ID 2-4093)

* Required field

GPON Traffic Management

This section describes E7 GPON traffic management capabilities, and includes the following topics:


- Overview of GPON Upstream Bandwidth Management
- Overview of the E7 GPON Upstream Traffic Flow
 - Dynamic Bandwidth Allocation (DBA)
 - DSCP-to-P-Bit Mapping
 - CoS-to-P-Bit Mapping
 - Selecting an E7 PON CoS Option
 - ONT Operation per E7 PON CoS Option
 - Applications for E7 PON CoS Options
- Overview of the E7 GPON Downstream Traffic Flow
 - Understanding Traffic Shaping and Policing
 - HSI Traffic Shaping and Policing Guidelines
- Queuing methods on ONT Ethernet ports
- Configuration guidelines and parameter descriptions for the following E7 GPON QoS features:
 - Ethernet Bandwidth Profiles
 - PON US CoS Profiles
 - PON CoS Map
 - DSCP Map
 - Match Lists
 - Service Tag Actions (selecting Match List, [Tag] Action, Outer/Inner Tags, and Outer/Inner P-Bit Source)

Overview of GPON Upstream Bandwidth Management

The Calix E7 GPON solution supports the ITU G.984 standard of 1.25 Gbps (1.244 Gbps) of upstream bandwidth on the PON. An E7 GPON card serves as an OLT, and is responsible for allocating upstream bandwidth to ONTs either as fixed allocated bandwidth when DBA is disabled, or dynamically (for non-fixed traffic classes) via the DBA algorithm when DBA is enabled.


On an ONT, the traffic-bearing entity that receives upstream BW allocations is called a Transmission Container (T-CONT). Different types of T-CONTs are used to prioritize traffic flows for different traffic classes of service (CoS) on the PON. All Calix ONTs support four industry-standard T-CONT types (CoS 1-4). On supported Calix ONTs, four additional, user-defined T-CONT types (User 1-4), which may use any forwarding type and PON CoS priority.

ONTs report the status of the queues associated with their T-CONTs to the OLT. Based on demand and the constraints of the DBA algorithm, the OLT facilitates sharing of the upstream bandwidth by allocating bandwidth grants to the T-CONTs. Excess bandwidth, either unconfigured or not currently required by a subscriber, is allocated to subscriber services based on the DBA priority order.



Overview of the E7 GPON Upstream Traffic Flow

This topic provides an overview of the E7 GPON upstream traffic flow and QoS features involved.



1. On an ONT Ethernet port, for ingress traffic to be accepted for a subscriber service, it must meet the pre-defined matching criteria for the service.
 - E7 GPON QoS Feature: **Match List (with Match Rules)** [selected in the Service Tag Action for a subscriber service]
2. Once frames meet the matching criteria, tag actions are performed to set the S-VLAN (outer tag), optionally the C-VLAN (inner tag), and the P-Bit(s).
 - E7 GPON QoS Features: **[Tag] Action, Outer/Inner P-Bit Source, Outer/Inner Tag** [selected in the Service Tag Action for a subscriber service]

Note: The outer P-Bit value may be directly specified as 0-7, mapped from the PON CoS Map, copied from an incoming tagged frame, or mapped from an incoming IP packet DSCP field. For more details, see *DSCP to P-Bit Mapping* (on page [248](#)) and *CoS to P-Bit Mapping* (on page [249](#)).

3. (Not shown) Once the tag actions have been performed, the tagged frames for all services are prioritized by P-Bit as they are transferred from the Ethernet MAC/switch to the GPON MAC of the ONT. For more details on the queuing technique used, see *Queuing Methods on ONT Ethernet Ports* (on page [264](#)).

Note: In E7 R2.2, the prioritization of frames is only accurate when the sum of the bandwidths specified for all services on a port is not exceeded. If the BW sum is exceeded, traffic is policed at the Ethernet port indiscriminately.

4. In the GPON MAC, for a given service, the frames are again priority queued as they are mapped to a GEM port for the service.
5. The frames for a given GEM port (or service) are then then classified and queued by T-CONT type, where T-CONT types are pre-defined for various PON traffic classes. Frames are mapped to T-CONT types (and PON traffic classes) depending on the PON CoS parameter for the subscriber service—mapped automatically based on P-bit value (using the PON CoS Map) or mapped by pre-selecting the T-CONT type for the service.
 - E7 GPON QoS Feature: **PON CoS > Derived, CoS 1-4, or User 1-4** [selected for the subscriber service]


Note: A given service must use only T-CONT type. Therefore, if a service contains multi-priority traffic, a specific T-CONT type must be selected, it should not be automatically mapped using PON CoS = Derived or Fixed. For more details on T-CONT types, see *Selecting an E7 PON CoS Option* (on page [250](#)).

6. On supported ONTs, rate-limiting is enforced per GEM port as traffic is mapped to a T-CONT type; otherwise, rate-limiting is enforced per T-CONT type via scheduling from the OLT.
 - E7 GPON QoS Feature: **Bandwidth Profile** [selected for the subscriber service]

Note: For ONTs where rate-limiting is only enforced per T-CONT type, if more than one service on an ONT is mapped to the same T-CONT type (same PON traffic class), the aggregate rate-limit of both services will be enforced. To isolate the bandwidth per service in this case, one of the services can be specifically mapped to a user-defined T-CONT type (a separate T-CONT type for scheduling purposes, but with the same PON traffic class). For more details, see *ONT Operation per E7 PON CoS Option* (on page [252](#)) and *Applications for E7 GPON CoS Options* (on page [255](#)).

7. Queued frames per T-CONT type are scheduled for transmission by the OLT using Dynamic Bandwidth Allocation (DBA).
 - E7 GPON QoS Feature: **DBA** (enabled by default)
8. (Not shown) The frames received at the OLT (via T-CONTs) are transferred from the GPON MAC to the Ethernet switch fabric, forwarded to the correct S-VLAN and Ethernet uplink, and prioritized by P-Bit in the transmit queue of the uplink interface.

The following diagram expands the view of E7 provisioning objects and parameters used for QoS. (For simplicity, only the parameters required for a data service are shown.)



After the requirements for your application have been determined, Calix recommends the following order of configuration. See the *Calix E7 GPON Applications Guide* for detailed configuration instructions. For configuration guidelines and parameter descriptions, see *Configuration Guidelines and Parameter Descriptions* (on page [266](#)).

1. BW Profile
2. PON US CoS Profiles and PON US CoS (for supported ONTs)
3. PON CoS Map customization, if required (default values recommended)
4. DSCP Map customization, if required (default values recommended)
5. Match List Rule(s) and Match List Profile
6. Service Tag Action, selecting Match List, [Tag] Action, Outer/Inner Tags, and Outer/Inner P-Bit Source
7. Subscriber Service, selecting BW Profile, PON CoS, and Service Tag Action

Dynamic Bandwidth Allocation (DBA)

Dynamic Bandwidth Allocation (DBA) is a feature that dynamically changes upstream GPON bandwidth assignments at the port level based on traffic class and demand. The GPON DBA feature is enabled by default.

DBA is a form of bandwidth management, and is essentially the same thing as statistical multiplexing, where the sharing of a Passive Optical Network (PON) adapts dynamically to traffic demands from the subscribers' use. Different network protocols implement DBA in different ways. These methods are typically defined in standards developed by standards bodies such as the ITU, IEEE, FSAN, or IETF. In the case of GPON, DBA is defined in specification ITU G.983.

Subscriber demands

Traffic in the Assured Forwarding classes is characteristically bursty, presenting the opportunity for re-assignment of its bandwidth allocation between bursts to lower priority services which indicate demand, such as those with Best Effort bandwidth. Additionally, some Best Effort traffic may experience light network load conditions, freeing up bandwidth for other high-demand Best Effort (non-guaranteed) subscribers actively using the PON up to their provisioned peak rates. If there is not enough bandwidth on the PON to meet their peak rates, DBA efficiently allocates available bandwidth between the subscribers that demand use of it, supporting over subscription.

Example: A carrier can provision all 32 ONTs with 1000M upstream best effort traffic, and subscribers can get this rate when the PON is not busy. As more subscribers request upstream capacity, the DBA algorithm will share the 1.2G bandwidth capacity among the ONTs requesting bandwidth grants.

DBA supports differentiating between guaranteed and non-guaranteed bandwidth provisioning, which is essential for business services. Businesses can purchase a guaranteed amount of bandwidth as well as an excess variable portion that they may use if it is available on the PON. DBA maximizes the use of the PON resources by utilizing unassigned, existing bandwidth in real-time, while also allocating unused guaranteed bandwidth to other subscribers for their excess or best effort traffic.

DBA for business services allows the operator to over provision the PON's upstream Peak Information Rate (PIR) bandwidth, while not over subscribing the Committed Information Rate (CIR) bandwidth. If an operator only sells CIR service, then DBA does not have as much value as most operators do not want to oversubscribe CIR traffic on the PON. However, the trend is to sell services with some component of CIR and PIR, in which case, the DBA algorithm performs the following sequence:

- a. Schedules the CIR (1.2Mbps max) component of traffic
- b. Returns to grant bandwidth requests for the PIR component of each service on the PON

Subscribers will be able to burst to their full PIR, while maintaining a CIR, allowing service providers to attain added revenue for the oversubscription of the PON.

Traffic class prioritization

Prioritization of traffic classes is also an important aspect of DBA, as not all subscribers and services are equal. For example, voice is more delay sensitive than video, yet video has more stringent guaranteed bandwidth requirements than data. Similarly, business services have more demanding SLAs than residential services and they pay higher rates for service guarantees. A number of traffic classes / priorities can be defined to meet customer and service requirements.

DBA uses the following underlying concepts and techniques to support the dynamic assignment of available bandwidth in a prioritized fashion:

- a. Traffic weighting allows a subscriber to receive a greater share of the excess bandwidth if they have been provisioned a larger relative amount of guaranteed bandwidth. If a business' subscribed guaranteed bandwidth is twice that of another, they would expect that it will receive twice the amount of excess bandwidth than the other. DBA's weighting mechanism supports this.
- b. Status reporting DBA uses queue depth information provided by the ONT's to make decisions on which ONT's to allocate bandwidth to at any given time. This enables DBA to grant bandwidth only to the ONT's that actually need it, thereby maximizing PON utilization. On the other hand, non-status reporting DBA, also called traffic monitored DBA, uses a historical approach to make its decision on which ONT's to allocate bandwidth. This requires DBA to always allocate some amount of bandwidth to each ONT in order to track its usage, thereby wasting PON resources.

How upstream bandwidth is allocated

The Calix E7 GPON solution supports the ITU G.984 standard of 1.25 Gbps (1.244 Gbps) of upstream bandwidth on the PON. An E7 GPON card serves as an OLT, and is responsible for allocating upstream bandwidth to ONTs either as fixed allocated bandwidth when DBA is disabled, or dynamically (for non-fixed traffic classes) via the DBA algorithm when DBA is enabled.

On an ONT, the traffic-bearing entity that receives upstream BW allocations is called a Transmission Container (T-CONT). Different types of T-CONTs are used to prioritize traffic flows for different traffic classes of service (CoS) on the PON. All Calix ONTs support four industry-standard T-CONT types (CoS 1-4). On supported Calix ONTs, four additional, user-defined T-CONT types are also available (User 1-4), which may use any forwarding type and PON CoS priority.

ONTs report the status of the queues associated with their T-CONTs to the OLT. Based on demand and the constraints of the DBA algorithm, the OLT facilitates sharing of the upstream bandwidth by allocating bandwidth grants to the T-CONTs. Excess bandwidth, either unconfigured or not currently required by a subscriber, is allocated to subscriber services based on the DBA priority order.

Status Reporting (SR) and Non-Status Reporting (NSR)

How the OLT learns that an ONT (i.e. subscriber) is trying to send traffic upstream depends on whether the DBA is Status Reporting (SR) or Non-Status Reporting (NSR).

- In **Status Reporting**, the ONT knows the “status” of its queues and can therefore send an accurate bandwidth request to the OLT. The OLT then grants an upstream time slot to the ONT, and the ONT sends the subscriber traffic upstream (a closed-loop system).
- In **Non-Status Reporting**, the OLT infers the status of the ONT queues solely by looking at the arrival of traffic from the ONT. Therefore, the NSR mechanism is not closed-loop, resulting in less efficient use of upstream bandwidth and more delay for low bit rate, time-sensitive traffic.

Configuration guidelines

- A GPON port is specified by card number and port number. For example, 1/1.
- Each GPON-4 card provides four GPON ports that subtend up to 64 ONTs each, for a card capacity of 256 ONTs.
- Each GPON-8 card provides eight GPON ports that subtend up to 64 ONTs each, for a card capacity of 512 ONTs.
- Each GPON-16x card provides 16 GPON ports that subtend up to 128 ONTs each, for a card capacity of 1024 ONTs.
- When removing all fibers from one card, wait two minutes before inserting the fibers into another card.
- The E7 GPON card supports up to 32 IGMP-enabled multicast VLANs on the system, which can all be on the PON.

- Dynamic Bandwidth Allocation (DBA) operation
 - DBA is set at the GPON port level, and is enabled by default
 - DBA supports up to 64 ONTs per PON.
 - DBA-managed traffic includes the Assured Forwarding and Best Effort traffic class types only.
 - The DBA algorithm operates on the configured bandwidth profiles for each subscriber and service, and the perceived demand from the ONTs.
 - All traffic in the DBA pool (Best Effort and Assured Forwarding classes) is re-allocated to the pool of available bandwidth when the service is not using its provisioned bandwidth.
- GPON bandwidth guidelines
 - GPON has a total upstream bandwidth of 1.2 Gbps.
 - Any one service can be provisioned up to 1 Gbps.
 - Expedited Forwarding bandwidth is limited to 600 Mbps, and is not part of the DBA pool of dynamically allocated bandwidth.
 - Committed Information Rate (CIR) traffic in the Assured Forwarding class is limited to 1.2 Gbps.
 - Assured Forwarding services scheduled by DBA have a minimum bandwidth allocation of 64 kbps per service, and a maximum of 1 Gbps per service.
 - Best Effort services scheduled by DBA have a maximum bandwidth allocation of 1 Gbps per service.
- E7 GPON ports have no interface association, unlike the E7 Ethernet ports.
- For traffic on a specific VLAN to be forwarded to and accepted on an E7 GPON port, the particular VLAN must be provisioned on the E7 and a service tag action specifying the VLAN must be included in the ONT service provisioning.
- In order to fully comply with the network design framework outlined by GPON C+ (ITU-T G.984.2 Amendment 2) a method for correcting errors over a single-path optical transmission line is used called Forward Error Correction (FEC). The E7 GPON port parameters of Upstream/Downstream FEC can be set to "auto" so the FEC is automatically enabled when a C+ GPON Optical Interface Module (OIM) is inserted into the GPON port socket for OLTs capable of supporting C+/FEC: E7-20 GPON-8x and GPON-16x, E7-2 GPON-8 and GPON-4r2.
- Class C+ optical interface modules can also be used in non-FEC capable E7 OLTs (E7-2 GPON-4 and E7-20 GPON-4x) to provide an extension of the B+ or Reach Extended B+ ODN reach, but FEC is required to comply with the full extent of C+ as defined in ITU-T G.984.2 Amendment 2, combining both the added optical power of the C+ optics, and the error correction algorithm provided by FEC.

- The E7-20 GPON-16x line cards introduced in software version R3.1 do not support PON hairpin on configured service VLANs. The following primary scenarios may be affected: 1) PWE3 endpoints on the same PON will not work, 2) Direct media RTP traffic between voice endpoints on the same PON will not work.

DSCP to P-Bit Mapping

For upstream ingress traffic on ONT Ethernet ports, the E7 provides the ability to translate IP packet DSCP values in untagged Ethernet frames to P-bits.

For an overview of IP packet DSCP values, see *DSCP Map* (on page [280](#)).

Configuration rules and guidelines

- The DSCP-to-P-Bit mapping operation is performed in the upstream direction only on the ONT.

Note: DSCP-to-P-Bit mapping for downstream traffic may be performed on E7 Ethernet uplink ports using Class Maps, Class Rules, and Policy Maps.

- The DSCP-to-P-Bit mapping operation is performed on untagged frames, and works with the "AddTag" and "Add2Tags" tag actions.
- Only one service per port can use the DSCP-P-bit mapping function. (An ONT Ethernet port supports only DSCP Map, and only one service on the port can use a “match all untagged” classification rule.)
- If another, second service on the same port has an untagged traffic match rule to match a specific OUI value (not supported for T-Series ONTs), traffic matching this OUI value will be classified into the second service with higher precedence than the “match all untagged” rule used to translate DSCP values to Pbits.
- GPON CoS considerations:
 - Any service using DSCP-to-P-Bit mapping must also be configured to use a specific T-CONT type, via an explicit PON CoS choice (not "Derived" or "Fixed.")
 - After translation from DSCP to P-Bit, upstream traffic for the service is prioritized by P-Bit as it is mapped to a single queue (or GEM port) on the ONT for upstream scheduling by the OLT.
 - In the event of congestion on the upstream bandwidth allocation, traffic in the queue with highest P-bit is transmitted first, thus preserving the DSCP priority set by the downstream CPE.

Configuration process

1. Create a DSCP Map (DSCP to P-Bit map).

Caution: Calix recommends to NOT assign any DSCP value to P-bit 7. P-bit 7 traffic from subscribers have the potential to overrun network control protocols (ERPS PDUs, RSTP BPDUs, etc.).


2. Assign the DSCP Map to an ONT Ethernet port.

Note: Each ONT port is independent and can have a different DSCP Map assigned.

3. Create a Service Tag Action with an Outer P-Bit Source of "Map a layer-3 Priority."
4. Create a subscriber service using the above Service Tag Action and using an explicit PON CoS (not "Derived" or "Fixed.")

Example diagram

In the example below, a data subscriber service is configured to add an outer tag with VLAN ID 100 and P-Bit value mapped from DSCP value. (Parameters/operations for rate-limiting not called out here.)



CoS to P-Bit Mapping

For upstream ingress traffic on ONT Ethernet ports, the E7 provides the ability to set the P-Bit value according to a pre-defined map of PON CoS to P-bit values.

Configuration process

1. Confirm that the E7 PON CoS map is defined correctly for your network (default values recommended).
2. Create a Service Tag Action with an Outer P-Bit Source of "From CoS" and select the desired CoS (1-4).

Note: During subscriber service creation, there is a "PON CoS" field, which has a different application.

3. Create a subscriber service using the above Service Tag Action.
4. Upstream traffic for the service will have P-Bits set according to the CoS-to-P-Bit mapping specified in the E7 PON CoS map.

Note: Where there is a P-Bit range per CoS, the lowest P-Bit value in the range will be used. For example, for CoS 1, the default P-Bit range is 0-2. If "From CoS > CoS 1" is selected for the Outer P-Bit Source, the P-Bit value on frames will be set to 0.

Selecting an E7 PON CoS Option

To support prioritized upstream service flows, the E7 GPON solution supports up to eight T-CONTs types per ONT. Four types (cos1-4) follow a system-default P-bit to CoS mapping model. For added flexibility, four other types (user1-4), supported by compatible ONTs, are user-configurable for any class of service. The combination of these eight T-CONTs supports multiple instances of the same traffic class per ONT to be treated independently for rate-limiting and shaping.

Note: The PON CoS options "derived" through "user-4" support the new GPON QoS behavior and features available in R2.2+.

Note: T-Series ONTs are supported in E7 R2.3+.

PON Class of Service options for an Ethernet service

PON CoS	Forwarding Type	Bandwidth	Description
derived	Derived from the priority of the incoming frames after tag actions have been performed, mapping to cos-1 to 4.	Aggregated for all services with the same PON CoS	With this option, incoming frames are mapped to cos-1 to 4 (T-CONT types) based on the P-Bit value indicated in the associated service-tag action. Note: This option cannot be used with P-Bit preservation and promotion. ONT operational notes: For P-Series GX ONTs: Multiple services with the same traffic class are rate-limited and shaped independently by GEM port, isolating the bandwidth for each service before sharing the same T-CONT. For P-Series GE and T-Series ONTs: Multiple services with the same traffic class have unique GEM ports, but are not rate-limited independently before sharing the same T-CONT. To achieve independent rate-limiting, the second service of the same traffic class should use user-1 to x to carry the service on a separate T-CONT.
cos-1	Fixed Best-Effort	Aggregated for all services associated with this system defined CoS level	These options represent standard T-CONT types with default PON CoS forwarding types. Incoming frames are mapped to the T-CONT type specified (cos-1 to 4). The provisioned services are required to have a bandwidth profile that matches the class of service. Bandwidth is assigned as aggregated from the multiple services and mapped to the ONT. ONT operational notes: Same as above.
cos-2	Fixed Assured Forwarding		
cos-3	Fixed Assured Forwarding		
cos-4	Fixed Expedited Forwarding		

user-1	Any	Aggregated or Settable. If Aggregated, the rate is derived from all services associated with the system defined CoS level.	These options represent user-defined T-CONT types that can be configured with any traffic class forwarding type. For compatible ONTs, this allows multiple services with the same traffic class to be carried by separate T-CONTs on the PON. Incoming frames are mapped to the T-CONT type specified (user-1 to 4). ONT compatibility: - For GX ONTs: Not supported. - For GE ONTs: user-1 to 4 are supported. - For T-Series ONTs: user-1 to 3 are supported.
user-2			
user-3			
user-4			
fixed	Derived from the priority of the incoming frames after tag actions have been performed, mapping to cos-1 to 4.	Aggregated for all services with the same PON CoS	This option represents the GPON QoS behavior and features prior to R2.2, and is intended for backwards compatibility with those earlier-provisioned services. With this option, incoming frames are mapped to cos-1 to 4 (T-CONT types) based on the P-Bit value indicated in the associated service-tag action. Note: This option cannot be used with T710G/T720G ONTs. Note: This option cannot be used with GPON-16x line cards when two VLANs on the same PON port are sharing the same bandwidth profile. Note: This option cannot be used with P-Bit preservation (outer P-bit set to "copy"). Per ONT Ethernet port, this setting provides four GEM ports for services, one for each system CoS (1-4) value. Multiple services that map to the same CoS share the same GEM port queue and aggregate bandwidth, where any one service can use all of the bandwidth, if there is no contention for it.

Quick QoS application examples:

- To support VoD over the top without rate limiting (where the VoD traffic P-Bit value is not explicitly defined by the service):
 - * Use "fixed."
- To rate limit two BE services independently:
 - * On GX SFU/MDU ONTs, use "derived" or "cos-1" for both services. The two services will be rate-limited independently by GEM port before sharing the same T-CONT.
 - * On GE ONTs, use "cos-1" (or "derived") for one service and "user-1" for the second service. Each service must be mapped to a different T-CONT.

ONT Operation per E7 PON CoS Option

The operational architecture of GPON ONTs differs based on the E7 PON CoS Option selected during subscriber service creation; specifically, if the "Fixed" option is selected versus any other non-Fixed option (Derived, CoS1-4, or User1-4).


This is illustrated below for P-Series 700GE ONTs and 700GX ONTs, with the simple use case of two data services mapping to the same traffic class.

PON CoS = Fixed

The "Fixed" PON CoS option has only been provided for backward compatibility with configurations in E7 R2.1. With the "Fixed" PON CoS option, ONTs operate as follows:

- Per ONT Ethernet port, four GEM ports are available for services; and per ONT, four T-CONT types are available. For two services that map to the same traffic class (sharing the same GEM port and T-CONT type), there is no bandwidth isolation per service.

The architecture within a 700GE ONT is shown below. (The 700GX ONT architecture is the same, except for a traffic shaper per GEM.)




E7 R2.2 with GE ONTs (fixed mode)

PON CoS = Non-fixed


If any non-fixed option is selected (Derived, CoS1-4, User1-4), ONTs will operate as follows:

- For 700GE ONTs, per Ethernet port, eight GEM ports are available for services; and per ONT, a total of eight T-CONT types are available (four standard, and four user-defined). For two services that map to the same traffic class, a user-defined T-CONT type may be selected for one of the services in order to achieve bandwidth isolation.
- For 700GX ONTs, per Ethernet port, eight GEM ports are available for services; and per ONT, the four standard T-CONT types are available. For two services that map to the same traffic class, individual shaping per GEM port (prior to mapping to a single T-CONT type) achieves bandwidth isolation.

The architecture within 700GE and 700GX ONTs is shown below.



E7 R2.2 with GE ONTs (non-fixed mode)



E7 R2.2 with GX ONTs (non-fixed mode)

Applications for E7 GPON CoS Options

The E7 supports the following GPON CoS applications:

- Multiple services, same traffic class (CoS), shared BW
- Multiple services, same traffic class (CoS), independent BW
- Multiple services, same traffic class (CoS), different P-Bits
- Multiple services, same traffic class (CoS), independent BW, prioritized scheduling on the PON
- Single service with prioritization of flows within the service

Multiple services, same traffic class (CoS), shared BW

This model is intended for more than one service being configured on a VLAN, where any service is with no explicit PON upstream CoS mapping, and is backwards compatible and upgrade compatible. Services provisioned with software version earlier than R2.2 will operate the same after upgrading to R2.2 or later.

This model assigns four T-CONT types, each representing an aggregated class of service for an ONT (BE, AF1, AF2, EF). The four T-CONT types are pre-assigned a class of service and the provisioned services are required to have bandwidth profiles that match the class of service. Also, bandwidth assigned to the T-CONT type is always aggregated from the multiple services mapped to it on the ONT.

Each ONT Ethernet port receives four upstream queues (or GEM ports) at service instantiation where the outer P-bit of the tag action maps to one of the four GEM ports, regardless of the VLAN. If two services are assigned to a single class of service, they will share the GEM port. A second service on a different ONT Ethernet port will be assigned a different GEM port, but will share the T-CONT type with service separation performed by the ONT GEM port shaper. The mapping between the GEM port and the T-CONT type is fixed for all services on the OLT and is defined by the E7 PON Class of Service configuration attribute.

Service Name	Traffic Class	T-CONT Type (PON CoS)	Downstream Bandwidth (PIR)	Upstream Bandwidth (CIR/PIR)	GEM Port	VLAN ID / P-Bit
Data-1	AF1	cos-2	10 Mb/s	1/1 Mb/s	GEM 3	30 / 3
Data-2	BE	cos-1	50 Mb/s	0/5 Mb/s	GEM 1	32 / 0
Data-3	BE	cos-1	10 Mb/s	0/1 Mb/s	GEM 1	34 / 0

- In Standard QoS Model T-CONT type (CoS-1):
 - Upstream BWP = Data-2 + Data-3 BWPs = 6 Mb/s
 - Downstream BWP = 60 Mb/s
 - Either service can use all of the bandwidth, if there is no contention for it.

- Traffic marked with a P-bit value is forced into a pre-determined upstream traffic class and bandwidth profile type.
- The fixed PON CoS service model does not support the following:
 - P-bit preservation or DSCP-to-P-bit mapping
- Ability to rate limit all P-bit flows within a single VLAN

Multiple services, same traffic class (CoS), independent BW

This model is intended for isolating the bandwidth for multiple services on an ONT Ethernet port that share the same traffic class. The same service class per ONT port can select different GEM ports and a different T-CONT type.

- PON US profile is created and mapped to user-1 at the ONT
- PON CoS of user-1 is selected when adding the Ethernet service

Service Name	Traffic Class	T-CONT Type (PON CoS)	Downstream Bandwidth (PIR)	Upstream Bandwidth (CIR/PIR)	GEM Port	VLAN ID / P-Bit
Data-1	AF1	cos-2	10 Mb/s	1/1 Mb/s	GEM 3	30 / 3
Data-2	BE	cos-1	50 Mb/s	0/5 Mb/s	GEM 2	32 / 0
Data-3	BE	user-1	10 Mb/s	0/1 Mb/s	GEM 1	34 / 0

Multiple services, same traffic class (CoS), different P-Bits

This model is intended for isolating multiple services on an ONT Ethernet port that share the same traffic class, where a service explicitly defines the outer VLAN, P-bit, and the PON upstream traffic class. The service can have any bandwidth profile, or share a single bandwidth profile with other services.

The specified P-Bit is not required to match one of the values defined for the CoS. The P-bit 0 (BE) service uses Software Egress Shaping Path when BE Rate < MAX Supported throughput.

- PON US profile is created and mapped to user-1 at the ONT
- Service match list has tagged rule for VLAN 34, P-bit 3
- Service tag action to use P-bit 3
- PON CoS of user-1 is selected when adding Ethernet service

Service Name	Traffic Class	T-CONT Type (PON CoS)	Downstream Bandwidth (PIR)	Upstream Bandwidth (CIR/PIR)	GEM Port	VLAN ID / P-Bit
Data-1	AF1	cos-2	10 Mb/s	1/1 Mb/s	GEM 3	30 / 3
Data-2	BE	cos-1	50 Mb/s	0/5 Mb/s	GEM 2	32 / 0
Data-3	BE	user-1	10 Mb/s	0/1 Mb/s	GEM 1	34 / 3

Multiple services, same traffic class (CoS), independent BW, prioritized scheduling on the PON

This model is intended for more than one service of the same class where one service (AF2) should have a higher priority than another (AF1), and they are configured as isolated services. The same service class per ONT port can select different GEM ports, a different T-CONT type, and priority of flow.

- PON US profile is created with priority 3 and mapped to user-1 at the ONT
- Service match list has tagged rule for VLAN 34, P-bit 4
- Service tag action to use P-bit 4
- PON CoS of user-1 is selected when adding Ethernet service

Service Name	Traffic Class	T-CONT Type (PON CoS)	Downstream Bandwidth (PIR)	Upstream Bandwidth (CIR/PIR)	GEM Port	PON Cos Priority	VLAN ID / P-Bit
Data-1	AF1	cos-2	1000 Mb/s	100/100 Mb/s	GEM 2	2	30 / 3
Data-3	AF2	user-1	1000 Mb/s	100/100 Mb/s	GEM 1	3	34 / 4

The recommended priority value of '3' will ensure AF2 is scheduled before AF1.

Single service with prioritization of flows within the service

This model is intended for a service that requires QoS within a VLAN and explicitly defines the PON upstream traffic class. This model allows a single service to be defined with P-bit preservation, while maintaining quality of service within the VLAN. Typically, this model is used within a TLAN or business service. There is a single rate for the entire VLAN service that can have any type of bandwidth profile (best-effort, assured, expedited). Also, a user-defined PON upstream CoS should be assigned to ensure service separation.

- Service tag action to preserve the P-bit value set in service provider supplied Residential Gateways (RGs):
 - Copies the subscriber-side Q-Tag P-Bits to a single C-tag or S-Tag, or to the inner C-Tag of a double-tagged Service.
 - For the outer tag P-bit, you can specify the value of "copy."
 - For the inner tag, you can specify the value of "preserve."

Service Name	Traffic Class	T-CONT Type (PON CoS)	Downstream Bandwidth (PIR)	Upstream Bandwidth (CIR/PIR)	GEM Port	Flow Type	VLAN ID / P-Bit
video-1	AF	cos-3	50 Mb/s	1/3 Mb/s	Traffic priority queued prior to mapping to a GEM port -- for example, GEM 3.	DHCP	30 / 0
						VoD C&C	32 / 2
						VoD	32 / 4
						MC	32 / 5

Overview of the E7 GPON Downstream Traffic Flow

This topic provides an overview of the E7 GPON downstream traffic flow and QoS features involved.

1. From the E7 GPON OLT, downstream tagged traffic is scheduled for transmission based on priority (P-Bit value), and broadcast on the PON.
2. Downstream traffic per GEM port (or service) is rate-limited separately, policed from the OLT to the PON interface. However, the policing rate for high speed internet (HSI) service is relaxed by 100 Mbps if an ONT is shaping the traffic.

From the ONT PON interface toward the ONT Ethernet port, traffic for a service is either shaped (by the ONT) or policed (by the OLT). For more details, see *Data Traffic Shaping and Policing Guidelines* (on page [260](#)).

- E7 GPON QoS Feature: **Bandwidth Profile** [selected for the subscriber service]


3. In the ONT Ethernet MAC, traffic is queued again by P-Bit, and tag actions are reversed as frames egress the ONT Eth port.
 - E7 GPON QoS Features: **[Tag] Action, Outer/Inner P-Bit Source, Outer/Inner Tag** [selected in the Service Tag Action for a subscriber service]

Understanding Traffic Shaping and Policing

Traffic shaping and policing are defined as follows:

- **Traffic shaping** buffers packets arriving at a rate in excess of the provisioned maximum rate and schedules the packets for transmission at the provisioned shaped rate. The result of traffic shaping is a smoothed packet flow rate, with limited packet loss, but an increase in packet delay.
- **Traffic policing** drops traffic exceeding the provisioned maximum rate limit, with a result that bursting traffic may appear saw-toothed with "flattened" crests and troughs. Packets are not subject to delay.

Shaping is used to manage traffic that crosses into reduced interface speeds (for example, 2.5 Gbps GPON into a 20 Mbps service rate), or to smooth traffic when another mechanism is managing traffic congestion (for example, TCP end-to-end flow control).




Throughput, the maximum speed at which data traverses a system, is affected by provisioned bandwidth limits and what rate-limiting method (shaping or policing) is used. While lab testing of throughput may be measured per RFC 2544 or ITU-T Y.1564, end subscribers may use TCP or UDP bandwidth performance monitoring tools. For single-subscriber Internet service below 200 Mbps, shaping provides more accurate results with TCP-based Ookla speed testing.

Note: When using the traffic policing model, P-Bit priority is ignored since packets are not queued. Setting the P-Bit value in a traffic shaping environment will have its desired effect.

Data Traffic Shaping and Policing Guidelines

For data services, the E7 provides shaping and policing functions on Calix P-Series 700GE/740GE/760GX ONTs and the 836GE RSG. When specific criteria are met for a data service (described in "Applying policing and shaping," below), shaping is automatically enabled for a single Class of Service on a single ONT Ethernet port for provisioned rates up to 1000 Mbps.

The following diagram summarizes the traffic shaping and policing behavior of P-Series GPON ONT/RSGs.



Notes:

- When the ONT/RSG is shaping downstream traffic, the E7 GPON subsystem relaxes (by 100 Mbps) the per-service policing normally applied to all services in the downstream direction. This results in the E7 policing traffic grossly exceeding the maximum rate, while the ONT/RSG shapes typical traffic bursts.
- The E7 OLT schedules all traffic upstream to ensure efficient and prioritized use of GPON bandwidth. Traffic is buffered in the ONT/RSG for periodic transmission, resulting in shaped traffic for all upstream traffic.

Detailed information on shaping/policing support per configuration, shaping and policing criteria, and speed-testing guidelines are provided below.

Traffic management recommendations

Calix recommends the following guidelines to achieve consistent and accurate performance for data services delivered from a Calix P-Series GPON ONT.

- Internet access service below 250 Mbps – traffic shaping
- L2 VPN / Transparent LAN (MEF E-LINE and E-LAN) services – traffic policing
- Very high speed Internet access (> 250 Mbps) – traffic policing

Applying shaping and policing

Consistent with the above recommendations, by default, the E7 GPON system will apply traffic shaping to a single internet access data service per ONT (select models listed below) when the following conditions are met:

- Shaping functionality enabled for E7 GPON
 - In the E7-level Work Area under **System > Provisioning**, the **Enable ONT Downstream Shaping** checkbox is selected (default setting is selected/enabled)
 - This system-level provisioning enables (or disables) traffic shaping on all ONTs on the E7 OLT.
- Full-bridge mode is not configured
- Bandwidth profile Peak Information Rate (PIR) < 1000 Mbps
- Data traffic is mapped to the lowest GPON CoS compared with all other services on the ONT/RSG
- VLAN settings
 - IGMP = flood (i.e. IGMP snooping disabled)
 - DHCP Snoop enabled
 - MAC FF enabled
 - IP Source Verify enabled
 - TLAN disabled

If all of the above criteria are met for a single* data service, then traffic shaping is applied. Otherwise, traffic policing is used by the system to rate limit data traffic flows.

***Note:** If a second data service on the ONT/RSG also meets the above criteria, the service on the lowest numbered ONT Ethernet port is shaped.

Traffic shaping on P-Series ONTs

700GE and 740GE ONTs, and 836GE RSGs support both traffic shaping and traffic policing.

Downstream traffic

Downstream internet access traffic is shaped or policed based on the criteria above and the ONT model.

Provisioned Service BW	700GE/836GE (L2 service)	700GE/836GE (L2 full-bridge)	700GE/836GE (L3 gateway)	740GE (L2 service)
0–250 Mbps	Traffic shaped	Traffic policed	Traffic shaped	Traffic shaped
250–1000 Mbps	Traffic policed	Traffic policed	Traffic policed	Traffic policed

Shaping for data services above the capabilities of the GPON system should be provided by a router/BNG upstream of the E7 OLT.

Upstream traffic

In a GPON FTTP system, upstream traffic is always shaped due to the mechanics of the shared N:1 GPON transport layer. Upstream traffic is buffered in the ONT and scheduled for transmission via the Status Reporting - Dynamic Bandwidth Algorithm (SR-DBA). By default, each service on the PON is mapped to a different GPON CoS and rate limited to individual limits.

Other ONTs

The 700GX and 760GX ONT families support a fixed traffic management capability.

- 760GX ONTs provide policing for all data services
- 700GX ONTs provide shaping for all services up to a 400 Mbps PIR


Queuing Methods on ONT Ethernet Ports

This topic describes the queuing methods on ONT Ethernet ports for the P-Series 700GE ONTs and the 836GE RSG.

For E7 R2.1, E7 R2.2, and AE R2.2, Calix employs a hybrid queuing method derived from two popular queuing technologies: "Strict Priority" (SP) and "Weighted Round Robin" (WRR).

WRR ensures that all queues are serviced during each cycle with a queuing algorithm used to rotate service among the eight possible queues (p-bit values). The rotation is based on the weights assigned to each queue. This method rotates service among the queues, forwarding a specific number of packets in one queue before moving on to the next. WRR uses a default set of queue weights.

An example of packet handling using the weighted round robin method is shown below:



6684

Depending on the platform, the following queue weight factors are used by default:

AE 2.2 / E7 2.1 / E7 2.2 Queue Weights		
Queue	P-bit	Queue Weight
0	0	1
1	1	2
2	2	4
3	3	8
4	4	16
5	5	32
6	6	64
7	7	128
AE 3.x Queue Weights		
Queue	P-bit	Queue Weight
0	0	1
1	1	1
2	2	1
3	3	8
4	4	32
5	5	49
6	6	49
7	7	49
Note: The lower the weight number, the lower the priority.		

When using queue 7, Strict Priority (SP) is always implemented by the ONT. SP queuing processes as many possible packets in a queue before moving on to the next lower queue. In other words, when using Queue 7 on the ONT, the ONT schedules traffic in this queue first and then processes the remaining queues. For this reason, Calix strongly recommends NOT using P-bit 7 for any services provisioned on the ONT.

IMPORTANT: Effective with E7 Release 2.3 and AE Release 3.1, GE ONTs use Strict Priority (SP) queuing for all queues.

Configuration Guidelines and Parameter Descriptions

This section includes the configuration guidelines and parameter descriptions for the following E7 GPON QoS features:

- Ethernet Bandwidth Profiles
- PON US CoS Profiles
- PON CoS Map
- DSCP Map
- Match Lists
- Service Tag Actions (selecting Match List, [Tag] Action, Outer/Inner Tags, and Outer/Inner P-Bit Source)

Ethernet Bandwidth Profile

An Ethernet bandwidth profile specifies the upstream and downstream bandwidth rates to apply to individual Ethernet services on ONT Ethernet ports. Typically, a single bandwidth profile is applied to many subscriber ports.

Each direction of a traffic shaper can be configured to provide a specific rate for the shaper. The following values are specified to designate the rate:

- **Committed Information Rate (CIR):** The minimum guaranteed rate for traffic that the ONT provides for the particular Ethernet service. CIR defines the average rate in bytes/s of packets up to which the network delivers and meets the performance objectives defined by the CoS Service Attribute. CIR MUST be ≥ 0 .
- **Committed Burst Size (CBS):** CBS limits the maximum number of bytes available for a burst of packets sent at the UNI speed to remain CIR-conformant.
- **Peak Information Rate (PIR):** The PIR value is the maximum bandwidth rate for upstream and downstream traffic that the ONT allows the service to use if bandwidth is available. PIR defines the average rate in bytes/s of packets up to which the network may deliver packets but without any performance objectives. PIR MUST be ≥ 0 .
- **Peak Burst Size (PBS):** PBS limits the maximum number of bytes available for a burst of packets sent at the UNI speed to remain PIR-conformant.

The PBS and the CBS and are measured in bytes and it is recommended that they be configured to be equal to or greater than the size of the largest possible IP packet in the stream.

Note: The CMS Global ONT Ethernet Bandwidth Profile includes parameters for Committed Rate for Downstream and Committed Burst Size for Downstream, however, those values are only applicable to AE ONT configurations and E7-20 GPON-16x, NOT E7 GPON.

The CIR and PIR values must be consistent with the following:

- The Class of Service (CoS) type as defined in the GPON CoS global configuration table or
- The P-bit being used, whether it is specified in the service tag action, or when provisioning the service if the service tag action referenced the "use in Service" option.

Dynamic Bandwidth Allocation (DBA) is a feature that dynamically changes GPON bandwidth assignments at the port level based on traffic class and demand.

If explicit CIR/PIR values are included for a PON CoS profile, these values take precedence over the Ethernet bandwidth profile associated to the service. Additionally, the downstream CIR is only implemented at the downstream E7-20 GPON-16x PON.

Examples:

- If the user-defined PON CoS (user 1-4) CIR is 0, there will be no guaranteed bandwidth, even if the service has set a CIR via the Ethernet bandwidth profile.
- If the user-defined PON CoS (user 1-4) PIR is less than the Ethernet bandwidth profile PIR, the service will only get the PON CoS PIR as the peak rate.

See GPON QoS and Creating a PON Upstream Class of Service Profile.

Because PON networks carry voice, video, and data traffic with different Quality of Service (QoS) requirements, including controlled latency and jitter (packet delay and packet delay variation), and loss characteristics, as well as other specifications defined by Service Level Agreements (SLAs). The total GPON bandwidth (upstream=1.2 Gbps, downstream=2.4 Gbps) is shared among the subscribers and mapped to the following commonly-defined Per-Hop Behavior (PHB) traffic classes:

- Default PHB (Per hop behavior)—typically best-effort traffic
Any traffic that does not meet the requirements of any of the other defined classes is placed in the default PHB. Typically, the default PHB has best-effort forwarding characteristics.
- Expedited Forwarding (EF) PHB—dedicated to low-loss, low-latency traffic
The IETF defines Expedited Forwarding behavior in RFC 3246. The EF PHB has the characteristics of low delay, low loss and low jitter. These characteristics are suitable for voice, video and other real-time services. EF traffic is often given strict priority queuing above all other traffic classes.
- Assured Forwarding (AF) PHB—gives assurance of delivery under prescribed conditions
The IETF defines the Assured Forwarding behavior in RFC 2597 and RFC 3260. Assured forwarding allows the operator to provide assurance of delivery as long as the traffic does not exceed some subscribed rate. Traffic that exceeds the subscription rate faces a higher probability of being dropped if congestion occurs.

The AF behavior group defines four separate AF classes. Within each class, packets are given a drop precedence (high, medium or low). The combination of classes and drop precedence yields twelve separate DSCP encodings from AF11 through AF43. Usually, traffic policing is required to encode drop precedence. Typically, all traffic assigned to a class is initially given a low drop precedence. As the traffic rate exceeds subscription thresholds, the policer will increase the drop precedence of packets that exceed the threshold.

- Class Selector PHBs—maintains backward compatibility with the IP Precedence field

Traffic Class Type	GPON CoS	E7 P-bits (default)	Service Category	E7 BW Profile	DBA Priority For Non-Fixed Bandwidth
Expedited Forwarding (EF)	4	5, 6, 7	Network control, voice, T1/E1 Example: CIR = 384k PIR = 0 P-bit = 5	CIR>0, PIR=0 (PIR is provisioned for "0" to represent that there is no PIR independent of the CIR) Committed Information Rate (CIR): Guaranteed level of bandwidth with low Delay and Delay Variation.	Fixed Bandwidth not subject to DBA scheduling (bandwidth is pre-scheduled and not re-allocated to other services when not used)
Assured Forwarding 1 (AF1)	3	4	Video Example: CIR = 512k PIR = 1M P-bit = 4	CIR>0, PIR>=CIR Committed Information Rate (CIR) + Peak Information Rate (PIR). Both guaranteed and non-guaranteed bandwidth levels. Moderate Delay and Delay Variation.	Committed Bandwidth + Excess Bandwidth Priority 1 (Highest Priority)
Assured Forwarding 2 (AF2)	2	3	Application signaling, TLAN	CIR>0, PIR>=CIR Committed Information Rate (CIR) + Peak Information Rate (PIR). Both guaranteed and non-guaranteed bandwidth levels. Moderate Delay and Delay Variation.	Committed Bandwidth + Excess Bandwidth Priority 2
Best Effort (BE)	1	0, 1, 2	High-speed internet Example: CIR = 0 PIR = 10M P-bit = 0	CIR=0, PIR>0 Peak Information Rate (PIR): No guaranteed bandwidth levels. Higher Levels of Delay and Delay Variation.	Best Effort Bandwidth Priority 3 (Lowest Priority)

Traffic in the Assured Forwarding classes is characteristically bursty, presenting the opportunity for re-assignment of its bandwidth allocation between bursts to lower priority services which indicate demand, such as those with Best Effort bandwidth. Additionally, some Best Effort traffic may experience light network load conditions, freeing up bandwidth for other high demand Best Effort subscribers. SR-DBA efficiently allocates available bandwidth, and supports oversubscription—significantly improving network performance, and allowing service providers to provide high bandwidth services to many subscribers.

Applying Ethernet bandwidth profiles

When creating services on ONT Ethernet ports, typically the service-tag action identifies an explicit P-bit. The P-bit determines the PON CoS as per the above table. A given service must use a bandwidth profile that is consistent with the services PON CoS. For example, if the service uses P-bit 0, 1, or 2, it is by definition a best effort service and must use a bandwidth profile where upstream CIR=0, and upstream PIR > 0.

How bandwidth is allocated

The Optical Line Terminal (OLT) provides GPON network terminations. The E7 GPON cards serve as the OLT, and are responsible for allocating bandwidth to ONTs either as fixed allocated bandwidth when DBA is disabled, or dynamically (for non-fixed traffic classes) via the DBA algorithm when DBA is enabled.

A transmission container (T-CONT) is an ONT object that represents a group of logical connections which appear as a single entity for the purpose of upstream bandwidth assignment on the PON. Each ONT can have up to four service T-CONTs assigned—one per PON class of service. On supported ONTs, four additional, user-defined TCONTs may be assigned. The ONT reports the status of the queues associated with the T-CONTs to the OLT. Based on demand and the constraints of the DBA algorithm, the OLT facilitates sharing of the upstream bandwidth by allocating bandwidth grants to the T-CONTs. Excess bandwidth, either unconfigured or not currently required by a subscriber, is allocated to subscriber services based on the DBA priority order shown in the table above.

Ethernet Bandwidth Profile Examples:

Best Effort (BE) profile

- Upstream CIR 0m
- Upstream PIR 50m
- Downstream PIR 100m

The Ethernet service MUST use P-bits 0, 1, or 2; or Cos 1 which is equivalent to 0-2. Best Effort means, if there is BW available, then up to 50M can be sent for this example. Best Effort is typically used for data services.

You might want to create three Ethernet Bandwidth profiles to support three tiers if Internet data service.

Example: Assured Forwarding (AF) profile

AF1:

- Upstream CIR 25m
- Upstream PIR 100m
- Downstream PIR 100m

AF2:

- Upstream CIR 50m
- Upstream PIR 200m
- Downstream PIR 1g

For AF1, the Ethernet service **MUST** use P-bit 3 or CoS 2, which is equivalent to P-bit 3. Since the AF definition is the same, it is the P-bit (3 or 4) that determines Assured Forwarding 1 or Assured Forwarding 2. A “committed” rate of 25M or 50M is guaranteed (in examples) and if there is extra bandwidth available, then up to 100M or 200M will be allowed. This profile can also be used for data or, with a P-bit of 4, for video but the cost to the end-user would increase, since there is a guarantee for part of the service.

You might want to create three Ethernet Bandwidth profiles to support three tiers if video service.

Example: Expedited Forwarding (EF) profile

- Upstream CIR 50m
- Downstream PIR 1g

A “committed” rate of 50M is guaranteed. No bandwidth above 50M will be available, but it will always have 50M assigned, even if no traffic is flowing. This profile is typically used for voice service. This profile could also be used for video if the end-user agrees to the cost, as this would be the most expensive.

See Creating a PON Class of Service.

DBA configuration guidelines

- The sum of guaranteed bandwidth (Committed Information Rate (CIR)), cannot exceed the total upstream bandwidth of the GPON interface (1.2 Gbps).
- For E7-20 GPON-16x only, the sum of guaranteed bandwidth (Committed Information Rate (CIR)), cannot exceed the total downstream bandwidth of the GPON interface (2.4 Gbps).
- DBA is enabled by default.
- DBA is set at the GPON port level.
- DBA supports up to 64 ONTs per PON.
- DBA-managed traffic includes the Assured Forwarding and Best Effort traffic class types only.
- All provisioned Expedited Forwarding bandwidth is fixed and statically pre-scheduled, and is not available for reuse by the DBA algorithm.
- Assured Forwarding services scheduled by DBA have a minimum bandwidth allocation of 64 Kbps per service, and a maximum of 1G per service.
- Best Effort services scheduled by DBA have a maximum bandwidth allocation of 1G per service.

- The DBA algorithm operates on the configured bandwidth profiles for each subscriber and service, and the perceived demand from the ONTs.
- Enable/disable DBA via the Command-line Interface (CLI).
- DBA does not support configuration via the Calix Management System (CMS) or the web interface.

Home Gateway services guidelines

When provisioning High Speed Internet (HSI) services on the routed WAN interface of a 700GE ONT or an 836GE RSG, downstream traffic shaping is supported with Bandwidth Profiles from 0-250 Mbps, and downstream traffic policing (by the OLT) is supported with Bandwidth Profiles from 250-1000 Mbps.

Using a Bandwidth Profile with an open bandwidth setting (CIR=0 and PIR=0) does not apply any bandwidth control, allowing the service to be shaped or policed by a router/BNG upstream from the OLT, up to approximately 1000 Mbps. With 700GEs/836GEs, this is the only way to achieve shaping for data services greater than 250 Mbps.

When provisioning HSI services on the routed WAN interface of 800 GigaCenters and GigaHubs, traffic shaping is supported from 0-1000 Mbps. An open bandwidth setting is not required or recommended for 800 GigaCenters/GigaHubs.

However, shaping is only supported on one interface of a device. As a result, if services are provided to multiple interfaces, such as HSI (wired and wireless) to the routed WAN interface and IPTV or private E-LAN or E-Line services to bridged Ethernet interfaces, the bridged interfaces will be policed, not shaped. However, for these services on bridged interfaces, shaping is not necessary or desired.

Ethernet bandwidth profile configuration guidelines

- A given service must use a bandwidth profile that is consistent with the service PON CoS. For example, if the service uses P-bit 0, 1, or 2, it is by definition a best effort service and must use a bandwidth profile where upstream CIR = 0, and upstream PIR > 0.
- The bandwidth limits for all ONT services (in the same class) of a single ONT Ethernet port are aggregated into a single CoS metered rate, therefore only one service per CoS per ONT port is allowed.
- The E7 ports have a single-rate traffic shaping function, differing from the MEF scheme and terms. For example, in MEF terms, Committed Burst Size (CBS) is a component of a two-rate traffic shaping system. Where you have:
 - CIR (Committed Information Rate)
 - CBS (Committed Burst Size) for the CIR buffer
 - EIR (Excess Information Rate)
 - EBS (Excess Burst Size) for the EIR buffer

However, with the E7 single-rate configuration settings, you have:

- Maximum Rate
- Burst Size

Consider the provisioned Maximum Rate to be the equivalent of the provisioned Committed Information Rate (CIR) rate in a two-rate system when the Excess Information Rate (EIR) value = zero. That is, $CIR + (EIR = 0) = PIR = Max$.

- Bandwidth across the E7 GPON subsystem is managed by the E7 OLT. Bandwidth is not affected by the number of ONTs on the PON, by the distance of the ONTs from the OLT, or the split ratio / design.
- Downstream, PON traffic management is much like the rest of the E7 core Ethernet architecture – services can be guaranteed bandwidth, but can also take advantage of excess bandwidth.
 - 8 P-bit priority classes are transported through the E7 core Ethernet transport and switching network.
 - Up to 2.5 Gbps P-bit priority-scheduled traffic is forwarded to the GPON MAC.
 - Up to 2.4 Gbps non-blocking traffic is sent down the PON to individual ONTs based on MAC learning.
 - The E7 OLT enforces the Peak Information Rate for each service.
 - Downstream traffic is dynamic and all services can burst to the maximum PON rate in the absence of higher priority traffic.
- For Downstream CIR, the committed bandwidth will not be released when a service is disabled or an ONT departs or is quarantined, although the bandwidth can be shared with other services. To release the committed bandwidth, remove the service or delete the ONT.
- Upstream on the PON, traffic is scheduled based on a request/grant cycle initiated by the ONT. Each provisioned service on the PON may have one of three components – Expedited, Guaranteed and Non-guaranteed. Expedited service has the lowest jitter and latency and is limited to roughly 50% of the PON upstream bandwidth capacity.
 - Expedited Forwarding (EF) class limit = 600 Mbps (composed solely of CIR traffic)
 - The Committed Information Rate (CIR) traffic limit = 1200 Mbps (includes the EF component)
 - If the total PON Peak Information Rate (PIR) traffic is over 1200 Mbps, all non-guaranteed traffic shares proportionately the bandwidth not used by the CIR traffic.
 - Up to 1.2 Gbps capacity
 - The OLT schedules all PON traffic based on service provisioning (not actual use). Once scheduled, traffic allocation is reserved and is not dynamically reallocated to other services if unused.
 - The ONT enforces upstream CIR and PIR rates for each service according the provisioned service bandwidth profile

- Best Effort and the Assured Forwarding non-guaranteed (PIR-CIR) traffic have equal priority and are oversubscribed as a class. Once the PON has been provisioned to its upstream capacity, additional non-guaranteed traffic Admitted to the PON will reduce all non-guaranteed traffic proportionately.
- T1 and VOIP services are required to use the Expedited Forwarding class.
- The EF traffic could be zero, if the sum of AF CIR is equal to 1200. There would be no capacity left for non-guaranteed traffic.
- All CIR traffic is scheduled at 100% of the requested bandwidth profile, up to an aggregate limit of 600 Mbps. Only Expedited Forwarding and Assured Forwarding class types in the GPON CoS global configuration have a CIR component.
- After CIR traffic is scheduled, the remaining PON bandwidth is shared proportionally based on the difference of the CIR and PIR.
- Downstream CIR only applies to E7-20 GPON-16x line card. This parameter has no effect on all other GPON cards (GPON-4, GPON-8, GPON-8x).
- When the PON is not over-subscribed in the upstream direction, all services get 100% of their PIR. When the PON is over-subscribed, no service gets 100% of the requested PIR.
- If the aggregate PIR bandwidth exceeds the PON bandwidth capacity, the services are scheduled a % less than their PIR based on the over-subscribed amount.

Example:

PON bandwidth = 1200 Mbps

3 ONTs with 1 service on each ONT

Rates are configured on the upstream portion of the services, relative to the oversubscribed amount:

ONT	CIR	PIR	Over (PIR — CIR)	Percentage
#1	50	800	750	750 / 2100 = 35.7%
#2	100	800	700	700 / 2100 = 33.3%
#3	150	800	650	650 / 2100 = 30.9%

2100 — Total

1200 (available) — 300 (Sum of CIR as shown in table) = 900 (capacity available for non-guaranteed traffic)

With the following values known, the calculations can be used:

- Capacity available for non-guaranteed traffic
- Correct percentages to apply to each of the services relative to the capacity left over for non-guaranteed traffic

ONT	Calculation	Expected Upstream Rate
#1	50 + (900 x 0.357)	371.3 Mbps
#2	100 + (900 x 0.333)	399.7 Mbps
#3	150 + (900 x 0.309)	428.1 Mbps

- Traffic moving through the E7 and downstream onto the PON is prioritized and scheduled according to the 8 Cos configuration in the Ethernet switch and then handed to the GPON OLT at an aggregate rate of 2.5 Gbps. In the OLT GPON MAC, the peak rate for the service is enforced (shaped) prior to sending the traffic down the PON to the ONTs. Upstream and downstream, the P-bits are mapped into one of the 4 GPON CoS global configuration.
- For any GPON service through the E7, the downstream P-bit must be set on all frames for a given ONT service (either upstream of the E7 or at ingress into the E7 using a policy map) to match the P-bit that is provisioned by the user in the service when it is added to the ONT.

The provisionable limits for each bandwidth rate are as follows:

- The minimum metered rate is 64 Kbps.
- The maximum metered rate is 1 Gbps.
- Rates can be selected in 64 Kbps increments up to 2 Mbps.
- Rates can be selected in 1 Mbps above 2 Mbps.
- Rates can be selected in 100 Mbps above 1 Gbps.
- Setting the metered rate to 0 Kbps disables the meter.
- **Note:** 700GX ONTs can rate limit traffic to a maximum rate of 400Mbps. Setting the rate limit for values higher than 400Mbps will disable the rate limiter.

PON Upstream Class of Service Profiles

PON upstream Class of Service (CoS) profiles allow you to expand on the 4 system-default instances of CoS to 8 priorities of scheduled service per ONT. This model supports multiple instances of the same traffic class per ONT to be treated independently for rate-limiting and shaping.

For example, if there are two Best Effort (BE) services provisioned on Port 1, you could assign one of the services to PON CoS **user-1** so that each service would have a dedicated T-CONT type in the upstream direction, rather than sharing a queue that is set to the aggregate bandwidth of both services, allowing any one service to use all of the bandwidth if there is no contention for it.

Service Name	Traffic Class	PON CoS	Downstream Bandwidth (PIR)	Upstream Bandwidth (CIR/PIR)	ONT Upstream Queue
Data-1	BE	cos-1	50 Mb/s	0/5 Mb/s	queue 2
Data-2	BE	user-1	10 Mb/s	0/1 Mb/s	queue 1

- The PON upstream CoS profile configuration and application process:
 - a. Create a PON US CoS profile with an identifying name for mapping.
 - b. Map the PON US CoS profile to an instance (user1-user4) at the ONT.
 - c. Select the PON CoS instance (user1-user4) that is mapped to a PON US CoS profile, when adding an Ethernet service to the ONT port.

Configuration guidelines

- There are still three classes of service (BE, AF, EF) for the user-defined PON upstream CoS profiles. However, from the Dynamic Bandwidth Allocation (DBA) scheduling perspective, a given GPON CoS can be prioritized so that there can be up to 4 AF classes of service and up to 4 BE classes of service, where $EF > AF_x > BE_x$ is the DBA scheduling priority. Within a given service class, the services at the highest priority will be allocated Excess bandwidth first and then the next highest priority.
- If explicit CIR/PIR values are included for a PON upstream CoS profile, these values take precedence over the Ethernet bandwidth profile associated with the service. Additionally, the upstream CIR is only implemented at the upstream PON.

Examples:

- If the user-defined PON CoS (user 1-4) CIR is 0, there will be no guaranteed bandwidth, even if the service has set a CIR via the Ethernet bandwidth profile.
- If the user-defined PON CoS (user 1-4) PIR is less than the Ethernet bandwidth profile PIR, the service will only get the PON CoS PIR as the peak rate.
- If the CIR and PIR are set on the user-defined PON CoS (user 1-4), the values should be considered a cap for all services associated with that PON upstream CoS profile. If the associated services are not managed to fit within these caps, you may see undesired or unexpected behavior.

For example, the following operation exists with the particular settings on the PON upstream CoS profile:

- PIR/CIR = aggregate
There is a total of 100M for three services of 60M, 20M, and 20M.
- PIR/CIR = 80M
There is a cap of the three services of 60M, 20M, and 20M, where the DBA prioritization would dictate which service receives the appropriate throughput.
- The following ONT-specific differences exist around a service, with respect to upstream shaping:
 - GX-series ONTs:
 - Four classes of service (cos 1-4) where the total = 4 T-CONT types
 - PON upstream CoS profile (user 1-4) application is NOT supported
 - Can rate limit the upstream bandwidth per individual service instance so that multiple instances of the same CoS share the upstream bandwidth according to the bandwidth profile of each individual service. One instance of a CoS cannot use all the bandwidth allocated to the PON CoS for that ONT by the Dynamic Bandwidth Allocation (DBA) function. Maximum bandwidth allocated per ONT CoS is the sum of the services bandwidth profiles defined with a matching CoS on a single ONT.

- GE-series ONTs:
 - Four fixed classes of service (cos 1-4) plus four user-defined classes of service (user 1-4) are supported for a total = 8 T-CONT types
 - Although using a pon-cos, other than fixed, enables upstream GEM port shapers that rate limit per service, the implementation of a dedicated T-CONT type is recommended for the most accurate upstream limiting.
- T-series ONTs (with E7 R2.3+):
 - Four system-defined (fixed) classes of service (cos 1-4) plus three user-defined classes of service (user 1-3) where the total = 7 T-CONT types
 - PON upstream CoS profile (user 1-3) application is supported where user-4 is NOT supported

Parameters

You can provision the following parameters for the PON upstream class of service profile:

Parameter	Description	Valid Options
Name	Name of PON upstream class of service profile.	text string
Type	Traffic class forwarding type for CoS. Typically, the following traffic classes, GPON CoS, and P-bit values are associated: • Expedited Forwarding (EF), CoS 4, P-bits 5,6,7 • Assured Forwarding 1 (AF1), CoS 3, P-bits 4 • Assured Forwarding 2 (AF2), CoS 2, P-bits 3 • Best Effort (BE), CoS 1, P-bits 0, 1, 2	expedited, assured, best-effort, blank=not used
Priority	Dynamic Bandwidth Allocation (DBA) priority where a given GPON CoS can be prioritized so that there can be up to 4 AF classes of service and up to 4 BE classes of service. Excess bandwidth, either unconfigured or not currently required by a subscriber, is allocated to subscriber services based on the DBA priority. • Where EF>AFx>BE is the DBA scheduling priority. • Within a given service class, the services at the highest priority will be allocated Excess bandwidth first and then the next highest priority.	1 (default), 2, 3, 4
CIR Bandwidth	Committed Information Rate (CIR) Bandwidth (Mb).	0-1200 Mb, aggregated (default)
PIR Bandwidth	Peak Information Rate (PIR) Bandwidth (Mb).	0-1200 Mb, aggregated (default)

PON Class of Service Map

Each E7 GPON provides a minimum of four classes of service for forwarding traffic. A service may be provisioned so that traffic is placed in a class based on the P-bit value in the VLAN tag. Since there are 8 possible P-bit values, there is not a one-to-one mapping between P-bit values and the classes of service. The configuration for each class defines the following:

- P-bit values that will be included in the class
- Forwarding type

Note: The upstream traffic on a service VLAN must have the same P-bit value as is set for downstream traffic. The P-bit value must also be consistent with the class of service type, as defined in the PON Class of Service (CoS) table.

The configurations apply to all PON ports in the system. Each ONT service can have independent CoS queue management, as appropriate for the services on a particular GPON port.

Note: Within an E7 network, ensure that the IPTV traffic has a high P-bit (priority) value for points of congestion. Calix recommends using a minimum P-bit value of 4 for video services. Typically, an upstream IP router sets the VLAN-ID and P-bit values at the transition to the Ethernet access network (for example, at the point before where an E7 connects to an edge router).

If necessary, you can apply a Policy Map to the E7 Ethernet (ingress) port to mark (re-mark) ingress traffic with a new P-bit value.

The default mapping between GPON CoS to P-bits is as shown:

GPON CoS	Scheduling Class Type	Default P-bits	Service
CoS 4	Expedited Forwarding	5, 6, 7	Network control, voice, T1
CoS 3	Assured Forwarding 1	4	Video
CoS 2	Assured Forwarding 2	3	Application signaling, TLAN
CoS 1	Best Effort Forwarding	0, 1, 2	High-speed internet

Note: Any change made to the GPON CoS global configuration that invalidates existing service provisioning will cause all GPON ports to be reset.

Configuration guidelines

- The P-bit value defined in the service-tag action must be consistent with the class of service type as defined in the PON Class of Service (CoS) table.
- PON CoS is a system-level configuration and applies to all PON ports in the E7 system.
- As a higher number P-bit denotes higher priority, so does a higher PON CoS value.
- All P-bits 0...7 must be represented in the mapping table.

- A P-bit can only belong to a single PON class of service.
- P-bits in lower classes of service must be lower than P-bits in higher classes of service.
- A lower PON CoS must have lower or equal Scheduling Class Type when compared to a higher CoS. The ordering of Scheduling Class Type values is Best Effort > Assured > Expedited.
- The P-bits can be mapped to fewer than all four CoS levels, however all CoS levels above the unused CoS must also be unused.
- Because the bandwidth limits for all ONT services (in the same class) of a single ONT Ethernet port are aggregated into a single Class of Service (CoS) metered rate, only one service per CoS per ONT port is allowed.

Parameters

You can provision the following parameters for the global GPON port class of service (CoS):

Parameter	Description	Valid Options
cos1-type	Forwarding type for CoS 1. "Not used" indicates no forwarding (CoS is unused).	expedited, assured, best-effort, not used
cos1-low-P-bit	Low bit value for CoS 1. "None" is the default behavior and indicates the P-bit value is not set.	0-7 none (default)
cos1-high-P-bit	High P-bit value for CoS 1. "None" is the default behavior and indicates the P-bit value is not set.	0-7 none (default)
cos2-type	Forwarding type for CoS 2. "Not used" indicates no forwarding (CoS is unused).	expedited, assured, best-effort, not used
cos2-low-P-bit	Low bit value for CoS 2. "None" is the default behavior and indicates the P-bit value is not set.	0-7 none (default)
cos2-high-P-bit	High P-bit value for CoS 2. "None" is the default behavior and indicates the P-bit value is not set.	0-7 none (default)
cos3-type	Forwarding type for CoS 3. "Not used" indicates no forwarding (CoS is unused).	expedited, assured, best-effort, not used
cos3-low-P-bit	Low bit value for CoS 3. "None" is the default behavior and indicates the P-bit value is not set.	0-7 none (default)
cos3-high-P-bit	High P-bit value for CoS 3. "None" is the default behavior and indicates the P-bit value is not set.	0-7 none (default)
cos4-type	Forwarding type for CoS 4. "Not used" indicates no forwarding (CoS is unused).	expedited, assured, best-effort, not used
cos4-low-P-bit	Low bit value for CoS 4. "None" is the default behavior and indicates the P-bit value is not set.	0-7 none (default)
cos4-high-P-bit	High P-bit value for CoS 4. "None" is the default behavior and indicates the P-bit value is not set.	0-7 none (default)

DSCP Map

To associate a DSCP Map with a service, do the following:

1. In the service tag action that is going to be applied to the data or video service, select **Map a Layer-3 priority** from the P-Bit Source list.
2. On the ONT port where you are going to be provisioning the data or video service, select the name of the DSCP profile to use from the DSCP/IP Precedence Profile list.

Note: The upstream traffic on a service VLAN must have the same priority value as is set for downstream traffic.

Note: If IP traffic on an Access Interface has a DSCP setting other than the values included in the assigned DSCP map, the system interprets this condition as a mismatch and uses a "Default" P-bit value from the assigned DSCP Map profile. For example, if you were using the system-default DSCP Map "Access" and the incoming frames had a DSCP value of 001 001, that value is not included in the assigned DSCP map, so the E7 would use the default P-bit value of 0. However, you can create a DSCP Map with a Default P-bit value of your choice, and then assign the DSCP Map to the ONT port.

The system-default DSCP profile "Access," has values as follows:

DSCP Value	Meaning	Decimal	Drop Probability	802.1p Priority-Bit
000 000	BE-CS0	0	n/a	0
001 000	CS1	8		1
001 010	AF11	10	Low	1
001 100	AF12	12	Medium	1
001 110	AF13	14	High	1
010 000	CS2	16		2
010 010	AF21	18	Low	2
010 100	AF22	20	Medium	2
010 110	AF23	22	High	2
011 000	CS3	24		3
011 010	AF31	26	Low	3
011 100	AF32	28	Medium	3
011 110	AF33	30	High	3
100 000	CS4	32		4
100 010	AF41	34	Low	4
100 100	AF42	36	Medium	4
100 110	AF43	38	High	4
101 000	CS5	40		5
101 110	EF	46		5
110 000	CS6	48		0
111 000	CS7	56		0
	Default			0

Best Effort (BE)
 Assured Forwarding (AF) gives assurance of delivery under prescribed conditions
 Expedited Forwarding (EF) dedicated to low-loss, low-latency traffic
 Default = P-bit assigned for a DSCP value that is not included in the DSCP Map

Note: Calix recommends NOT assigning P-bit value 7 to any DSCP value.

Match Lists for GPON Services

The ONT performs the following two actions:

- **Classification** – the service match lists are for matching and classifying traffic into services.
- **Marking** – the service tag actions are for marking the traffic with the appropriate VLAN and P-bit required by the service.

A match list defines a set of criteria (matching rules) for classifying ONT traffic into services. A match list can contain both "tagged" and "untagged" match rules, up to 12 tagged rules and up to 16 untagged rules for each ONT Ethernet port. The match list defines how the ONT port classifies subscriber traffic into a service VLAN by associating it in a service tag action that is then associated with an ONT Ethernet service.

- **Untagged match rules** can match on the following:
 - A portion of the source MAC address as indicated by the Source MAC and Source MAC Mask attributes of a video Set-Top box (not supported with T-Series ONTs)
 - The Ethernet type to distinguish video and High-Speed Internet (HSI) untagged traffic (not supported with P-Series and GigaCenter ONTs)
 - All untagged traffic by using the system default match list "all-untagged"
 - PVC VPI and VCI port numbers (only applies to VDSL2 services)
- **Tagged match rules** can match on any combination of the following:
 - Outer tag
 - VLAN-ID
 - P-bit values
 - Tag Protocol Identifier (TPID) can also be specified at the E-Series egress Ethernet port interface, but defaults to 0x8100

Example of multiple match rule use for classifying traffic:

A service provider wants to use a single ONT Ethernet port to connect to a home network with a combination of IPTV and data traffic. The IPTV traffic is identified by using the OUI value of the video set top box.

Note: The T-Series ONTs do not support configuring a match rule with an OUI filter.

This combination of rules can accomplish the example task:

1. An untagged traffic match rule would be applied with a specific MAC OUI match criteria.
2. A second untagged match rule would be applied with a "match any" criteria.

Using a match list and service tag action, Calix ONT's support matching the traffic from all set-top boxes (STB) of a particular manufacturer by indicating the STB OUI in the first half of the specified Source MAC attribute followed by zeros, and then providing a mask for the STB OUI by indicating ff:ff:ff:00:00:00 in the Source MAC Mask attribute.

Example entries that match all STB source MAC addresses of a particular manufacturer:

- ADB set-top boxes: 00:03:91:00:00:00
- Amino set-top boxes: 00:02:02:00:00:00

Note: If STBs from multiple vendors are being used, create additional untagged match rules, entering the OUI for each STB vendor.

Alternatively, the ONT's support matching the traffic from a specific host by indicating the MAC address in the Source MAC attribute, and then providing a mask for the address by indicating ff:ff:ff:ff:ff:ff in the Source Mac Mask attribute.

Configuration guidelines

- If a residential gateway (RG) at the subscriber's premise classifies traffic into VLANs, then for each service type, create a match list with a tagged rule that matches the service VLAN.
- Because only video services are supported on the P-series FB interface, it is not necessary to specify the OUI of the STB, and the all-untagged match list applies.
- Up to 32 OUIs are supported per match list.
- Ethertype is the only untagged traffic classifier supported by the T-Series ONT's.
- The ONT's do not support tagged and untagged match rules on the same ONT for video services. Only one classification rule is allowed for video services per ONT.

Match list and rule parameters

You can provision the following parameters for match list and rules:

Parameter	Description	Valid Options
Name*	Name of service match list.	text string
S-VLAN (Outer Tag)	VLAN ID of the outer tag (tagged match rules only). VLANs can be specified by name or by numeric VLAN ID, which ranges from 1-4094. In addition, "untagged" indicates that only untagged traffic should be matched and "ignore" indicates that the VLAN ID should not be examined.	text string, 1-4094, untagged, ignore (default)
Outer Pbit	P-bit value of the outer tag (tagged match rules only). P-bit values are in the range 0-7. Alternately "pbit-none" means the P-bit value is not considered. If this parameter is not specified, "p-bit-none" is the default behavior.	0-7, ignore, p-bit-none(default)
Source MAC	Source MAC address (untagged match rules only).	six hexadecimal digits in the range 0-FF, optionally separated by colons Default = ignore
Source MAC Mask	Source MAC mask (untagged match rules only). - source MAC mask (x:x:x:x:x) - none (no MAC mask) - oui (mask OUI fields (FF:FF:FF:FF:FF:FF)) Note: The T-Series ONTs do not support configuring a match rule with an OUI filter.	six hexadecimal digits in the range 0-FF, optionally separated by colons, none (default), oui
Ethertype	Ethernet type to match (untagged match rules only). - any = default - pppoe = 0x8863 or 0x8864 (for HSI) - arp = 0x0806 (for video) - ipv4 = 0x0800 (for video) - ipv6 = 0x86DD (for video) Note: Not supported for P-Series and GigaCenter ONTs.	any (default), pppoe, arp, ipv4, ipv6

*Required field

Service-Tag Actions for GPON Services

The ONT service-tag actions are for marking the traffic with the appropriate VLAN and P-bit required by the type of service.

Service-tag actions are applied to subscriber service provisioning as follows:

- Video and data services provisioning on an ONT Ethernet port includes a reference to a service-tag action.
- Voice services provisioning on an ONT Voice port implicitly references a Voice Service IP Host that includes a reference to a service-tag action.

The E7 and the ONT together can apply the following tag actions assigned to an ONT port service.

Add a single tag	Adds a single outer tag to the matched traffic. • Can be applied to tagged and untagged traffic. • Supported for DSCP-to-P-bit translation. • Supported for Full Bridge (FB) interface services. • NOT supported for Residential Gateway (RG) interface services.
Add 2 tags	Adds an inner and outer tag to the matched traffic*. • Can be applied to untagged traffic, only. • Supported for DSCP-to-P-bit translation. • NOT supported for FB or RG interface services.
Add and change tag	Adds an outer tag to the matched traffic and changes the inner tag. • Can be applied to tagged traffic, only. • Supported for RG interface services. • NOT supported for FB interface services. • NOT supported for DSCP-to-P-bit translation.
Change tag	Changes the outer tag on the matched traffic. • Can be applied to tagged traffic, only. • Supported for FB or RG interface services. • (Prior to R2.6.30) NOT supported for DSCP-to-P-bit translation. • R2.6.30 introduces support for DSCP-to-P-bit translation.

*To properly process ingress double tags, the GE network interface (uplink) must be configured as a Trunk role.

Determining the service-tag action values

When a service is provisioned at an ONT Ethernet port, a match list and service-tag action specify the classification and marking of packets from the subscriber port into the service VLAN.

You have the following options for determining the added- or changed-tag values:

- **Specify the VLAN value explicitly.**

This option is used for a VLAN-per-service provisioning model (N:1). A service carried on an N:1 VLAN applies to multiple subscriber ports, where a single match list and tag action can indicate the service.

- **Reference the VLAN value from the service provisioning on the ONT port.**

This option is used for a VLAN-per-port provisioning model (1:1). A service carried on a 1:1 VLAN is the same for each subscriber except the customer tag is unique for each single subscriber port, so you can define the match list and tag action pair such that multiple subscriber ports can reference it. This is accomplished by indicating "Specified in Service" for the Output Tag value in the tag action, and then the subscriber-specific value for the output tag is defined when the service is provisioned on the port.

Note: The service provisioning process follows the creation of a service-tag action.

You also have the following options for determining the service-tag action P-bit value, assigning it to a Class of Service (CoS):

- **Specify the P-bit value explicitly.**

Enter the P-bit value to use for P-bit marking the outer tag.

- **Specify the CoS.**

Enter the CoS ID to use for P-bit marking the outer tag to the lowest P-bit value associated with this CoS. This value identifies the PON class of service.

GPON CoS	Scheduling Class Type	Default P-bits	Service
CoS 4	Expedited Forwarding	5, 6, 7	Network control, voice, T1
CoS 3	Assured Forwarding	4	Video
CoS 2	Assured Forwarding	3	Application signaling, TLAN
CoS 1	Best Effort Forwarding	0, 1, 2	High-speed internet

- **Reference a Layer-3 Priority map.**

Use a DSCP table to map untagged frames coming from the subscriber to a P-bit value. From the ONT Ethernet port, you can select the system-default profile ("access") or specify another custom profile. See Layer 3 Priority to P-Bit Mapping. Also, see Creating a Policy Map for L3 Priority Mapping for instructions on using a class map to match the DSCP bits of untagged frames of ingress traffic on Trunk and Edge Ethernet (GE/10GE) ports, and then using a policy map to assign a corresponding P-bit (priority) value into the classified traffic.

Configuration rules:

- Only one service per port can use the DSCP-Bit mapping function. However, a second service of untagged traffic on the same port that matches a specific OUI value is supported, where that OUI-matching service takes precedence for any matching frames.
- Layer-3 priority maps work with add-tag and add-2-tags tag actions.
- If you select "Map a layer-3 Priority," in the service tag action, then when creating the service, you must select a PON CoS parameter of User1-4 or CoS1-4. Do not leave the system-default parameter of "derived" as there is no single unique P-bit specified from which to derive the intended CoS.

- The T-Series ONTs do not mark the P-bit of broadcast and multicast frames based on DSCP. T-Series ONTs handle all broadcast / multicast packets from the subscriber as best-effort and mark them with the default P-bit of the service.

Provisioning sequence:

- Create a DSCP-to-P-Bit map, if you do not want to use the system-default.
- Assign the map to the ONT port, or leave the system-default map "access" assignment.
- Create a service-tag action with an Outer P-Bit Source of **Map a layer-3 Priority**.
- Create a service using the defined service-tag action.

The following table shows the 802.1p to DiffServ DSCP pre-defined, system-default mapping table ("access").

Default P-bits	DiffServ Code Point**	Service
0	CS6, CS7 (48-56)	Network control
0	EF (46)	Voice, NSP, CES
5	CS5 (40)	Network Management
4	CS4 & AF41-43 (32-38)	Video
3	CS3 & AF31-33 (24-30)	Data3
2	CS2 & AF21-23 (16-22)	Data2
1	CS1 & AF11-13 (8-14)	Data1
0	0	Best Effort

**The E7 supports alphanumeric DSCP values that denote forwarding classes such as AF21, AF31 in addition to numeric value assignment.

- **Preserve the existing P-bit value upstream.**

Indicate that the E7 preserves the P-bit value set in service provider supplied Residential Gateways (RGs).

- Copies the subscriber-side Q-Tag P-Bits to a single C-tag or S-Tag, or to the inner C-Tag of a double-tagged Service.
- For the outer tag P-bit, you can specify the value of "copy".
- For the inner tag, you can specify the value of "preserve".

- **Promote the existing P-bit value upstream.**

Indicate that the E7 promotes the P-bit value set in service provider supplied Residential Gateways (RGs).

- The P-bit of the inner tag is promoted to the P-bit of the outer tag, allowing the network QoS to handle the priority of the frame consistently.
- For the inner tag P-bit, you can specify the value of "same-as-outer".

Note: The upstream traffic on a service VLAN must have the same P-bit value as is set for downstream traffic. The P-bit value must also be consistent with the class of service type as defined in the PON CoS table.

Note: For VDSL2 and GPON subscribers (if PON CoS = Fixed), downstream traffic with P-Bit values not explicitly defined by the service(s) on the port will be passed through with unmetered bandwidth.

VLAN Traffic flow in GPONs

Since E7 GPON Ethernet ports have no interface association, unlike the E7 Ethernet ports, rather than adding a port interface to a VLAN membership to enable traffic flow, a service-tag action must be created that specifies the VLAN. This service-tag action is then referenced when the service is provisioned on the ONT port. The VLAN must already be created on the E7 for a GPON port to pass traffic carried on a VLAN.

- To view the VLANs associated with specific GPON ports, click **E7 > GPON-x > GPON:** on the Navigation Tree, and then click the **VLANS** tab.
- To view the VLANs associated with an ONT, click **ONTS** on the Navigation Tree, and then click the **VLANS** tab.
- To view the ONT services associated with a specific VLAN, click **VLANS** on the Navigation Tree, click the particular VLAN from the list that appears in the Work Area, and then click the **Service Associations** tab.

The ONT Ethernet ports support the same VLAN services as the E7 Ethernet ports, including native IEEE 802.1Q VLAN tagging and IEEE 802.1ad VLAN stacking (Q-in-Q). The following service models are supported.

- VLAN per subscriber Ethernet port (1:1)
- VLAN per Service (N:1)
- Ethernet virtual circuits (VLAN trunks)
- Transparent LAN Services

Within the E7 GPON OLT and ONT subsystem, MAC learning and switching occurs within a single, outer VLAN ID, referred to as the VLAN C-tag. Packets on the PON always have a C-tag, which consists of a provisioned VLAN ID and p-bit. The ONT is responsible for classifying subscriber traffic and adding/modifying the C-tag.

Configuration guidelines:

- The P-bit value specified in the service-tag action must be consistent with the class of service type as defined in the PON Class of Service (CoS) table.
- The ONT requires that SIP traffic be set to Expedited Forwarding. Therefore, the P-bits set in the service tag action must map to the Expedited Forwarding CoS.
- Each service-tag action references a service-match list (classify) and specifies a tag action (mark and tag).
- Each tag action is associated with a specific service.
- The ONTs do not support tagged and untagged match rules on the same ONT for video services. Only one classification rule is allowed for video services per ONT.
- The logical components within the RSG are functionally similar to a standard GPON ONT with a subtended 3rd party residential gateway (RG): The ONT receives tagged traffic from the downstream RG and sends it upstream single-tagged where the OLT may or may not add a second tag.
- Services provisioned on a Full Bridge interface only support:
 - add-tag
 - change-tag
- Services provisioned on a Residential Gateway interface only support:
 - change-tag
 - add-and-change-tag
- The following tag actions are supported for VLANs with IGMP enabled:
 - Add-tag
 - Change-tag
- Services provisioned with a tag action Outer P-bit source of **Map a layer-3 Priority** (DSCP translation to P-bit values) support the following:
 - add-tag
 - add-2-tags
 - Only one service per port can use the DSCP-Bit mapping function. However, a second service of untagged traffic on the same port that matches a specific OUI value is supported, where that OUI-matching service takes precedence for any matching frames.
 - Each port is independent where it can use a different DSCP map.
 - GPON services only support DSCP to P-bit mapping, (while xDSL services support DSCP or IP Precedence for a service, not both).
 - A PON CoS parameter of User1-4 or CoS1-4 must be used. Do not leave the system-default parameter of “derived” as there is no single unique P-bit specified from which to derive the intended CoS.

- All of the traffic classified by the DSCP map and P-bit-marked according to the DSCP map table is put into a single queue for upstream scheduling by the OLT.
- In the event of congestion on the upstream bandwidth allocation, traffic in the queue with the highest P-bit is transmitted first, thus preserving the DSCP set by the downstream router.
- Multiple services and service-tag actions can be applied to a single ONT Ethernet port.
- The all-untagged match rule cannot be applied to multiple service-tag actions on the same port. For example, one of the untagged match rules must point to a specific OUI. The system applies the tag action that has the most specific match rule.
- The same service-tag action can be used on multiple ONT Ethernet ports.
- Each E7 shelf, or E7 Modular Chassis system has a capacity of 256 service tag actions. For configurations that require large numbers of VLANs, Calix recommends selecting the **Specified in Service** value for the "Outer Tag" parameter, and then specifying the VLAN ID when provisioning the service.
- The upstream and downstream values specified in the Ethernet Bandwidth Profile and P-bit values defined in the service-tag action must be consistent with the class of service type as defined in the global CoS table.
- For traffic to flow on the VLAN specified in a service-tag action, the VLAN and uplink must already be created on the E7.
- For a GPON port to pass traffic carried on a VLAN, a service-tag action must be created specifying the VLAN and the VLAN must already be created on the E7.
- If no matching criteria are assigned to a tag action, then the tag action is performed on all packets entering the GPON port. That is, when a match list has a single untagged rule with the outer tag set to "ignore" and the Outer P-bit set to "Pbit-none."
- At least one tag must be specified in the match criteria because the system cannot change the VLAN ID for untagged traffic.
- Traffic must have P-bits assigned in order to be scheduled (prioritized) on the GPON.
- The "change-tag" action changes the VLAN ID in the outermost tag.
- For double-tagged packets, the outer VLAN ID must match a provisioned VLAN. The inner tag does not have to match a provisioned VLAN ID.
- Edge interfaces can only use single-tag actions when DHCP Snooping is enabled on the associated service VLAN.
- To properly process ingress double tags, the GE network interface (uplink) must be configured as a Trunk role.
- The E7 reserves four VLAN values for system operation. The default values for these VLANs are 1002, 1003, 1004 and 1005. You can change these values to another set of four consecutive VLANs, if required.
- Although the "Native VLAN" (VLAN 1) does not apply for traffic arriving at the E7 GPON port, the default VLAN 1 is off limits for user provisioning.
- IGMP snooping is provisioned on a VLAN basis.

- IGMP snooping is only enabled on the outer VLAN ID as multicast traffic is typically transported throughout the network with a single VLAN ID tag.
- For double-tagged VLANs, the same inner tag cannot be used with different outer tags. Any given (inner) tag received upstream on the PON can only push a single (outer) tag. For example, if VLAN 100 is received on the PON, tag 200 can be pushed, and then tag 200 will be popped in the downstream.
- Once a tag is used as a single-tagged VLAN, it cannot also be used in double-tagged VLANs. For example, if VLAN 300 is single tagged:
 - It cannot be used as the inner tag of a double-tagged VLAN because the double-tagged VLAN would require VLAN 300 to push an outer tag.
 - It cannot not be used as the outer tag in a double-tagged VLAN because then it would get popped in the downstream direction.
- For video, the specified P-bit must match the CoS type of bandwidth profile.
- The ONT does not support heterogeneous tagging in the MDU case for video service.
- In double-tagged VLANs, if the inner tag is applied to multiple ONT subscriber ports, it is an N:1 VLAN.
- For single-tagged VLANs, if the outer tag is applied to multiple ONT subscriber ports, it is an N:1 VLAN.
- If a given VLAN, single- or double-tagged is applied to a single ONT subscriber port, it is 1:1 VLAN.
- VLAN Association to voice services:
 - The IP Host object references a service tag action that indicates the voice service VLAN by one of the following methods:
 - An IP host can directly specify an outer tag, when the referenced service tag action states "Specified in Service" for the Outer Tag parameter.
 - This allows you to apply a different VLAN tag for the same service type (SIP, TDM gateway, H.248 gateway, MGCP, or T1 PWE3) on each ONT.
 - An IP host can reference a service tag action that specifies the service VLAN ID for the provisioned voice services.
 - This allows you to apply the same VLAN tagging for all ONTs that use the referenced service tag action object.
 - The VLAN assignment must be different between the VoIP IP host and the subscriber port. Sharing VLANs between a subscriber port and a VoIP IP host is not supported.

- For RG specific (Gateway) services (as opposed to half-bridge [HB] or full-bridge[FB]), only “Change Tag” and “Add and Change Tag” actions are supported. The concept of provisioning this way was derived from the same topology that is observed today in deployments that utilize external non-Calix RGs. This is especially used in scenarios where a third-party RG is downstream of a regular GPON ONT. From a deployment perspective, the HGW ONT then supports two logical components; a regular GPON ONT and an RG. In this scenario, the ONT would receive tagged traffic from the downstream RG and the ONT would either change the received tag and send it upstream single-tagged or either change the received tag, send it upstream, and the OLT would then add the second tag.

Example service-tag action configurations for data services

VLAN-per-service

Applies the specified service-tag action to the matched traffic.

Matches all untagged traffic on the ONT Ethernet port.

- Adds VLAN ID 300
- Marks with P-bit 0

The upstream traffic on a service VLAN must have the same P-bit value as is set for downstream traffic.

VLAN-per-subscriber

Applies the specified service-tag action to the matched traffic.

Double tags all untagged traffic arriving on the ONT port.

- Outer tag is explicitly set to 300
- Inner tagged is defined when adding the service to a port
- Upstream traffic is marked with P-bit 0

The upstream traffic on a service VLAN must have the same P-bit value as is set for downstream traffic.

Parameters

You can provision the following parameters for service tag actions:

Parameter	Description	Valid Options
Name*	Descriptive name of service tag action.	text string
Action*	Specifies the tag action to perform if the match criteria is met.	Add Tag, Add 2 Tags, Change Tag, Add and Change Tag
Match List*	Name and ID of the service match list to use for this tag action.	Any established service match list
S-VLAN (Outer Tag)	VLAN ID for the new outer tag. The specified VLAN must already be provisioned in the system. Alternately, select use the tag value specified in the ONT port service provisioning "use-svc" indicates that the VLAN ID will be specified when you provision a service on the ONT port.	Specified in Service, Enter Value (1-4093)
C-VLAN (Inner Tag)	Name of VLAN for the new inner tag. The specified VLAN must already be provisioned in the system. Alternately, "use-svc" indicates that the VLAN ID will be specified when you provision a service on the ONT port. Note: This parameter only applies if the selected action is "Add 2 Tags" or "Add and Change Tag."	Specified in Service, Not Used, Enter Value (1-4093)
Outer P-Bit Source	Specifies where to derive the P-bit value for the outer tag. You can either promote the P-bit from the subscriber tag, or set an explicit P-bit value in the tags. The upstream traffic on a service VLAN must have the same P-bit value as is set for downstream traffic. The P-bit value must also be consistent with the class of service type, as defined in the PON Class of Service (CoS) table. Note: If you select "Map a layer-3 Priority," then, when creating the service, you must select a PON CoS parameter of User1-4 or CoS1-4. Do not leave the system-default parameter of "derived" as there is no single unique P-bit specified from which to derive the intended CoS.	Specify P-bit (default), Map a layer-3 priority, From CoS
Outer P-Bit	Value to use for P-bit marking the outer tag when the Outer P-bit Source is set to "Specify P-Bit." Copy indicates that the E7 copies the P-bit value set in service provider supplied Residential Gateways (RGs). This option copies the subscriber-side Q-Tag P-Bits to a single C-tag or S-Tag, or to the inner C-Tag of a double-tagged Service. Note: If you select "copy," then, when creating the service, you must select a PON CoS parameter of User1-4 or CoS1-4. Do not leave the system-default parameter of "derived" as there is no single unique P-bit specified from which to derive the intended CoS.	0-7, copy
Inner P-Bit Source	Specifies where to derive the P-bit value for the inner tag. You can either preserve the P-bit from the subscriber tag, or set an explicit P-bit value in the tags. Note: This parameter only applies for GPON when the selected action is "Add 2 Tags" or "Add and Change Tag."	same-as-outer (default), Specify P-Bit
Inner P-Bit	Value to use for P-bit marking the inner tag when the Inner P-bit Source is set to "Specify P-Bit." Preserve indicates that the E7 preserves the P-bit value set in service provider supplied Residential Gateways (RGs). This option copies the subscriber-side Q-Tag P-Bits to a single C-tag or S-Tag, or to the inner C-Tag of a double-tagged Service. Note: This parameter only applies if the selected action is "Add 2 Tags" or "Add and Change Tag."	0-7, preserve
CoS	CoS ID to use for P-bit marking the outer tag when the P-bit Source is set to "Specify CoS." This value identifies the PON class of service.	1-4

Parameter	Description	Valid Options
Untagged P-Bit	When untagged upstream frames with the preserve or copy parameters set for inner or outer P-bits, this "untagged P-bit" value sets the P-bit on the untagged frame. Note: This parameter only applies for GPON when the selected action is "Add 2 Tags" or "Add and Change Tag."	0-7
For E7	(Viewable in CMS only.) Indicates that the global CMS profile is for use with the E7.	Y, N
For AE ONT	(Viewable in CMS only.) Indicates that the global CMS profile is for use with AE ONTs.	Y, N
Local S-VLAN (Outer Tag)	(Viewable in CMS only.) If enabled, this field allows the global CMS profile to be downloaded/synchronized successfully even if the S-VLAN ID (Outer Tag) value is provisioned differently on the local E7 node (via a local E7 profile with the same name). This allows a service provider's OSS to provision a service without managing S-VLANS per E7; a single global profile can be used across all E7s even if the S-VLANS differ from network to network. Y: During profile sync, CMS will not check the S-VLAN (Outer Tag) in the local E7 profile. N: During profile sync, CMS will check the S-VLAN (Outer Tag) in the local E7 profile, and successful synchronization will only result if the S-VLAN (Outer Tag) values are the same.	Y, N (default)

*Required fields

xDSL Traffic Management

The E-Series provides per-service hierarchical QoS to deliver uncompromised triple play services to subscribers. A powerful collection of classification, policing, queuing and scheduling algorithms let operators manage per-subscriber and per-service traffic flows to maintain priority, delay and loss differentiation within the E-Series network.

The E-Series supports the ability to classify ingress traffic based on the P-bit, VLAN, OUI (Source MAC for E7) or DSCP (and IP Precedence) values delivered by the modem or Residential Gateway (RGW). The P-bit values are defined per IEEE 802.1p, which is the specification used to give layer 2 Ethernet switches the ability to prioritize traffic. The E-Series supports a standard queuing structure, which defines 8 priority queues (0-7) and specifies that the highest priority queue as 7, with the lowest priority queue as 0. This quality of service scheme will allow service providers to determine and manage the priority of traffic within the network throughout the entire Ethernet network.

As part of the classification of traffic the E-Series will examine the P-bit or VLAN value depending on the provisioned classifications.

Egress queues on the DSL port deliver the traffic from the network using a strict priority queuing scheme (with minimum COS bandwidth guarantees if desired). Strict priority queuing ensures that all traffic (in the queues) with the highest priority will be delivered through the network first, followed by lower priority traffic.

Traffic Management and Rate Limiting on Ingress

In order to ensure proper Ethernet network management, the E-Series has been designed to support rate limiting and traffic shaping. The primary application for this functionality is to ensure level agreements and provide control over the amount of bandwidth utilized for service delivery, especially for TLS type services. This functionality allows the E-Series to limit the rate at which a subscriber can send traffic to the network and prevents a negative impact on the network from one high bandwidth user. This also enables the service provider to offer tiered services and gain additional revenue from customers requiring more bandwidth.

The E-Series uses two traffic control constructs that are well known in the data switching industry, called the class map and policy map. The class map defines traffic into groups, such as VOIP traffic. The policy map then matches the classes from the class map with information about how much bandwidth or what priority is desired for that grouping of traffic.

Traffic Shaping

Traffic shaping ensures that traffic leaving the E-Series network does not exceed a particular bandwidth requirement. When traffic is shaped it is sent at a particular rate by controlling the queuing of packets. If the queue overflows, the traffic is dropped at the tail of the queue. The E7 supports the ability to shape traffic on a per Class of Service (CoS) basis with up to 8 COS queues supported.

The ability to shape on a per VLAN basis allows for a service provider to use the same physical uplink port for different businesses or residential customers and shape each customer's traffic to their allotted limit. The benefit of shaping is that the bursting nature of the traffic is reduced and as long as there isn't a substantially large burst that overflows a queue, traffic is not dropped.

Rate Limiting

The E-Series also supports rate limiting, which is similar to traffic shaping and is sometimes referred to as policing because it ensures traffic does not exceed a specified bit rate. Rate limiting is used for traffic entering the E-Series network from a downstream device. It enables traffic to be controlled at a specified rate and keeps a subscriber from exceeding their data rate contract. Traffic shaping is also used to smooth traffic (not drop); enabling a downstream device (i.e. modem) to process the traffic without overflowing traffic queues.

A side effect of rate limiting is that connection oriented protocols, such as FTP, will send an acknowledge (ACK) to indicate receipt of a packet. In certain scenarios, the ACK can actually be dropped causing a retransmission of the original packet even though it may have been successfully delivered. If FTP is being used to determine the actual throughput of an interface, the results can be distorted by the retransmission of traffic making the throughput look substantially lower than the actual capability of the interface.

Ethernet Bandwidth Profile for xDSL Services

An Ethernet bandwidth profile is used to shape data traffic at a consistent rate by specifying the upstream and downstream bandwidth rates to apply to individual Ethernet services on xDSL ports, even when the modem is trained at a higher rate. Typically, a single bandwidth profile is applied to many subscriber ports. Yet, for each level of data service, you need an corresponding Ethernet Bandwidth Profile.

The Ethernet bandwidth profile on E-series VDSL2 cards and nodes rate limits traffic as defined by the VLAN and the P-bit value settings applied to the Ethernet Service. For example, if a bandwidth profile is applied to a service with a tag action that adds an outer VLAN ID of 901 and a P-bit value of 0, traffic is rate limited for both upstream and downstream against VLAN 901/P-bit 0. All traffic downstream that matches this VLAN/P-bit combination is rate limited. If video is sent downstream with a different VLAN and/or P-bit value, it is not rate limited. Note that if the P-bit value is not specifically assigned by the service, all traffic in the VLAN is rate limited according to the setting in the Ethernet bandwidth profile.

Each direction of a traffic shaper can be configured to provide a specific rate for the shaper. Ethernet bandwidth profiles specify two settings to designate the rate:

- **Committed Information Rate (CIR):** The amount of committed (guaranteed) bandwidth for upstream traffic on the interface.
CIR ≥ 0
- **Peak Information Rate (PIR):** The maximum amount of bandwidth for upstream or downstream traffic, in excess of the CIR that the interface supports, if available. PIR is non-guaranteed bandwidth.
PIR > 0

Commonly-defined Per-Hop Behavior (PHB) traffic classes:

- **Default PHB (Per hop behavior)**—typically best-effort traffic
A default PHB is the only required behavior. Essentially, any traffic that does not meet the requirements of any of the other defined classes is placed in the default PHB. Typically, the default PHB has best-effort forwarding characteristics. The recommended DSCP for the default PHB is '000000'(0).
- **Expedited Forwarding (EF) PHB**—dedicated to low-loss, low-latency traffic
The IETF defines Expedited Forwarding behavior in RFC 3246. The EF PHB has the characteristics of low delay, low loss and low jitter. These characteristics are suitable for voice, video and other real-time services. EF traffic is often given strict priority queuing above all other traffic classes. The recommended DSCP for expedited forwarding is '101110' (46).
- **Assured Forwarding (AF) PHB**—gives assurance of delivery under prescribed conditions
The IETF defines the Assured Forwarding behavior in RFC 2597 and RFC 3260. Assured forwarding allows the operator to provide assurance of delivery as long as the traffic does not exceed some subscribed rate. Traffic that exceeds the subscription rate faces a higher probability of being dropped if congestion occurs.
The AF behavior group defines four separate AF classes. Within each class, packets are given a drop precedence (high, medium or low). The combination of classes and drop precedence yields twelve separate DSCP encodings from AF11 through AF43. Usually, traffic policing is required to encode drop precedence. Typically, all traffic assigned to a class is initially given a low drop precedence. As the traffic rate exceeds subscription thresholds, the policer will increase the drop precedence of packets that exceed the threshold.
- **Class Selector PHBs**—maintains backward compatibility with the IP Precedence field

Traffic Class Type	E7 P-bits (default)	Service Category	E7 BW Profile
Expedited Forwarding (EF)	5, 6, 7	Network control, voice, T1/E1	CIR>0, PIR=CIR (PIR is provisioned for "0" to represent that there is no PIR independent of the CIR) Committed Information Rate (CIR): Guaranteed level of bandwidth with low Delay and Delay Variation.
Assured Forwarding 1 (AF1)	4	Video	CIR>0, PIR>0, PIR>=CIR Committed Information Rate (CIR) + Peak Information Rate (PIR). Both guaranteed and non-guaranteed bandwidth levels. Moderate Delay and Delay Variation.
Assured Forwarding 2 (AF2)	3	Application signaling, TLAN	CIR>0, PIR>0, PIR>=CIR Committed Information Rate (CIR) + Peak Information Rate (PIR). Both guaranteed and non-guaranteed bandwidth levels. Moderate Delay and Delay Variation.
Best Effort (BE)	0, 1, 2	High-speed internet	CIR=0, PIR>0 Peak Information Rate (PIR): No guaranteed bandwidth levels. Higher Levels of Delay and Delay Variation.

Configuration guidelines

- You can create up to 100 custom Ethernet bandwidth profiles.
- The CIR and PIR values must be consistent with the traffic priority settings in the following:
 - Policy maps applied to ingress ports
 - Service tag actions applied to xDSL services
 - The priority of traffic as defined in the Layer 3 DSCP table or IP precedence table
 - or
 - The P-bit value specified in the service tag action
- The E7 calculates a burst size for the traffic rate limiter when the value is not set for this parameter.
 Auto-calculated burst size = peak rate * 80msec for the given rate (PBS = CBS)
 This allows the flow that is being rate limited to exceed its bandwidth profile up to CIR within the 80msec window. The established guideline for this parameter is to reflect double the Round Trip Delay (RTD) of the network, with 40 ms used as the default RTD. If the burst size is set to a value less than the MTU size of the interface, the system adjusts the burst size to the MTU size.
- The simplest approach to account for committed bandwidth and guarantee its delivery, is to allocate guaranteed egress bandwidth equal to the sum of the ingress CIR values. Note there are multiple congestion points that need to be accounted for on the VDSL2 card.
- For VDSL2-based cards, the minimum-shapeable CIR/PIR for a transmit profile is a dynamic value determined by the provisioned maximum downstream (DS) train rate of the port (~3% of the maximum DS rate). For example:

Proprietary Information: Not for use or disclosure except by written agreement with Calix.

Maximum DS Train Rate	Minimum DS Shapeable Rate
50M	1.5M
40M	1.2M
30M	900K
20M	600K
10M	300K
5M	150K
1.5M	45K

Note: If the CIR/PIR is lower than the minimum-shapeable rate, it is not enforced; downstream traffic is allowed up to the minimum rate.

Parameters

You can provision the following parameters for Ethernet bandwidth profiles:

Parameter	Description	Valid Options
Name*	Name of the bandwidth profile.	
Committed Rate for Upstream	Committed information rate for upstream traffic. Where rates may be specified as follows: - In 64 kbps increments up to 2 Mbps - In 1 Mbps increments between 2 Mbps to 1000 Mbps 0 kbps disables the meter Use "m" suffix for Mb/s or "g" for Gb/s in whole number increments.	0-2048kbps, 0-1000Mbps, 1Gbps kbps in 64k increments
Peak Rate for Upstream	Peak information rate for upstream traffic. Where rates may be specified as follows: - In 64 kbps increments up to 2 Mbps - In 1 Mbps increments between 2 Mbps to 1000 Mbps 0 kbps disables the meter Use "m" suffix for Mbps or "g" for Gbps in whole number increments.	0-2048kbps, 0-1000Mbps, 1Gbps kbps in 64k increments
Peak Rate for Downstream	Peak information rate for downstream traffic. Where rates may be specified as follows: - In 64 kbps increments up to 2 Mbps - In 1 Mbps increments between 2 Mbps to 1000 Mbps 0 kbps disables the meter Use "m" suffix for Mbps or "g" for Gbps in whole number increments.	0-2048kbps, 0-1000Mbps, 1Gbps kbps in 64k increments
Downstream Peak Burst Size	Peak burst size for downstream traffic. - Use "k" suffix for Kbytes (4 to 16000kbps) - Use "m" suffix for Mbps (0 to 16 mbps) - Use whole number increments	4-16000 Kbps

Parameter	Description	Valid Options
Upstream Committed Burst Size	Committed burst size for upstream traffic. - Use "k" suffix for Kbytes (4 to 16000kpbs) - Use "m" suffix for Mbps (0 to 16 mbps) - Use whole number increments	4-16000 Kbps
Upstream Peak Burst Size	Peak burst size for upstream traffic. - Use "k" suffix for Kbytes (4 to 16000kpbs) - Use "m" suffix for Mbps (0 to 16 mbps) - Use whole number increments	4-16000 Kbps

*Required field

DSCP Map

To designate the DSCP profile values as the ingress traffic access list, do the following:

1. In the service tag action that is going to be applied to the data or video service, select **Map a Layer-3 priority** from the P-Bit Source list.
2. On the xDSL port associated interface where you are going to be provisioning the data or video service, select the name of the DSCP profile to use from the DSCP/IP Precedence Profile list.

Note: The upstream traffic on a service VLAN must have the same priority value as is set for downstream traffic.

Note: If IP traffic on an Access Interface has a DSCP setting other than the values included in the assigned DSCP map, the E7 system interprets this condition as a mismatch and uses a "Default" P-bit value from the assigned DSCP Map profile. For example, if you were using the system-default DSCP Map "Access" and the incoming frames had a DSCP value of 001001, that value is not included in the assigned DSCP map, so the E7 would use the default P-bit value of 0. However, you can create a DSCP Map with a Default P-bit value of your choice, and then assign the DSCP Map to xDSL port interfaces.

The system-default DSCP profile "Access," has values as follows:

DSCP Value	Meaning	Decimal	Drop Probability	802.1p Priority-Bit
000 000	BE-CS0	0	n/a	0
001 000	CS1	8		1
001 010	AF11	10	Low	1
001 100	AF12	12	Medium	1
001 110	AF13	14	High	1
010 000	CS2	16		2
010 010	AF21	18	Low	2
010 100	AF22	20	Medium	2
010 110	AF23	22	High	2
011 000	CS3	24		3
011 010	AF31	26	Low	3
011 100	AF32	28	Medium	3
011 110	AF33	30	High	3
100 000	CS4	32		4
100 010	AF41	34	Low	4
100 100	AF42	36	Medium	4
100 110	AF43	38	High	4
101 000	CS5	40		5
101 110	EF	46		5
110 000	CS6	48		0
111 000	CS7	56		0
	Default			0

Best Effort (BE)

Assured Forwarding (AF) gives assurance of delivery under prescribed conditions

Expedited Forwarding (EF) dedicated to low-loss, low-latency traffic

Default = P-bit assigned for a DSCP value that is not included in the DSCP Map

Match Lists for xDSL Services

The VDSL2 card performs the following two actions:

- **Classification** – the service match lists are for matching and classifying traffic into services.
- **Marking** – the service tag actions are for marking the traffic with the appropriate VLAN and Priority required by the service.

A match list defines a set of criteria (matching rules) for classifying xDSL traffic into services. A match list can contain both "tagged" and "untagged" match rules, up to 12 tagged rules and up to 16 untagged rules for each xDSL port. The match list defines how the xDSL port classifies subscriber traffic into a service VLAN by associating it in a service tag action that is then associated with an xDSL Ethernet service.

- **Untagged match rules** can match on the following:
 - A portion of the source MAC address as indicated by the Source MAC and Source MAC Mask attributes of a video Set-Top box (not supported with T-Series ONTs)
 - The Ethernet type to distinguish video and High-Speed Internet (HSI) untagged traffic (not supported with P-Series and GigaCenter ONTs)
 - All untagged traffic by using the system default match list "all-untagged"
 - PVC VPI and VCI port numbers (only applies to VDSL2 services)
- **Tagged match rules** can match on any combination of the following:
 - Outer tag
 - VLAN-ID
 - P-bit values
 - Tag Protocol Identifier (TPID) can also be specified at the E-Series egress Ethernet port interface, but defaults to 0x8100

Example of multiple match rule use for classifying traffic:

A service provider wants to use a single ONT Ethernet port to connect to a home network with a combination of IPTV and data traffic. The IPTV traffic is identified by using the OUI value of the video set top box.

Note: The T-Series ONTs do not support configuring a match rule with an OUI filter.

This combination of rules can accomplish the example task:

1. An untagged traffic match rule would be applied with a specific MAC OUI match criteria.
2. A second untagged match rule would be applied with a "match any" criteria.

Using a match list and service tag action, Calix ONTs support matching the traffic from all set-top boxes (STB) of a particular manufacturer by indicating the STB OUI in the first half of the specified Source MAC attribute followed by zeros, and then providing a mask for the STB OUI by indicating ff:ff:ff:00:00:00 in the Source MAC Mask attribute.

Example entries that match all STB source MAC addresses of a particular manufacturer:

- ADB set-top boxes: 00:03:91:00:00:00
- Amino set-top boxes: 00:02:02:00:00:00

Note: If STBs from multiple vendors are being used, create additional untagged match rules, entering the OUI for each STB vendor.

Alternatively, the ONT's support matching the traffic from a specific host by indicating the MAC address in the Source MAC attribute, and then providing a mask for the address by indicating ff:ff:ff:ff:ff in the Source Mac Mask attribute.

Configuration guidelines

- If a residential gateway (RG) at the subscriber's premise classifies traffic into VLANs, then for each service type, create a match list with a tagged rule that matches the service VLAN.
- Because only video services are supported on the P-series FB interface, it is not necessary to specify the OUI of the STB, and the all-untagged match list applies.
- Up to 32 OUIs are supported per match list.
- Ethertype is the only untagged traffic classifier supported by the T-Series ONT's.
- The ONT's do not support tagged and untagged match rules on the same ONT for video services. Only one classification rule is allowed for video services per ONT.

Match list and rule parameters

You can provision the following parameters for match list and rules:

Parameter	Description	Valid Options
Name*	Name of service match list.	text string
S-VLAN (Outer Tag)	VLAN ID of the outer tag (tagged match rules only). VLANs can be specified by name or by numeric VLAN ID, which ranges from 1-4094. In addition, "untagged" indicates that only untagged traffic should be matched and "ignore" indicates that the VLAN ID should not be examined.	text string, 1-4094, untagged, ignore (default)
Outer Pbit	P-bit value of the outer tag (tagged match rules only). P-bit values are in the range 0-7. Alternately "pbit-none" means the P-bit value is not considered. If this parameter is not specified, "p-bit-none" is the default behavior.	0-7, ignore, p-bit-none(default)
Source MAC	Source MAC address (untagged match rules only).	six hexadecimal digits in the range 0-FF, optionally separated by colons Default = ignore

Parameter	Description	Valid Options
Source MAC Mask	Source MAC mask (untagged match rules only). - source MAC mask (x:x:x:x:x:x) - none (no MAC mask) - oui (mask OUI fields (FF:FF:FF:FF:FF:FF)) Note: The T-Series ONTs do not support configuring a match rule with an OUI filter.	six hexadecimal digits in the range 0-FF, optionally separated by colons, none (default), oui
Ethertype	Ethernet type to match (untagged match rules only). - any = default - pppoe = 0x8863 or 0x8864 (for HSI) - arp = 0x0806 (for video) - ipv4 = 0x0800 (for video) - ipv6 = 0x86DD (for video) Note: Not supported for P-Series and GigaCenter ONTs.	any (default), pppoe, arp, ipv4, ipv6

*Required field

Service Tag Actions for xDSL Services

Applying service tag actions to service provisioning performs the following two actions:

- **Classification** – the service match lists are for matching and classifying traffic into services.
- **Marking** – the service tag actions are for marking the traffic with the appropriate VLAN and P-bit required by the service.

This topic describes how to create service-tag actions that can be used for many subscriber ports that require the same actions to be performed. All services configured on an xDSL port exist in a VLAN-tagged service flow, requiring an associated service-tag action.

Service-tag actions are applied to subscriber service provisioning as follows:

- Video and data services provisioning on an xDSL port includes a reference to a service-tag action.
- Voice services provisioning on a Voice (POTS) port references an IP Host that sets the output service VLAN and points to a system default service-tag action that specifies the traffic as untagged and assigns the appropriate priority value.

The E7 can apply the following tag actions assigned to an xDSL port service.

Add a single tag	Adds a single outer tag to the matched traffic • Can be applied to tagged and untagged traffic. • Supported for DSCP/IP Precedence-to-P-bit translation.

Add 2 tags	Adds an inner and outer tag to the matched traffic* • Can be applied to untagged traffic, only. • Supported for DSCP/IP Precedence-to-P-bit translation.
Add and change tag	Adds an outer tag to the matched traffic and changes the inner tag • Can be applied to tagged traffic, only. • NOT supported for DSCP/IP Precedence-to-P-bit translation.
Change tag	Changes the outer tag on the matched traffic • Can be applied to tagged traffic, only. • Supported for DSCP/IP Precedence-to-P-bit translation.

*To properly process ingress double tags, the GE network interface (uplink) must be configured as a Trunk role.

Determining the tag-action values

When a service is provisioned at an xDSL port, a match list and service-tag action specify the classifying and marking of packets from the subscriber port into the service VLAN.

You have the following options for determining the added- or changed-tag values:

- **Specify the VLAN value explicitly.**

This option is used for a VLAN-per-service provisioning model (N:1). A service carried on an N:1 VLAN applies to multiple subscriber ports, where a single match list and tag action can indicate the service.

- **Reference the VLAN value from the service provisioning on the xDSL port interface.**

This option is used for a VLAN-per-port provisioning model (1:1). A service carried on a 1:1 VLAN is the same for each subscriber except the customer tag is unique for each single subscriber port, so you can define the match list and tag action pair such that multiple subscriber ports can reference it. This is accomplished by indicating "Specified in Service" for the output tag in the tag action, and then the subscriber-specific value for the output tag is defined when the service is provisioned on the port.

Note: The service provisioning process takes place after the creation of a service-tag action.

You also have the following options for determining the service-tag action P-bit for services provisioned on xDSL ports:

- **Specify the P-bit value explicitly.**

Enter the P-bit value to use for P-bit marking the outer tag.

- **Reference a Layer-3 Priority map.**

Use a DSCP table or an IP Precedence table to map to a P-bit value. From the xDSL port associated interface, you can select to use the system-default profiles ("access") or specify another custom profile. See Layer 3 Priority to P-Bit Mapping.

- Only one service per port can use DSCP-Bit mapping function.
- Works with Add-Tag, Add-2-Tags, and Change-Tag tag actions.
- Either DSCP or IP Precedence can be used for a service, not both.

Provisioning sequence:

- a. Create a DSCP-to-P-Bit map or create an IP Precedence to P-Bit Mask, if you do not want to use the system-default.
- b. Assign the map to the xDSL port associated interface, or leave the system-default map "access" assignment.
- c. Create a service-tag action with an Outer P-Bit Source of **Map a layer-3 Priority**.
- d. Create a service using the defined service-tag action.

- **Preserve the existing P-bit value upstream.**

Indicate that the E7 preserves the P-bit value set in service provider-supplied CPEs by specifying a P-bit value of "copy."

Note: The upstream traffic on a service VLAN must have the same priority value as is set for downstream traffic.

The following table shows the 802.1p to DiffServ DSCP pre-defined, system-default mapping table ("access").

Default P-bits	IP Precedence	DiffServ Code Point**	Service
0	7	CS6, CS7 (48-56)	Network control
0	6	EF (46)	Voice, NSP, CES
5	5	CS5 (40)	Network Management
4	4	CS4 & AF41-43 (32-38)	Video
3	3	CS3 & AF31-33 (24-30)	Data3
2	2	CS2 & AF21-23 (16-22)	Data2
1	1	CS1 & AF11-13 (8-14)	Data1
0	0	0	Best Effort

**The E7 supports alphanumeric DSCP values that denote forwarding classes such as AF21, AF31 in addition to numeric value assignment.

VLAN Traffic flow in xDSL ports

For xDSL ports, rather than adding a port interface to a VLAN membership to enable traffic flow, a service-tag action must be created that specifies the VLAN and traffic priority. This service-tag action is then referenced when the service is provisioned on the xDSL port. The VLAN must already be created on the E7 for an xDSL port to pass traffic carried on a VLAN.

- To view the VLANs associated with specific xDSL port, on the Navigation Tree, click the xDSL port of interest, and then click **Associated Interface > VLANs**.
- To view the xDSL services associated with a specific VLAN, on the Navigation Tree, click **VLANs**, click the particular VLAN from the list that appears in the Work Area, and then click the **Service Associations** tab.

The xDSL ports support the same VLAN services as the E7 Ethernet ports, including IEEE 802.1Q VLAN tagging and IEEE 802.1ad VLAN stacking (Q-in-Q). The following service models are supported.

- **N:1/VLAN-per-service model:** each service is assigned a dedicated VLAN where multiple subscriber ports are assigned to the VLAN for a single service. This model is often referred to as N:1. An additional tag can be added to limit the size of a group for improved reliability. A service carried on an N:1 VLAN applies to multiple subscriber ports, allowing a single match list and tag action to be used to denote the service.

- Subscriber tag = 110
- Match tag = 110
- Change tag = 200 (service VLAN)

- **1:1/VLAN-per-port model:** the traffic frames from each subscriber port are assigned a unique VLAN ID (tag). The tag is applied to a frame before the traffic is aggregated for transport. An additional tag can also be added to group subscriber ports in logical categories.

A service carried on a 1:1 VLAN is the same for each subscriber, except the customer tag is unique per subscriber. You can define the match list and tag action pair such that multiple subscriber ports can reference it. This is accomplished by assigning an attribute for the Output Tag in the tag action that indicates the value of the output tag is subscriber specific (Specified in Service). The customer specific tag is included in the service definition when provisioning the service.

- Subscriber tag = 110
- Match tag = 110
- Change tag = 110
- 1:1 VLAN-per-Port, Single Tagged – Subscribers on an E7 system are provisioned with individual Customer Tags (C-Tags). Add tag, or change tag action.
- 1:1 VLAN-per-Port, Double Tagged – Subscribers on an E7 system are provisioned with individual Customer Tags (C-Tags). The S-Tag can be added by the service tag action. Add 2-tags, or change-and-add tag action.

The 1:1 model does not require any new or different traffic handling techniques to isolate and manage subscriber traffic. With the 1:1 service delivery model, each subscriber's traffic (except broadcast video) is sent down a dedicated single VLAN on a per subscriber basis. When ATM encapsulation is used for ADSL/ADSL2+ modes, the single ATM PVC maps to a single VLAN.

Configuration guidelines:

- The P-bit value specified in the service tag action must be consistent with the DiffServ Code or IP Precedence as defined in the default profiles, or custom profiles referenced at the xDSL interface.
- If the specified P-Bit source is "copy," the P-bit value cannot be specified because the P-bit value of the incoming traffic is honored.
- Each service tag action references a service-match list (classify) and specifies a tag action (mark and tag).
- Each service tag action is associated with a specific service.
- Multiple services and service tag actions can be applied to a single xDSL port. However, if multiple services on the same xDSL port use the same outer VLAN ID, the priority must be specified for each service in the associated service tag action using the **P-bit Source = Specify P-Bit**. Using the P-Bit Source = Map a layer-3 priority is not supported for this scenario.
- The same service tag action can be used on multiple xDSL ports.
- Each E7 shelf, or E7 Modular Chassis system has a capacity of 256 service tag actions. For configurations that require large numbers of VLANs, Calix recommends selecting the **Specified in Service** value for the "Outer Tag" parameter, and then specifying the VLAN ID when provisioning the service.
- For traffic to flow on the VLAN specified in a service tag action, the VLAN and uplink must already be created on the E7.
- For an xDSL port to pass traffic carried on a VLAN, a service tag action must be created specifying the VLAN and the VLAN must already be created on the E7.
- For installations using multiple PVC assignments, individual PVCs may be mapped to specific VLANs that are applied to the port. For example, installations that have multiple PVC assignments (e.g. 0/33, 8/33, 0/35) may be mapped to individual data services (e.g. "Data 1", "Data 2", "Data 3") that would then be applied to the port.
- If no matching criteria are assigned to a tag action, then the tag action is performed on all packets entering the xDSL port. For example, when a match list has a single untagged rule with the outer tag set to "ignore" and the Outer P-bit set to "Pbit-none."
- Traffic must have a priority assigned in order to be scheduled on the xDSL port.
- The "change-tag" action changes the VLAN ID in the outermost tag.
- For double-tagged packets, the outer VLAN ID must match a provisioned VLAN. The inner tag does not have to match a provisioned VLAN ID.
- Edge interfaces can only use single-tag actions when DHCP Snooping is enabled on the associated service VLAN.

- Within the xDSL subsystem, MAC learning and switching occurs within a single, outer VLAN ID, referred to as the VLAN C-tag. Packets on the xDSL port always have a C-tag, which consists of a provisioned VLAN ID and priority for classifying subscriber traffic and adding/modifying the C-tag.
- The E-Series reserves four VLAN values for system operation. The default values for these VLANs are 1002, 1003, 1004 and 1005. You can change these values to another set of four consecutive VLANs, if required.
- Although the "Native VLAN" (VLAN 1) does not apply for traffic arriving at the xDSL port, the default VLAN 1 is off limits for user provisioning.
- IGMP snooping is provisioned on a VLAN basis.
- IGMP snooping is only enabled on the outer VLAN ID as only single-tag multicast VLANs are supported.
- The following tag actions are supported for VLANs with IGMP enabled:
 - Add-tag (only supported for untagged traffic)
 - Change-tag
- Services provisioned with DSCP or IP Precedence translation to P-bit values only support:
 - add-tag
 - add-2-tags
 - change-tag (R2.6.30 and higher)
 - One service per port can use the DSCP or IP Precedence to P-bit mapping function
 - Either DSCP or IP Precedence can be used for a service, not both
- Once a tag is used as a single-tagged VLAN, it cannot also be used in double-tagged VLANs. For example, if VLAN 300 is single tagged:
 - It cannot be used as the inner tag of a double-tagged VLAN because the double-tagged VLAN would require VLAN 300 to push an outer tag.
 - It cannot be used as the outer tag in a double-tagged VLAN because then it would get popped in the downstream direction.
- In double-tagged VLANs, if the inner tag is applied to multiple xDSL subscriber ports, it is an N:1 VLAN.
- For single-tagged VLANs, if the outer tag is applied to multiple xDSL subscriber ports, it is an N:1 VLAN.
- If a given VLAN, single- or double-tagged is applied to a single xDSL subscriber port, it is 1:1 VLAN.

Example service-tag action configurations for data services

VLAN-per-service

Applies the specified service-tag action to the matched traffic.

Matches all untagged traffic on the xDSL port.

- Adds VLAN ID 300
- Marks with P-bit 0

The upstream traffic on a service VLAN must have the same P-bit value as is set for downstream traffic.

VLAN-per-subscriber

Applies the specified service-tag action to the matched traffic.

Double tags all untagged traffic arriving on the xDSL port.

- Outer tag is explicitly set to 300
- Inner tagged is defined when adding the service to a port
- Upstream traffic is marked with P-bit 0

The upstream traffic on a service VLAN must have the same P-bit value as is set for downstream traffic.

Before starting

Before starting the service tag action creation process, check that the following conditions are met:

- All VLANs that are to be specified in the service tag action have already been created.
- The match list that is to be specified in the tag action has already been created.
- The DSCP or IP Precedence values are defined, if using a map other than the system default.

Parameters

You can provision the following parameters for service tag actions:

Parameter	Description	Valid Options
Name*	Descriptive name of service tag action.	text string
Action*	Specifies the tag action to occur.	Add Tag, Add 2 Tags, Change Tag, Add and Change Tag
Match Criteria List	Name and ID of the service match list to use for this tag action.	Any established service match list
Outer Tag	VLAN ID for the new outer tag. The specified VLAN must already be provisioned in the system. Alternately, "use-svc" indicates that the VLAN ID will be specified when you provision a service on the VDSL port.	Specified in Service, Enter Value (1-4093)
Inner Tag	Name of VLAN for the new inner tag. Alternately, "use-svc" indicates that the VLAN ID will be specified when you provision a service on the VDSL port. Note: This parameter only applies if the selected action is "Add 2 Tags" or "Add and Change Tag."	Specified in Service, Not Used, Enter Value (1-4093)
P-Bit Source	Specifies where to derive the P-bit value for the outer tag, selecting one of the following: - Specify P-bit - Sets an explicit P-bit value in the tags. Note: The upstream traffic on a service VLAN must have the same priority as is set for downstream traffic. - Map a layer-3 priority - Maps to either a DSCP map or IP Precedence values. Note: The DSCP map or IP Precedence map is specified in the xDSL port interface.	Specify P-bit ^(default) , Map a layer-3 priority, From CoS (GPON only)
P-Bit	The P-bit value to use for P-bit marking the outer tag when the P-bit Source is set to "Specify P-bit." Note: When the value of "copy" is selected, the E7 honors the P-bit value of the incoming traffic and passes the existing P-bit value upstream.	0-7, copy 0 (default)
For E7	(Viewable in CMS only.) Indicates that the global CMS profile is for use with the E7.	Y, N
For AE ONT	(Viewable in CMS only.) Indicates that the global CMS profile is for use with AE ONTs.	Y, N
Local S-VLAN (Outer Tag)	(Viewable in CMS only.) If enabled, this field allows the global CMS profile to be downloaded/synchronized successfully even if the S-VLAN ID (Outer Tag) value is provisioned differently on the local E7 node (via a local E7 profile with the same name). This allows a service provider's OSS to provision a service without managing S-VLANS per E7; a single global profile can be used across all E7s even if the S-VLANS differ from network to network. Y: During profile sync, CMS will not check the S-VLAN (Outer Tag) in the local E7 profile. N: During profile sync, CMS will check the S-VLAN (Outer Tag) in the local E7 profile, and successful synchronization will only result if the S-VLAN (Outer Tag) values are the same.	Y, N (default)

*Required fields

Chapter 5

Subscriber Services

This chapter provides engineering and planning information regarding E-Series subscriber services.

Topics Covered

This chapter covers the following topics:

- Ethernet Elements
- IPv4 Security Features
- IPv6 Support
- MEF Compliant Service
- Data Services
- Video Services
- VoIP Services
- RF Overlay Video Service
- T1/E1 PWE3 Service

Ethernet Elements

Many services use the same Ethernet and IP objects regardless of the final configuration. This section describes the following objects and their configurable parameters:

- Ethernet Interfaces
- Ethernet Ports
- VLANs
- Tag Actions on an Ethernet Interface
- Service Tag Actions on an ONT or xDSL Ethernet Interface

Ethernet Interfaces

Each GE and 10GE port, as well as each link aggregation group (LAG), is mapped to an Ethernet port interface within the E-Series. An Ethernet port interface is a logical construct to which VLAN memberships and tag actions are assigned, QoS is applied, and topology protocols are enabled. Two interface roles are defined to combine common configuration options and improve usability of the system.

Interface names

- Eth interfaces (Non-LAG related interfaces) share the same name as the E7 Ethernet Ports (card 1/Eth port g1, for example 1/g1)
 - g(port number) = GE(port number)
 - x(port number) = 10GE(port number)
- LAG interfaces are named at the time of creation

Interface roles

Each GE and 10GE interface in the has one of the following configuration-role types:

Trunk: A port connecting to other equipment belonging to the service provider or to another service domain with consistent VLAN tagging levels. These ports may also be referred to as Network ports or Provider ports in industry standards. These ports support outer VLAN tag plus MAC switching.

Examples of trunk ports are 10G ERPS transport ports and GE uplinks. Trunk ports can be configured for link aggregation, RSTP, or ERPS.

Edge: A port facing customer equipment or facing reduced functionality devices, alternative administrative domains, or managed CPE. Generally, this Ethernet port interface is where all classification is first performed on ingress traffic (if customer facing). The Ethernet interface is also expected to add, replace, or remove one or more VLAN tags on edge traffic. RSTP, and LAG networking protocols are supported.

Examples of an edge port would include GE ports to managed CPE, a GE/10GE port to external equipment which may use different tagging levels, or GPON ports.

Access: A port facing untrusted customer equipment or other devices serving subscribers. Generally, this port interface is where individual subscriber services are defined and enforced (bandwidth limits, security, multicast profiles). Networking protocols are not supported.

Examples of an access port would be point-to-point connections to subscribers or other devices serving subscribers.

VLAN support

By default, every Ethernet port with a trunk or edge interface on the unit is a member of VLAN 1, the Native VLAN. The Native VLAN is available to pass any untagged traffic. You can provision an Ethernet port interface to use a different existing VLAN as the Native VLAN.

To forward untagged traffic on Ethernet ports with an access interface, an add-tag action must be applied to untagged frames, assigning the traffic to a designated VLAN.

Tagged traffic that does not match any of the tagging criteria is dropped.

For modular chassis nodes, any VLAN created on the system is automatically mapped to the Stacking Ports. The remaining port interfaces in the system must be a VLAN member for traffic to pass on the VLAN through the interface.

Interface role configuration guidelines

	Trunk	Edge	Access
E7-2	X	X	X
E7-2 MCC	X	X	X
E7-2 MCE			X
E7-20 SCP	X	X	
E7-20 GE-24			X
ONT Ethernet Port			X
Tag Actions		X	X
Native VLAN	X	X	
Networking Protocol (RSTP, ERPS*, LAG)	X	X	

*Edge ports only support RSTP and LAG networking protocols, NOT ERPS.

Link aggregation groups

The E-Series uses IEEE 802.3ad Link Aggregation (LAG) to bond multiple GE or 10GE ports into a single link aggregation group with a single logical Ethernet interface. Link aggregation expands the bandwidth of the logical interface beyond the capacity of a single link and also provides link redundancy between devices, even if one of the links fails. The ports that comprise a LAG can be on the same line card, or on two different line cards in the same E7 shelf.

Link aggregation provides two principle values to the network operator:

- Point to point link protection / redundancy between network elements
- Bandwidth expansion of the logical interface beyond the capacity of a single link

BPDU Guard

Spanning Tree BPDU Guard feature is commonly implemented in Ethernet switches to prevent the participation of an interface in a spanning tree. The BPDU Guard protects the network from unwanted spanning tree BPDUs that may impact network stability.

BPDU Guard is enabled and disabled on an interface basis. When the feature is enabled, the interface can detect received spanning tree protocol BPDUs, causing the interface to be put into a disabled state and report a system alarm. The operator must manually intervene to disable and then re-enable the interface to put the interface back into service. If the interface receives another BPDU it will be disabled again.

When BPDU guard is enabled, spanning protocol is disabled – the two are mutually exclusive. The default configuration for RSTP and BPDU guard is as follows:

• Stand alone system shelf:	RSTP enabled on GE/10GE ports, BPDU guard disabled
• Modular Chassis Control shelf:	RSTP enabled on GE/10GE ports, BPDU guard disabled
• Modular Chassis Expansion shelf:	RSTP disabled on GE/10GE ports, BPDU guard enabled

Configuration guidelines

Follow these guidelines when configuring an Ethernet port interface or LAG interface:

- An Ethernet interface always exists, cannot be deleted, can be modified, and is associated with a specific companion Ethernet port (GE or 10GE).
- A port is always a member of exactly one interface, even when it is being used in a standalone manner.
- A LAG interface can be created, deleted, and modified.
- The ports that comprise a LAG can be on the same line card or on two different line cards in the E7 same shelf, but, must be of the same module type. For example, either all XFP or all SFP+.
- For E-Series (E5-48/E5-48C) products, LAG supports only two ports in a group, either ports 1 and 2, or ports 3 and 4. LAG group may NOT include ports across the pair boundary. For example, a LAG with ports 1 and 3, or 2 and 3 are not supported.

- You can configure only one of the following attributes on a given VLAN on a given interface:
 - Trunk interfaces:
 - VLAN member
 - Native VLAN
 - Edge interfaces:
 - VLAN member
 - Tag-action
 - Native VLAN
 - Access interfaces:
 - VLAN member
 - Tag-action
 - Trunk and Edge interfaces are always associated with at least one VLAN, through the native VLAN attribute (VLAN 1, by default). Access interfaces do not support a Native VLAN, however, tag actions can be used to assign untagged traffic to a VLAN ID.
 - All network connections (ERPS, RSTP, LAG, DHCP servers, multicast routers, IGMP-enabled video servers, network-facing routers) are made on Trunk or Edge interfaces, only.
 - Edge interfaces can only use single-tag actions when DHCP Snooping is enabled on the service VLAN. Calix recommends that uplink interfaces, use the Trunk role.
 - Interfaces can be associated with additional VLANs through memberships or tag-actions.
 - If an Ethernet uplink Edge interface with tag actions (for double-tagged traffic) is on the same card as Ethernet downlink Edge interfaces (that are members of the outer VLAN), ingress double-tagged traffic from the downlink ports will not flow upstream, resulting in a service interruption. As a workaround, for this application, ensure that one of the following configurations are used:
 - The Ethernet uplink Edge port is on a different card than downlink Edge ports.
 - Ethernet downlink ports are configured with the Trunk interface role.
- Note:** For this application, if active/standby LAG is used for the uplink, you must ensure that Ethernet downlink ports are configured with the Trunk role.
- When a VLAN has DHCP snooping and Option 82 relay enabled, an Ethernet interface can be directly added to the VLAN membership, but using a tag action to associate an Ethernet interface to such a VLAN is not supported.
 - BPDU Guard and RSTP cannot be enabled on an interface, simultaneously.

- For Policy maps:
 - For Trunk interfaces, any policy map assignment is allowed.
 - For Edge and Access interfaces, if a policy map contains a two-tag classification, the edge or access link must have a tag-action that adds the outer tag being matched by the class rule.
- For an RSTP network, Calix recommends setting the following parameters as shown:
 - The VLAN IGMP Mode = snoop-suppress or proxy
 - The IGMP profile Router Learning Mode = static-dynamic
 - The IGMP profile Router Solicit On Topology Change = Y (enabled)
 - The interface NOT be designated as a static router port through the VLAN membership.

The Interface Quality Audit (IQA) function periodically checks the number of File Check Sequence (FCS) errors received as a percentage of total frames received on an interface. An interface that exceeds the provisioned thresholds can be set to generate an alarm, switch traffic to an alternate path, or force the interface to an OOS state where operator intervention is required to bring the interface to an operational state by manually disabling the interface, and then re-enabling the interface.

Ethernet Interface Roles and Applications

In general, subscriber services are deployed from the Access Domain, the logical portion of the E-Series platform with subscriber-facing ports. The Access Domain is designed to be managed through the logical Network Domain that securely connects to a service provider's back-office systems for the following:

- Enable flow-through provisioning
- Automate provisioning through the Calix Management System (CMS)
- Direct provisioning (CLI or web GUI) through the local management (MGT) ports


In the Networking Domain, services are managed in the aggregate, with the purpose of delivering high-speed traffic to and from the service network. In the Access Domain services are managed as individual Layer 2 flows. At the edge of the network, individual services are classified and rate limited, and may receive additional Layer 3 service processing (DHCP Relay Agent information insertion, subscriber service isolation, ARP proxy, etc.). Depending on the type of equipment or services that the E-Series is aggregating, this "service intelligence" may be located in the ONT, at a DSL access node, or at the Ethernet interface.

Interface roles in the domains

Ethernet interface roles are used within the E-Series to streamline configuration and standardize functions by applications that correspond to the operation in the Access and Network Domains.

- **Network Domain interface roles:** Associates the interface with the logical Networking Domain of the system and enables features consistent with an Ethernet aggregation switch. Configuration changes are infrequent and managed by a network engineer.
 - **Trunk:** Used primarily for connections to other service provider owned networking equipment.
 - **Edge:** Used primarily for connections at administrative boundaries where service VLANs are aligned / changed.
- **Access Domain interface roles:** Associates the interface with the logical Access Domain of the system and enables features consistent with Calix subscriber Access interfaces (e.g. VDSL2 ports, ONT Ethernet ports). Configuration changes are frequent and managed with flow-through provisioning (CMS or OSS).
 - **Access:** Used primarily for point-to-point connections to subscribers or other devices serving subscribers.

The following diagram shows how interface roles are used within the E7 to connect to subscribers (ONTs), Ethernet-based DSL access nodes like the E3, or other networked equipment like the E7, and it shows connections to the IP services network. Interface roles of Trunk, Edge, and Access are applied to each interface to optimize functionality, enforce network security, and scale E7 system resources.



E7 Ethernet Interface Roles and Applications

Interface role planning guidelines

	Trunk	Edge	Access
E7-2	X	X	X
E7-2 MCC	X	X	X
E7-2 MCE			X
E7-20 SCP	X	X	
E7-20 GE-24			X
ONT Ethernet Port			X

	Trunk	Edge	Access
Tag Actions		X	X
Native VLAN	X	X	
Networking Protocol (RSTP, ERPS, LAG)	X	X	

Access Interface Features

In many cases, the device providing the physical service provider edge is also the service intelligence edge. These devices have full subscriber service awareness and enforce all bandwidth, multicast, subscriber isolation, and denial of service (DOS) components of the service offering. Examples of this architecture include the following:

- Calix 700 family of GPON and Active Ethernet (AE) ONTs
- Calix E3 and E5 DSL access nodes

The E7 supports the aggregation of intelligent access nodes like the 700 ONTs and E3 VDSL2 access nodes.

In some applications, the device providing the physical service provider edge is not the service intelligence edge. Examples of this application include the following:

- Simple “media converter” ONUs that act as the WAN-to-LAN interface at the subscriber premises, where all traffic coming in the WAN port is switched to one LAN port without any significant higher-level service intelligence.
- Business service Ethernet network interface devices (“E-NID”) can provide advanced performance monitoring for the service provider, but may not provide rate limiting or service intelligence.

To facilitate the aggregation of devices that are not designed to provide a secure subscriber access interface, the E7 platform will evolve over several releases to support a full suite of service intelligence features equivalent to the GPON ONT Ethernet interfaces and VDSL2 ports. In addition to standardizing access features across all products and technologies, support for native access interfaces allows the E7 to proxy the intelligent edge for media converter ONTs and other limited functionality Ethernet access node devices.

The following table identifies the features and attributes of the E7 platform Access Domain, a space that includes ONT Ethernet ports, VDSL2 ports (R2.1 and later), and line card Ethernet interfaces. Temporary differences between the Calix ONTs and the line card Access interfaces are highlighted.

- VLAN and MAC switching
 - TR-101 1:1 per-subscriber and N:1 per-service VLAN models
 - IEEE 802.11 VLAN tags; IEEE 802.1AD Q-in-Q VLAN tags
 - Service provider VLAN tag range of 1-4093
 - Customer VLAN tag range of 1-4094

- No restrictions on VLAN ID for any service
- 9000 byte MTU on the E7; MTU size may be limited by the ONT or access node
- Tag Actions
 - Match criteria: tagged, untagged, priority tagged
 - Add-tag, change-tag, add-Add, add-change
 - Native VLAN is not supported (all untagged traffic must have an explicit tag action)
- Bandwidth profiles
 - Defined with committed information rate (CIR) and peak information rate (PIR)
 - Line cards currently support only a max-rate bandwidth profile
- Traffic engineering and QOS
 - Bi-directional, per-service rate limiting
 - Line cards currently rate limit service on ingress, with CoS rate limiting on egress. Class and Policy Maps are used on ingress to apply rate limits.
- Networking protocol support
 - Point-to-point connections
 - RSTP, LACP, and ERPS are not supported on Access interfaces
- Subscriber isolation and denial of service (DOS) protection
 - DHCP snooping and relay agent (option 82) on subscriber services
 - IP source verification
 - Broadcast and multicast traffic rate limiters
 - DHCP servers and multicast sources blocked
- Multicast services
 - IGMPv2 proxy

Access Subscriber Isolation

The Access Interfaces do not currently support per-VLAN MAC Forced-Forwarding (MAC FF) or any proprietary method to force subscriber traffic up to the Layer 3 service network. Therefore, all traffic within the system on Access interface ports will be learned and forwarded to other interfaces when operating in an N:1 VLAN-per-service delivery model.

To provide subscriber isolation, operators must do one or both of the following:

- MAC Forced-Forwarding must be used on all ONT's and subtended access nodes to force subscriber traffic up to the access router. Enabling MAC FF at the subscriber edge is a best practice and as a general practice, should not be left to aggregation systems.
- Per-subscriber 1:1 VLAN model must be used so that each subscriber is in their own broadcast domain.

System Scalability

The Access interface role in part helps to scale system resources for the E7-20 and E7-2 modular chassis, which may ultimately support 1000's of subscribers per system. In the Access Domain, each component of the system – ONTs, GE line cards, and VDSL2 line cards (R2.1 and later) – are isolated from each other with respect to the delivery of secure access services. The scalability benefits of Access Domain component isolation include the following:

- DHCP snooping for each subscriber occurs on the line card rather than on a centralized device, distributing the real-time DHCP control processes within the system.
- IGMP snooping and multicast replication occurs on the line card, as well as on the E7-20 SCP and E7-2 MCC, distributing the real-time IGMP control processes within the system.
- Networking protocols are not enabled on Access interfaces, allowing each card to manage individual services without the high overhead of coordination with other cards in the system.
- Fault isolation is improved and an individual card failure does not impact access services on other cards in the system.

Ethernet Ports

Ethernet ports and the associated Ethernet interfaces always exist and can only be modified. LAG interfaces and their association with Ethernet ports can be created, deleted, and modified. See *Ethernet Interfaces* (on page [312](#)).

The following physical characteristics of the underlying ports are interdependent as there are many interactions and limitations in their use:

- Port speed
- Duplex setting
- Flow control
- Interface type

ETH-Port names:

- g(port number) = Gigabit Ethernet Ports (GE)
- x(port number) = 10Gigabit Ethernet Ports (10GE)

Port speed

The Ethernet Port Speed has the following parameters available:

- Auto
- Fixed speed settings
- Module-Rate

Provisioning the port for a fixed speed setting (10G, 2.5G, 1G, 100M, 10M) means that speed will be forced to the value specified. The meaning of Auto and Module-Rate is described below.

Auto vs. Auto-negotiation

Auto-negotiation allows the link partners to arrive at an agreement about the speed, duplex, and flow control parameters they both support. On an auto-negotiated link there is a guarantee that the duplex and flow control options selected by the two link partners are compatible with each other.

Auto-negotiation is based on the link partners exchanging code pages that encode the capabilities they wish to advertise. The end result of the negotiation is the largest common subset of the abilities advertised by the two sides. For example, if one side supports both full and half duplex, and the other side only supports half duplex, then the outcome will be a half duplex link they can both support.

- Auto-negotiation is supported over twisted pair 1000BASE-T links and fiber 1000BASE-X links.
- Auto-negotiation is not supported over 2.5G or 10G links and some types of 1000BASE-X.

The meaning of the Auto port speed option depends on whether the link supports auto-negotiation:

- If the link supports auto-negotiation, the Auto setting means that the links will auto-negotiate the speed while advertising the duplex and flow control parameters specified. The end result of the negotiation may differ from the provisioned value depending on the link partners' advertised abilities.
- If the link does not support auto-negotiation, the Auto setting means that the highest data rate supported by the pluggable is used, and the port speed is forced to that data rate. For example, on a card that supports 2.5G, an SFP port with a 2.5G SFP provisioned for Ethernet Speed = Auto will link up at 2.5G. The same port with a 1G SFP will link up at 1G.

The Auto setting is supported on all pluggable ports.

Forced vs. Auto-negotiated Links

Speed is forced on a port if one of the following conditions exists:

- It is provisioned for a fixed speed (10G, 2.5G, 1G, 100M, 10M)
- It is provisioned for Ethernet Speed = Module-Rate
- It is provisioned for Ethernet Speed = Auto on a link that does not support auto-negotiation

Note: On forced links, there is NO guarantee that the attributes selected by the link partners are compatible.

You must ensure that the options selected by both link partners are in fact compatible. There is no way to detect this condition when linking up. For example, if the link partner only supports half-duplex, and one side selects full-duplex and the other side selects half-duplex, the link will be up, but it will not pass data reliably. The end of the link must be manually provisioned for half-duplex.

The solution for connecting to devices that do not support auto-negotiation is to provision the port manually for the speed, duplex, and flow control options that are compatible with the options supported by the other side.

Forced Speed Settings Supported	10M	100M	1G	2.5G	10G
Copper modules in SFP/CSFP ports	X	X	X		
Direct-Attach cables in SFP/CSFP ports			X		
Direct-Attach cables in SFP+ port			X		X
Copper modules in SFP+ ports			X		
1G fiber modules in SFP or SFP+ ports			X		
2.5G* fiber modules in SFP ports				X	
10G fiber modules in SFP+ ports					X

*2.5G modules are supported in SFP ports of 10GE-4 and GPON-4, only.

Module-Rate

Module-Rate is an option for provisioning port speed. It is primarily relevant for SFP+ ports that support both 10G and 1G modules. If the port is set to Module-Rate, the bit rate of the installed module is forced as the port speed. No auto-negotiation takes place with this setting. Module rate is not supported for XFP ports.

Note: See the configuration guidelines above for the forced speed setting supported for various ports.

Duplex mode for the port

The Ethernet port has the following duplex settings:

- **half-duplex** uses Carrier Sense Multiple Access (CSMA) to detect collisions and recover from them.
- **full-duplex** transmits and receives at the same time.
- **auto** operation depends on the type of link and the provisioned speed parameter:
 - If the link is auto-negotiated, the duplex attribute is negotiated with the other side.
 - If the link is forced, the setting defaults to full-duplex.

Flow control

Flow control is a method to apply back pressure to a transmitter that is outrunning the receiver's capacity to process incoming data. The flow control settings are as follows:

- **tx-pause** sends pause packets to the partner link, when needed.
- **rx-pause** honors the partner link's pause packets and stops transmitting, when asked.

- **tx-rx** sends pause packets and honors the partner link's pause packets.
- **none** does not send pause packets and does not honor the partner link's pause packets.
- **auto** setting:
 - If the link is auto-negotiated, the pause attribute is negotiated with the partner link.
 - If the link speed is forced to a set value, tx-rx is the default.

Default port settings

The default port speed for Ethernet ports are as follows:

1G:	Auto
10G SFP+:	Module-Rate
10G XFP:	10-Gig

The Duplex and FlowControl attributes default to “auto” for all front-panel port types.

E7-2 Backplane Ports

When a front-facing port is reassigned as a back-facing port, its port settings are no longer controlled by provisioning. When the port is reassigned to front-facing, its attributes revert to the default (all auto) setting.

Interaction with ERPS

Ports may be assigned as either primary or secondary ERPS ports. When this assignment, the system insists that the port be set to duplex=full, pause=none. Otherwise, the ERPS command fails. While the port belongs to an ERPS ring, the system does not allow changes to the duplex and pause attributes. Once the port is removed from the ERPS ring, the attributes can be reprovisioned to any value.

Configuration guidelines

- An Ethernet port is always a member of exactly one interface, even when it is being used in a standalone manner.
- An Ethernet port can be either assigned to the Ethernet interface associated with the port, or assigned to an existing LAG interface.
- The interface provisioning (for example, VLAN membership) applies to the port when the interface is assigned to the port.
- When a port interface is added to an ERPS ring, the port attribute of Duplex defaults to **full** and the Flow Control setting defaults to **none**. If the port interface is removed from the ERPS domain, these attributes can again be modified.
- Destination lookup failure (DLF) based rate limiting should never be used on aggregation network elements as this naturally happens when the access network goes through a topology change. The use of this feature at the access node is also not advised as the rate limiter indiscriminately applies to all services including business services.

DLF max rate is the ingress maximum rate for unknown unicast or destination lookup failure (DLF) traffic (packets/seconds). Use "k" and "m" to multiply rate. DLF applies to unicast packets where the bridges lookup the destination MAC address in their learning tables and cannot find it (a lookup miss) thus floods the packet to the broadcast domain until the packet hits a bridge that knows (learned) the destination MAC address or the packet is received by the destination device.

- The E7-2 supports 2.5 Gbps pluggable module interfaces in the SFP ports of the 10GE-4 and GPON-4 cards.
- The 2.5GE interfaces support equivalent functions and networking protocols as the GE and 10GE interfaces.
- Auto-negotiation is supported over twisted pair 1000BASE-T links and some fiber 1000BASE-X links. It is not supported over 2.5G or 10G links.
- When connecting to devices that do not support auto-negotiation, provision the E7 port manually for the speed, duplex, and flow control options that are compatible with the options supported by the other side.

VLAN Support Overview

The E-Series provide standards-based IEEE 802.1Q VLAN tagging, and Q-in-Q VLAN stacking support based on IEEE 802.1ad Bridged VLANs - 802.1Q Amendment 4, Provider Bridges. VLAN tagging was developed as a means to allow multiple networks to transparently traverse the same physical network. This is sometimes called Ethernet trunking or VLAN trunking. In addition, 802.1Q can be used to create transparent LANs, virtual LANs, and to separate traffic by VLANs, ensuring end user traffic segregation and network security. When VLANs are provisioned as a unique customer or port identifier, VLAN C-tags (customer tag) are utilized to create a VLAN per customer / port association. See *VLAN-Per-Port Provisioning Model* (on page [340](#)) for details.

The VLAN tagged frames are identified as having a tag by utilizing the Ethertype field, which is part of a standard Ethernet frame and is part of the Ethernet packet overhead. The Ethertype field is set by default for VLAN tagging (0x8100). This value is a software provisionable value in the E-Series since some switches and routers require specific values that don't follow the standard recommendation for Ethertype designations.

High-performance Ethernet switch: Each line card includes a high-performance Ethernet switch that powers network transport, aggregation, and service delivery. All services and packet flows within the system are managed with native IEEE 802.1Q VLAN tagging and IEEE 802.1ad VLAN stacking (Q-in-Q) support.

VLAN Support

Native VLAN and untagged traffic

By default, every Ethernet port with an associated interface role of Trunk or Edge is a member of VLAN 1, the Native VLAN. The Native VLAN is available to pass any untagged traffic. You can provision an Ethernet port interface to use a different existing VLAN as the Native VLAN.

Ethernet ports with an associated interface role of Access do not support a Native VLAN, however, tag actions can be used to match untagged traffic to a VLAN ID.

Note: If you are configuring the E-Series for in-band management, create a VLAN dedicated to management traffic. Do not put management traffic on the Native VLAN (default is VLAN 1).

Note: For Ethernet services, VLAN tags provisioned on an Edge or Access link interface are only passed.

VLAN tagging and provisioning models

The E-Series provides standards-based VLAN tagging, and Q-in-Q VLAN stacking support. VLAN tagging was developed as a means to allow multiple networks to transparently traverse the same physical network. When a service is provisioned at an xDSL or ONT subscriber port, a match list and service-tag action specify the classification and marking of packets from the subscriber port into the service VLAN.

VLAN-per-service provisioning model (N:1): When VLANs are provisioned to separate traffic onto VLANs based on the type of service carried in the traffic. For example, IPTV traffic is carried on a separate VLAN from data and voice traffic. A service carried on an N:1 VLAN applies to multiple subscriber ports, where a single match list and tag action can indicate the service. To isolate subscriber traffic, you must enable MAC Forced Forwarding (MACFF).

VLAN-per-port provisioning model (1:1): When VLANs are provisioned as a unique customer or port identifier, VLAN C-tags (customer tag) are utilized to create a VLAN per customer / port association. A service carried on a 1:1 VLAN is the same for each subscriber except the customer tag is unique per subscriber, you can define the match list and tag action pair such that multiple subscriber ports can reference it. This is accomplished by having a special value for the output tag in the tag action, indicating that the value of the output tag is subscriber specific. The customer-specific tag is contained in the Service object.

VLAN stacking or Q-in-Q provisioning model: The ability to add multiple VLAN tags enables the following functionality:

- Expands the addressable VLAN space from 4094 VLANs to over 16 million VLANs

- Allows logical separation and trunking of VLANs through a network by using a VLAN tag to group a larger range of VLAN tags together

The most common way to use VLAN stacking is by inserting two tags on the traffic. These tags are typically referred to as the inner tag or C-tag and the outer tag or S-tag. As stated previously, the C-tag, also known as the customer tag is used to uniquely identify a customer, typically is used on a per port basis. The S-tag, also known as the service-provider tag is used to logically group C-tags together.

Metro Ethernet Forum (MEF)/Transport LAN (TLAN) business service models:

MEF/TLAN service can transparently trunk business traffic across a network to other locations, typically a remote office or secondary business location. The traffic received from the business may be untagged, single-tagged, or double-tagged. The E-Series adds an outer tag to all frames to create a private switched LAN with two or more end points.

VLAN Traffic flow in VDSL2 and GPON Cards

Since xDSL and GPON Ethernet ports have no interface association, unlike the Ethernet ports, rather than adding a port interface to a VLAN membership to enable traffic flow, a service-tag action must be created that specifies the VLAN. This service-tag action is then referenced when the service is provisioned on the xDSL or ONT port. The VLAN must already be created on the system for an xDSL or GPON port to pass traffic carried on a VLAN.

- To view the VLANs associated with specific GPON or xDSL port, click the port of interest on the Navigation Tree, and then click the **VLANS** tab or the **Associated Interface > VLANS** tabs for xDSL.
- To view the VLANs associated with system ONTs, click **ONTS** on the Navigation Tree, and then click the **VLANS** tab. Or, for a specific ONT, select the ONT ID first, and then click the **VLANS** tab.
- To view the services associated with a specific VLAN, click **VLANS** on the Navigation Tree, click the particular VLAN from the list that appears in the Work Area, and then click the **Service Associations** tab.

VLAN with IGMP Snooping and multicast IP addresses

The E-Series supports industry standard IETF RFC 4541 IGMP snooping of multicast video leave and join requests sent between the set-top box and the video distribution network.

IGMP snooping enables the E-Series to determine which ports should be a member of a multicast group. It can provide a local response if the channel already exists on the system or it will forward those requests upstream to a multicast router or another IGMP snooping Ethernet switch. When multicast channels are received from the network, the E-Series forwards the channels to all ports that are currently members of the requested IGMP membership group. The ability to perform snooping can be enabled on a per VLAN basis. If multicast traffic is received on a VLAN where IGMP snooping is not enabled (flood), the traffic is handled as broadcast traffic and sent to all ports on the VLAN.

- Passive IGMP snooping (snoop-suppress), augmented with report suppression reduces the number of general query and group specific reports sent to the multicast router (querier).

Note: The suppress-snoop parameter does not apply to the E7-20, as it only supports proxy for IGMP.

- Active IGMP Snooping (IGMP Proxy), becomes capable of acting as a general querier and thus takes an active role in maintaining the IGMP multicast network. When the IGMP mode is set to proxy, an IGMP profile must be referenced from the VLAN. In addition to general querier timing parameters, the following IGMP attributes per multicast VLAN are supported:
 - Robustness parameter (number of times E-Series sends out a general query)
 - Last member query interval
 - Immediate leave option
 - Static multicast router interface
 - IGMPv3
 - Router learn mode

IGMP Proxy provides additional scalability to the IGMP network, reducing the IGMP signaling load on the multicast router.

Additional multicast addressing considerations

In accordance with RFC 4541 (section 2.1.2), the E-Series automatically passes multicast IP addresses in the 224.0.0.x address range (defined as link-local) through the system on VLANs with IGMP Snooping mode set to snoop-suppress or proxy (snooping enabled). For example, the following 224.0.0.x multicast addresses are reserved for specific routing protocols, and forwarded automatically on the E-Series without a corresponding join request.

HSRP HELLO	224.0.0.2	EIGRP	224.0.0.10
DVMRP	224.0.0.4	PIM	224.0.0.13
OSPF ALL RTR	224.0.0.5	VRRP	224.0.0.18
OSPF DES RTR	224.0.0.6	HSRP	224.0.0.102
RIP	224.0.0.9	MLS ALL	224.0.0.106
		SNOOPERS	

If reserved 224.0.0.x multicast addresses are assigned to video channels in your lineup, the E-Series will flood these channels to all ports on IGMP Snoop enabled VLANs, which could potentially consume available bandwidth on subscriber links and result in tiling. Therefore, Calix recommends NOT using 224.0.0.x (or any of the 32 multicast IP addresses that are not unique at the L2/MAC address level) for video channel assignments. See *Calix Quick Tip E7 QT-12-003* for more information.

For VLANs with IGMP Snooping set to flood (no snooping), the E-Series passes all multicast traffic transparently.

Service Security

The E-Series supports the following service security features that are defined when the VLAN is created. If using a N:1 VLAN model for services, Calix recommends enabling MAC Forced Forwarding (MAC FF) and IP Source Verify (IPSV) when creating the VLAN, where appropriate. The VLAN Per Service model is susceptible to MAC spoofing unless MACFF and IPSV are enabled.

- **IP Source Verification** ensures that only data from IP addresses learned by DHCP snooping or static provisioning are allowed to ingress ONT Ethernet ports or xDSL ports.
- **MAC Forced-Forwarding** ensures that traffic from one subscriber interface cannot be sent directly to another, reducing the chance that malicious traffic can be transmitted between ports.

Reserved VLANs

To ensure a wide range of access deployments, the E-Series supports an extremely flexible VLAN support model. For future development support, the E-Series reserves four VLAN values. The E-Series requires these VLANs always be identified. The default values for these VLANs are 1002, 1003, 1004, and 1005. It is possible to change these four values to be some other set of four consecutive VLANs, if required using the basic system settings.

AE ONT Discovery

When a Calix AE ONT is attached to an E7 untrusted Ethernet port, the E7 system can send the Discovery event as an SNMP notification (trap or inform message) to CMS regarding the newly discovered AE ONT, as configured by the “snmp trap-dest” provisioning.

SYSTEM MANAGEMENT RSTP ERPS POLICIES MULTICAST DHCP ETHERNET OAM PROFILES PPPoE TEMPLATE									
INBAND INTERFACE SNMP CRAFT USERS RADIUS ACTIVE SESSIONS									
CONFIGURATION COMMUNITY USERS TRAP DESTINATION									
CREATE DELETE ACTION REFRESH APPLY									
ID	SNMP VERSION	ADMIN STATE	PORT	V3 USER	V2C COMMUNITY	HOST NAME	INFORM		
1	v2c	enabled	162		public	10.208.53.34	N		

Note: CMS automatically adds its IP address as a trap destination when adding an E7 to CMS.

The VLAN configured at the untrusted Ethernet port must have both the AE ONT Discovery function and the DHCP Snooping function enabled for the E7 to initiate the SNMP notification of an AE ONT discovery.

- When you enable DHCP Snooping and AE ONT Discovery on a VLAN, the following sequence occurs:
 - a. DHCP messages received by the E7 DHCP Snooping agent at “untrusted” Ethernet interfaces are examined to identify messages that are originated by the Calix AE ONT (DHCP Message with Option 60: Vendor class identifier value = “Calix ONT”).

- b. A lease is granted by the DHCP Server to the AE ONT, triggering the AE ONT Discovery event where the ONT is dynamically linked with specific GE port it reported on, and will show up on the navigation tree.

The event report includes the following fields:

- ♦ MAC address of the AE ONT
- ♦ VLAN
- ♦ Interface where the AE ONT is discovered

- c. The lease is first added to the binding table maintained by the DHCP Agent.

Note: If the lease is already present in the DHCP lease table (assuming DHCP Snooping is enabled) and AE ONT Discovery is enabled after a DHCP exchange, no event is reported. Therefore, it is important that you enable both AE ONT Discovery and DHCP Snooping when first configuring the VLAN attributes of the AE ONT management VLAN.

Note: The AE ONT discovery feature only supports Calix AE ONT release 2.0 or higher with dynamically assigned IP addresses and with DHCP Option60 (Vendor class identifier value = "Calix ONT").

- When you do NOT enable the AE ONT Discovery on a VLAN, yet you want configured AE ONTs linked frequently, you will need a separate, more frequent MAC/IP Host only scheduled inventory task run often (daily) to collect the MAC information for linking.

VLAN Association to voice services

The IP Host object references a service tag action that indicates the voice service VLAN by one of the following methods:

- An IP host can directly specify an outer tag, when the referenced service tag action states "Specified in Service" for the Outer Tag parameter.

This allows you to apply a different VLAN tag for the same service type (SIP, TDM gateway, H.248 gateway, MGCP, or T1 PWE3) on each ONT.

- An IP host can reference a service tag action that specifies the service VLAN ID for the provisioned voice services.

This allows you to apply the same VLAN tagging for all ONTs that use the referenced service tag action object.

Configuration guidelines

- VLAN-Per-Service Model (N:1):
 - The VLAN Per Service model requires fewer VLANs than the VLAN Per Subscriber model.
 - To isolate subscriber traffic, you must enable MAC Forced Forwarding (MACFF). The VLAN Per Service model is susceptible to MAC spoofing unless MACFF and IPSV are enabled.

- This model provides the most bandwidth-efficient way of distributing multicast content to all access nodes for video service.
- Isolation of multicast traffic provides the flexibility to size a video VLAN without having to address unnecessary multicast replication on the access network.
- The P-bit value defined for a service VLAN must be consistent with the Class of Service (CoS) type.
- Deploy VoIP service using the VLAN-Per-Service model; it is a captive VLAN that is not open to subscribers.
- Limit broadcast traffic and Address Resolution Protocol (ARP) requests by configuring smaller subnets, each on a unique VLAN.
- Do not create a large VLAN space with a large associated subnet space.
- Set the broadcast domain size so that devices in the domain are not overwhelmed with packets during peak traffic events (for example, following a power outage all network devices may request IP addresses simultaneously).
- VLAN-Per-Subscriber Model (1:1):
 - Each subscriber is assigned a unique VLAN, and VLANs belonging to the same access node are grouped under a single VLAN.
 - Because MAC addresses are learned on a per VLAN basis, the VLAN Per Subscriber model provides MAC address isolation, protecting the network from MAC collisions, with the following exceptions:
 - A broadcast domain is limited to a single access node and the router, making it easier to manage the broadcast load.
- Even if no interface is currently using VLAN 1 as the Native VLAN, it is still off limits for user provisioning, including use as the Management VLAN or ERPS control VLAN.
- The E-Series reserves four VLAN values for system operation. The default values for these VLANs are 1002, 1003, 1004 and 1005. You can change these values to another set of four consecutive VLANs, if required.
- For an Ethernet port interface to pass traffic carried on a VLAN, the interface must be added to the VLAN membership, or specified as the target in a tag action.
- You can configure only one of the following attributes on a given VLAN on a given Ethernet port interface:
 - VLAN membership
 - Tag-action (edge and access interfaces only)
 - Native VLAN (edge and trunk interfaces only)
- For an xDSL or GPON port to pass traffic carried on a VLAN, a service tag action must be created specifying the VLAN.
- For Home Gateway services, the first service VLAN is required to be assigned to High Speed Internet services. It has additional attributes which define the management mode, WAN protocol, and required configuration or credentials.

- The T-Series supports four independent WAN VLANs (High-Speed Internet, IPTV, VoIP, TR-069). The RG configuration file "NAT enabled" attribute controls which of these VLANs are routable from the LAN. By default, only the HSI WAN is routable. Static/Default routes assigned to other WAN's require NAT to be enabled before you can access them from the LAN.
- The E7 Management VLAN needs to have a VLAN ID that is unique and different from any service VLAN, including the RG management profile VLAN used for TR-069.
- DHCP Snoop/Proxy, IP Source Verify, MAC Forced-Forwarding, PPPoE, and Option 82/LDRA are not supported on the Management VLAN.
- For GPON or xDSL VLANs that need to carry IPv6 traffic, enable the VLAN TLAN parameter for transparent passthrough of IPv6 traffic.
- Option 82/LDRA control is supported at the VLAN level. For complete details, see Configuring DHCP Relay Option 82 and LDRA.
- An operator may manually delete a lease entry from the table. However, if an IP Source Verification is enabled on the VLAN, deleting the DHCP lease entry results in the subscriber traffic being dropped until the subscriber's IP host requests, and is granted, a new IP address using DHCP.
- When DHCP snoop/proxy is enabled on a subscriber VLAN, the E-Series drops all DHCP server communication originating from subscriber Ethernet interfaces. Calix recommends DHCP snoop/proxy be enabled for all residential subscriber services, and that redundant DHCP servers are included in the equipment configuration.
- In the VDSL2 and GPON subsystems, DHCP leases cannot be learned without also applying a limit to the number of learned leases on the port. Each Ethernet port has an associated Port Security Profile that can limit DHCP leases in a range of 1-16. A security profile used by an xDSL port must have a DHCP lease limit value of 10 or less.
- Only perform redundant storage of 32,000 DHCP leases.
- When DHCP Snoop/Proxy is enabled, DHCP requests for VLANs do not appear in a port mirror session.
- The craft management ports cannot use IP addresses from the same subnet where DHCP Snoop/Proxy is enabled.
- DHCP Snoop/Proxy is not supported on management VLANs.
- DHCP Snoop must be enabled on the VLAN to support the AE Discovery Event operation.
- Simultaneous operation of DHCP Snoop/Proxy and PPPoE are not supported. When a PPPoE profile is selected, the DHCP features are disabled. See Configuring PPPoE Operation for a Data Service for more information.
- For E-Series VDSL2 and GPON, IP and MAC addresses may be dynamically learned using either DHCP Snoop/Proxy or manually provisioned with static IP/MAC addresses. Services with static subnets (without MAC address specification) may also be provisioned for IP Source Verification, but are bound to the port only by IP address.
 - DHCP and Static IP host: Binds IP and MAC address to a Port

- Static IP Subnet: Checks individual host IP addresses to the subnet and binds the subnet to a port
 - The E-Series GPON systems have the following limits on the number of IP stations (active DHCP leases, static IP hosts, and dynamically learned hosts within a static subnet):
 - ♦ The ONT supports up to 256 “IP stations” per ONT Ethernet port.
 - ♦ The ONT supports up to 1024 total “IP stations” per ONT.
 - The following capacities apply to Static IP Addresses/Subnets within the VDSL2 subsystem.
 - ♦ 1 Static IP subnet can be provisioned per xDSL Ethernet service.
 - ♦ 4 Static IP hosts can be provisioned per xDSL Ethernet service.
 - ♦ 8 Static IP hosts/subnets can be provisioned per xDSL port, across all services.
 - ♦ 256 Static addresses total per port (includes static addresses and sizes of static subnets)
 - ♦ 16 DHCP Leases per port (defined in the Ethernet Security Profile with a default setting of 8).
- Disabling MAC Learning causes traffic within the VLAN to be flooded to all egress interfaces with membership in the VLAN, which can degrade throughput.
- Disable MAC Learning for MEF point-to-point Ethernet Private Line (EPL) and Ethernet Virtual Private Line (EVPL) services to allow service delivery without the E7 adding any subscriber MAC addresses to the E7 forwarding table.
- Disabling MAC Learning should only be done with point-to-point services and is NOT appropriate for a multi-point Ethernet-LAN (ELAN) service or other VLAN-per-service applications.
- MACFF can only be used in conjunction with DHCP Snoop/Proxy, where you first enable DHCP Snoop/Proxy, wait a lease cycle, and then enable MACFF.
- For T-series ONT peer-to-peer voice service, PON Hairpin or MAC-FF must be enabled.
- Enable MACFF to limit broadcast traffic on the VLAN domain when using the VLAN Per Service model. It is not required for the VLAN per Port (single and double-tagged, Q-in-Q) data model or PPPoE because this VLAN model by definition limits the broadcast domain to a single access node and the router and makes it easier to manage the broadcast load. However, MACFF should still be enabled on these VLANs when the subscriber edge does not include a residential gateway.
- For VDSL2 or GPON, MAC Forced Forwarding (MAC FF) can be used with the following IP hosts:
 - IP hosts learned via DHCP Snoop/Proxy
 - IP hosts statically provisioned (IP and MAC)
 - IP hosts statically provisioned (IP address only)
 - IP subnets statically provisioned (single IP address/MAC)

- MAC Forced Forwarding supports a single source MAC address with multiple source IP addresses, when the MAC address is not specified.
- Neither MAC Forced Forwarding or IP Source Verify is supported for Residential Gateway (RG) service VLANs.
- For VDSL2 or GPON, MAC FF can be used with static IP subnets.
- The GPON system has the following limits for the PON Hairpin feature:
 - PON Hairpin capacity limitation of 2.4Gbps of traffic on a per GPON-4 card basis. This is equivalent to 50% of the total upstream PON bandwidth capacity (using 4 PON interfaces). There is no limitation on the amount of traffic that can be switched between GPON OLT ports on the same GPON-4 line card or between E7 line cards.
 - The 2.4Gbps capacity is shared by GPON-4 OLT ports, thus 2 of the PON ports can have 100% of their capacity hairpinned. If capacity is equally distributed each PON port can have 600Mbps of hairpinned traffic.
 - If the 2.4Gbps capacity is exceeded then the GPON-4 subsystem will go into congestion and will drop packets.
 - The VLAN PON Hairpin function is not compatible with DHCP Snoop/Proxy, IGMP snoop, MAC Forced-Fwd, or IP Source Verify.
 - The MVR Unicast VLAN must have TLS/PON Hairpin disabled on that VLAN for services to work correctly.
- IGMP snooping is provisioned on a VLAN basis and is recommended for VLANs carrying video services and should be avoided for VLANs in a Transparent LAN Service (TLS).
- IGMP snooping is only enabled on the outer VLAN ID as multicast traffic is typically transported throughout the network with a single VLAN ID tag.
- All nodes in an ERPS or G.8032v2 ring must have the same IGMP Snooping provisioning on the video VLAN for traffic to flow--either all with snooping (snoop-suppress, proxy) or all without snooping (flood).
- T-series and L2 IGMP-controlled multicast:
 - T-Series supports multiple VLANs, but only supports a single service provider and requires channel independence across VLANs.
 - Different VLANs on different LAN ports share query domains and are treated as a single video service by the ONT.
 - The ONT has a single proxy agent and does not consider the transport VLAN when filtering channels. The only relevant information is the layer-3 IP address.

- The IGMP mode setting affects VLANs as follows:
 - Snoop-suppress mode only uses IGMPv2.
 - Proxy mode references an IGMP profile that specifies the IGMP version setting. You can edit the system-default IGMP profile or create a custom IGMP profile. Other VLANs that use the same IGMP profile are not affected by the version setting unless they are in proxy mode.
 - Snoop-Suppress or Proxy mode can reference a P-bit value for IGMP traffic that passes through the system, allowing the traffic type to be treated differently as it passes through the network. The P-bit value set in a service-tag action, such as change-tag, can override the provisioned IGMP P-bit value.
- The default behavior for data services is to filter all multicast traffic upstream from an ONT Ethernet port or xDSL port, unless the TLAN parameter is enabled on the VLAN. For VLANs that are carrying IPv6 multicast traffic, enabling (selecting) this feature allows a passthrough of the IPv6 traffic.
- Channels not joined by a subscriber on the ONT are dropped at the ONT. Channels not joined by any ONT on the PON are dropped at the GPON OLT port.
- The GPON card supports up to 32 IGMP-enabled multicast VLANs on the system, which can all be on the PON.
 - For 1:1 provisioning model, the C-tag must be unique for each subscriber.
 - For N:1 provision model, the C-tag must be unique for each service.
 - For a double-tagged model, each of the inner tags must be unique.
 - For a single-tagged model, each of the tags must be unique.
- A VLAN can only be deleted if it is NOT referenced by a tag action, a service tag action, a provisioned service, or a voice service IP Host.
- E-Series system supports the unrestricted use of C-Tag VLAN ID values in combination with any S-Tag value, across any service end-point on an E7 system, regardless of PON port, line card, or shelf.

With the following exceptions:

- A C-Tag (or single-tagged S-Tag) may not be re-used on an individual subscriber port (UNI).
- An S-tag may not be re-used for a C-tag of another service on the same PON or VDSL2 port. For example, when a VoIP service uses S-tag 200 for all subscribers, and a Data services uses S-tag 100, the VoIP S-tag may not be re-used as a Data service C-TAG 200 (100, 200) on a given PON or VDSL2 port. The recommended practice is to ensure the S-Tag range is outside of the subscriber C-Tag range.
- An MVR VLAN may not match the outer tag specified by any service-tag action or Ethernet service.
- An MVR VLAN may not match the inner tag specified by any service-tag action or Ethernet service.

Parameters

You can provision the following parameters for VLANs:

Parameter	Description	Valid Options
ID*	VLAN ID for the VLAN you are creating.	2-4093 (Except for 1002-1005 which are reserved for E7 operation.)
Name	Name of VLAN that is descriptive of the VLAN use. For example, VoIP or Management.	String of 31 characters
IGMP Mode	IGMP mode for the VLAN. Calix recommends that VLANs used for video services be set to suppress-snoop or proxy (enable snooping). snoop-suppress (not supported on E7-20): Enables IGMP snooping (with report suppression) on the VLAN to manage the multicast group memberships of subscriber ports. The IGMP capability ensures that only multicast channels which are joined by a particular set-top box (STB) appear on the subscriber network. All nodes in an ERPS ring must have the same IGMP Snooping provisioning on the VLAN for video traffic to flow--either all enabled or all disabled. A VLAN with IGMP mode = snoop-suppress only uses IGMPv2. Within the IGMP protocol, the querier generates the following messages: --A General Query message is generated where all devices which are joined to any group respond with a Report message. --A Group Specific Query message is generated where the specific group responds with a Report message. The messages and reports can generate a lot of IGMP activity in a large network. Report suppression allows intermediate devices, between the querying router and the multicast consumer host, to suppress duplicate Report messages within the response timer window, in order to reduce the amount of IGMP traffic that must be processed higher in the network. Disabling Report Suppression allows more Report messages to be allowed up through the network. proxy: Causes the E7 to act as an IGMP v2 proxy for all STBs, sending join/leave requests to the upstream IGMP router/content provider as required. This provides a more robust IGMPv2 implementation, scaling to the maximum number of subscribers and channels. When the IGMP mode is set to proxy, an IGMP profile must be referenced from the VLAN on all E-Series that have a video VLAN with the IGMP Proxy setting. The IGMP version (v2 or v3) is specified in the associated IGMP profile. The E7-20 only supports proxy mode. flood: Disables IGMP snooping and proxy. If multicast traffic is received, IGMP traffic is forwarded through and multicast traffic is forwarded to all member ports on the VLAN. The G.8032v2 ring video only supports flooded IGMP mode. Note: Recommended for VLANs not carrying multicast traffic.	snoop-suppress proxy flood (default)
IGMP Profile	Name of the IGMP profile to associate with the VLAN. You can select the system default profile or another existing IGMP profile. Note: This profile only applies if the IGMP mode is set to proxy.	system-default (default) any existing profile

Parameter	Description	Valid Options
DHCP Mode	Allows the selection of three DHCP modes: • Snoop: Enables DHCP snooping to track all DHCP activity on the VLAN and create/update a table of DHCP leases granted. • Proxy: Enables DHCP snooping as well as a DHCPv4 Proxy agent for added security. To complete the DHCPv4 Proxy (or Layer 3 DHCP Relay) configuration, you must also select a DHCP Proxy Service profile for this same client-side VLAN and configure a server-side VLAN with DHCP Mode = none. • None: Disables DHCP Snoop or DHCPv4 Proxy on the client-side VLAN. Notes: • Throughout this guide, references to DHCP Snoop or snooping also apply to DHCP Proxy. • Calix recommends DHCP snoop/proxy be enabled for all residential subscriber services, and that redundant DHCP servers are included in the equipment configuration. • Maximum number of service VLANs with DHCP snoop/proxy enabled, per VDSL2 line card = 48, per GPON card = 256.	none (default) snoop proxy
MAC Forced-Fwd	(VDSL2 and PON only) Whether to enable MAC Forced Forwarding (MAC FF) which screens upstream packets at the ONT Ethernet port or xDSL port and only allows through those packets with a destination MAC address (DMAC) that matches the upstream access router. MACFF can only be used in conjunction with DHCP Snoop/Proxy, where you first enable DHCP Snoop/Proxy, wait a lease cycle, and then enable MACFF. Calix recommends MAC Forced Forwarding be enabled for all residential and business internet access services. MAC FF supports the following: • IP hosts learned via DHCP Snooping • IP hosts statically provisioned (IP and MAC) • IP hosts statically provisioned (IP address only) • IP subnets statically provisioned (single IP address/MAC) Note: Maximum number of service VLANs with MAC Fored-Forwarding and/or IP-Source-Verify enabled, per VDSL2 line card = 48. Note: For T-series ONT peer-to-peer voice service, PON Hairpin or MAC-FF must be enabled.	cleared (disabled) (default) selected (enabled)
IP Src Verify	(VDSL2 and PON only) Whether to bind the IP address and MAC address to the physical ONT Ethernet port or xDSL port, preventing subscribers from assigning an IP address to a device and passing traffic on it. IP Source Verification for Static IP hosts requires MAC FF be enabled. Calix recommends IP Source Verification be enabled for all residential and business internet access services. Note: Maximum number of service VLANs with MAC Fored-Forwarding and/or IP-Source-Verify enabled, per VDSL2 line card = 48.	cleared (disabled) (default) selected (enabled)

Parameter	Description	Valid Options
MAC Learning	Controls MAC address learning. This parameter control only applies to standalone E7-2 systems. - E7-20 and E7-2 in Modular Chassis mode only support MAC Learning enabled. - E7-2 system cannot be set to Modular-Chassis mode while any VLAN has MAC Learning disabled. Calix recommends leaving MAC Learning in the default enabled mode. - Disabling MAC learning causes traffic within the VLAN to be flooded to all egress interfaces with membership in the VLAN. - Disabling MAC Learning and DHCP Snooping, and enabling MAC Forced-Fwd, causes all downstream packets to be flooded. - Enabling DHCP Snoop, sets MAC Learning to disabled by default.	cleared (disabled) selected (enabled) (default)
AE Discovery Event	Whether the E7 sends an event to CMS whenever a new Calix AE ONT is discovered on an untrusted port. When enabled, an event is sent for any new lease added to the binding table entries belonging to AE ONTs. If the lease for the AE ONT already existed in the table at the time of enabling the feature, an event will not be sent. Note: DHCP Snoop must be enabled on the VLAN to support the AE Discovery Event operation.	cleared (disabled) (default) selected (enabled)
TLAN	(VDSL2 and PON only) Whether to enable PON or xDSL transparent LAN service which disables Ethernet Security Profiles, MAC Forced Forwarding, DHCP Lease Limits, or Upstream Broadcast Limits that were applied to the ONT Ethernet port or xDSL port with the associated VLAN. Note: The default behavior for data services is to filter multicast traffic upstream from an ONT Ethernet port or xDSL port, unless the TLAN parameter is enabled on the VLAN. For GPON or xDSL VLANs that need to carry IPv6 traffic, enable the VLAN TLAN parameter for transparent passthrough of IPv6 traffic.	cleared (disabled) (default) selected (enabled)
PON Hairpin	(PON only) Whether to enable PON hairpin which allows traffic to flow upstream from an ONT Ethernet port and back downstream on the same PON to another ONT Ethernet port. Hairpin of services in the PON requires additional resource allocation in the GPON subsystem and should not be enabled unless hairpin service is required. Note: The PON Hairpin function is not compatible with DHCP Snoop, IGMP snoop, MAC Forced-Fwd, or IP Source Verify. Note: Supported for TLAN and T1/E3 PWE3 services. Note: For T-series ONT peer-to-peer voice service, PON Hairpin or MAC-FF must be enabled.	cleared (disabled) (default) selected (enabled)
IGMP P-Bit	The P-bit value for IGMP traffic that passes through the system, allowing the traffic type to be treated differently as it passes through the network. Note: Only applies when the IGMP mode is set to Snoop-Suppress or Proxy. The P-bit value set in a service-tag action, such as change-tag, can override the provisioned IGMP P-bit value.	pbit-4 (default) pbit-0 through pbit-7
DHCP Proxy Service	Required for the configuration of DHCPv4 Proxy (Layer 3 DHCP Relay) on the client-side VLAN. The selected profile determines the DHCP Proxy Agent IP interface and IP address of the DHCP server(s).	valid profile name
Option 82	Enables or disables Option 82/LDRA at the VLAN level. Note: If Option 82 is disabled at the system level, enabling it at the VLAN level will not have any effect. The Option 82 enable/disable control at the VLAN level is only effective if Option 82 is enabled at the system level.	cleared (disabled) selected (enabled) (default)

Parameter	Description	Valid Options
Access Identifier Profile for Ethernet	Access Identifier Profile for Ethernet/xDSL subscribers, specifying the Circuit and Remote ID format for Option 82 content insertion. Note: A profile must be selected for both Ethernet/xDSL and GPON subscribers, even if only one type of subscriber is served by the VLAN. In this case, any profile may be used for the subscriber type that is not present, or a "null" profile (with Circuit and Remote ID = None) may be created and used for this purpose.	valid profile name
Access Identifier Profile for GPON	Access Identifier Profile for GPON subscribers, specifying the Circuit and Remote ID format for Option 82 content insertion. (See the above note.)	valid profile name
Mobility	Enables or disables the ability for client devices to move freely between different ONTs on the same PON or card. (Not supported across multiple cards.) - disabled (unselected) = mobility off Station movement is blocked in the GPON MAC when moving between ONTs on the same PON. There is a 5-minute expiry timer before a new MAC can be learned in the GPON MAC. - enabled (selected) = mobility on for a VLAN Station movement is enabled in the GPON MAC that is provisioned for the specified VLAN. A device is allowed to move from one ONT to another without a 5-minute delay. Security features (DHCP, MACFF, IPSV) must be disabled when using the mobility feature. Services must be HSI only. No multicast profiles can be applied to ONT services.	cleared (disabled)(default) selected (enabled)
PPPoE Profile	Assign a previously-created PPPoE profile. When a PPPoE profile is selected, the DHCP features are disabled. Setting a VLAN PPPoE profile to "none" passes through all PPPoE traffic, transparently. If a PPPoE profile is used with PPPoE snoop, a list of all the active sessions and statistics are available, and the PPPoE stack is enabled, which passes through PPPoE traffic transparently as long as the Clients/BRAS are operating normally (illegal packets will be dropped). See Creating a PPPoE Profile for instructions. Note: Simultaneous operation of DHCP Snooping and PPPoE are not supported.	any available PPPoE profile

*Required field


VLAN-Per-Service Provisioning Model

As shown in the following diagrams, the VLAN-per-service provisioning model separates traffic onto VLANs based on the type of service carried in the traffic. For example, IPTV video traffic is carried on a separate VLAN from data and voice traffic. This provisioning model is the easiest model to configure and use but does not provide segregation of user's traffic, which is often desired especially for data traffic to ensure security and tracking of subscriber activity. In this VLAN provisioning model, the VLANs may have an S-tag added to the traffic by the E-Series as shown in the diagram. The main reason for the S-tag would be to segregate traffic on the network into smaller groups, which provides added security benefit. For example, if a broadcast storm occurs on the network, the amount of traffic that is affected is limited because the network has been logically divided into smaller groups.

The E-Series supports the following VLAN capabilities when configured for the VLAN-per-service provisioning:

Proprietary Information: Not for use or disclosure except by written agreement with Calix.

- S-tag (and/or P-bit) addition to traffic on a per-drop basis
- S-tag (and/or P-bit) addition to traffic based on the value of the C-tag
- S-tag (and/or P-bit) addition to traffic based on the DSCP (IP Precedence) value of the traffic
- Traffic switching based on the newly added S-tag, and prioritizing and rate limiting the traffic based on P-bit
- Traffic switching based on the S-tag on the traffic received from the downstream port, and prioritizing and rate limiting the traffic based on P-bit



VLAN-per-Service Model

VLAN-Per-Port Provisioning Model


The the VLAN-per-port model is based on each port of a broadband access platform being mapped into a unique VLAN prior to Ethernet aggregation and transport. The VLAN tag is passed upstream to the Ethernet switch, router, or broadband remote access server (BRAS). This deployment method is very similar operationally to the traditional ATM model, where a layer 2 ATM PVC maps each customer (VCI) on a DSLAM's ATM uplink (VPI) for transport back to a broadband RAS for protocol conversion and subscriber management.

The VLAN-per-subscriber model provides the following features and advantages:

- Every subscriber has their own dedicated VLAN.
- Individual customer's traffic is segregated, providing security for data traffic and subscriber activity tracking.
- Used to emulate ATM provisioning models in Ethernet.
- Traffic is mapped into an upstream VLAN trunk using an outer VLAN tag (S-tag).
- Centralized service and subscriber management is required.


Double-tagged

- ▶ Subscribers have common outer tag
- ▶ Each subscriber gets a unique inner tag
- ▶ Only service VLAN needs to be created



The graphics below show the E7 system's VLAN-per-port (subscriber) provisioning models from the perspective of network, logical, and VLAN flow.

This provisioning model provides 4094 VLANs per the IEEE 802.1Q standard. The E-Series allows the combination of the VLAN per port model and the ability to add S-tags to incoming traffic to simplify transport and switching. For additional information, see *VLAN Stacking (Q-in-Q)* (on page [342](#)).



VLAN-per-Subscriber (Port) Model

As shown in the above diagram, the E-Series supports the VLAN-per-subscriber (port) provisioning model and allows the full range of VLANs to arrive on a downlink port and be forwarded to the upstream ports. Unless VLAN stacking is utilized, there are only 4094 usable VLANs per the IEEE 802.1Q standard. Therefore, it can easily be seen that the aggregation capabilities in this scenario are limited, making VLAN stacking essential.

The E-Series supports the following VLAN capabilities when configured for VLAN-per-port provisioning:


- VLAN switching based on the outer VLAN tag received from the downstream device, with no modifications made to the VLAN tags.
- Outer tag (S-tag) addition to the VLAN-per-port traffic, then switch based on the newly added S-tag.
- Outer tag (S-tag) addition on all received VLANs on an E-Series interface, regardless of originating subscriber port, then switch based on the newly added outer S-tag.
- VLAN tag translation that changes VLAN tags to the same value or a new value based on pre-identified criteria. The E-Series limits the number of VLANs that can be translated to an aggregate limit of 768.
- A VLAN tag addition using the 802.1p bit value to classify the incoming frame.
- Set the value of the VLAN S-tag and/or P-bit based on the DSCP (IP Precedence) value of the IP header of the traffic received from the subscriber. The ability to support the provisioning of the S-tag and/or P-bit value enables support for devices that cannot implement VLANs or P-bit priority Ethernet standards independently.
- The value setting of the Ethertype on a per port basis. The default Ethertype value is set to 0x8100 for the VLAN S-tag. The default value can be changed to be 0x8100, 0x9100 or 0x88a8. The ability to support the provisioning of the Ethertype value enables support for devices that implemented Q-in-Q prior to the standard being fully specified.

VLAN Stacking / Q-in-Q

As shown in the following diagram, the E-Series allows the combination of the VLAN-per-subscriber (port) model and the ability to add S-tags to incoming traffic to simplify VLAN transport / switching. Some 3rd-party devices may not have the ability to add S-tags before an E-Series handoff. In these cases, the E-Series can add the tag. The tag added to this traffic can be specified to be different, based on what port the E-Series receives the traffic, or be common for all traffic sent upstream.

The diagrams below show the typical use of the C-tag and S-tag stacking in a residential 1:1 per port model. The S-tag represents the DSLAM in the network and the C-tag represents DSL port.

VLAN Q-in-Q single tag




VLAN Q-in-Q Single tag Model

VLAN Q-in-Q double tag

External VLAN double tagging allows the E-Series to switch on VLAN (transport only). The diagrams below show the E-Series addition of an outer S-tag to previously tagged frames. This scenario is typical of aggregation of older IP DSLAMs that do not have the ability to add an outer S-tag.

Another way VLAN stacking can be used is when the aggregated device has the ability to send traffic with two VLAN tags. In this example, the E-Series only needs to switch the traffic to the appropriate port without any modifications to the VLAN tags.



VLAN Q-in-Q Double tag Model

Hybrid VLAN Provisioning Model

The E-Series allows VLAN-per-port traffic to be on the same interface as VLAN-per-service traffic. This configuration is typically used by service providers that want to maintain data traffic security and customer isolation, yet utilize the VLAN-per-service model for IPTV and VoIP. Typically, IPTV networks place all multicast channels on a single VLAN, enabling a simplified provisioning model and interoperability with multiple middleware vendors.

The E-Series supports the VLAN provisioning model for the following scenarios:

- When there are different S-tag stacking requirements for a single drop based on the type of tag received
- When the VLAN-per-service aggregation requires different S-tags to be added for different C-tags

VLAN Memberships for Interfaces

This topic describes the association of Ethernet port interfaces, LAG interfaces, and ERPS domains to a VLAN, thereby creating a VLAN members. A VLAN member enables traffic on a specific VLAN to be forwarded to or accepted on an Ethernet port interface or ERPS domain. For any VLAN-tagged traffic to flow through an Ethernet interface, either the interface must be a member of the tagging VLAN, or the interface must have a corresponding VLAN tag-action associated with it. When you create a tag-action, the interface automatically becomes a member of the VLAN specified in the tag-action.

VLAN Traffic flow in GPONs

GPON ports have no interface association, unlike the Ethernet ports. Therefore, rather than adding a port interface to a VLAN membership to enable traffic flow, for a GPON port to pass traffic carried on a VLAN, a service-tag action must be created specifying the VLAN. Also, the VLAN must already be created on the E-Series.

- To view the VLANs associated with specific GPON ports, click the node > **GPON-X** > **GPON**: on the Navigation Tree, and then click the **VLANS** tab.
- To view the VLANs associated with an ONT, click **ONTS** on the Navigation Tree, and then click the **VLANS** tab.
- To view the ONT services associated with a specific VLAN, click **VLANS** on the Navigation Tree, click the particular VLAN from the list that appears in the Work Area, and then click the **Service Associations** tab.

Configuration guidelines

- You can configure only one of the following attributes on a given VLAN on a given interface:
 - Trunk interfaces:
 - VLAN member
 - Native VLAN
 - Edge interfaces:
 - VLAN member
 - Tag-action
 - Native VLAN
 - Access interfaces:
 - VLAN member
 - Tag-action
- When you create a tag-action, the interface automatically becomes a member of the VLAN specified in the tag-action. See example below.
- Trunk and edge interfaces are always associated with at least one VLAN, through the native VLAN attribute (VLAN 1, by default). Access interfaces do not support a Native VLAN, however, tag actions can be used to assign untagged traffic to a VLAN ID.
- An interface role (edge or trunk) cannot be modified when it is a member of a VLAN.
- ERPS domains can only be associated with VLANs through becoming a member of the VLAN. All VLANs must be created on each node of the ERPS ring and be associated with the ERPS domain membership. (Do NOT associate VLANs to underlying ERPS port interfaces.)
- When you add VLAN router interfaces, you are setting the static multicast router location on specific ports, allowing the system to know which interface has the multicast router. You can add many interfaces to the static location. For example, if the system is connected to the multicast router via RSTP, you would want to add the static location to both ports (for instance, 1/g1 and 2/g1). RSTP will block one port at a time. When the port switches to the other port, there would still be connectivity to the multicast router.
- When the IGMP profile has the Router Learning Mode configured for 'static-only', IGMP Proxy will not allow a static multicast router ('mrouter') interface to be a multicast destination. If the interface could become a multicast destination in the event of a network topology change, the interface should not be configured as a mrouter interface.
- Even if no interface is currently using VLAN 1 as the Native VLAN, it is still off limits for user provisioning, including use as the Management VLAN or ERPS control VLAN.
- For modular chassis nodes, any VLAN created on the system is automatically mapped to the Stacking Ports. The remaining port interfaces in the system must be a VLAN member for traffic to pass on the VLAN through the interface.

Example application configuration for tag-action versus VLAN members:

For incoming E7 traffic where the voice and video services are single-tagged and the per port data VLANs are double-tagged:

- VLAN-per-port for data traffic
- Single voice service VLAN
- Single video service VLAN

Configure the following:

1. Create a voice VLAN and a video VLAN and make each interface a member of each VLAN.
2. Create a no-match add-tag action to the interface for all remaining VLANs to add an outer data service tag.

Before starting

Before starting the VLAN member creation process, check that the following conditions are met:

- A VLAN must be created prior to adding VLAN member interfaces.
- The Ethernet LAG interface or ERPS domain must exist to make it a VLAN member.

VLAN Tag Action to an Ethernet Interface

This topic describes VLAN tag actions that can be added to an edge link interface. A tag-action allows you to add or change VLAN tags on packets arriving on a specified GE or 10GE interface. You can use a set of matching criteria to select which packets to transform. Or, packets can be transformed unconditionally. While tag actions are typically described as being applied to packets received (ingress traffic) on an Ethernet port interface, the reverse of the tag action is applied to packets transmitted (egress direction) from the interface. For example, if an Add-Tag action is applied to an interface, that same tag is stripped from packets leaving the interface.

The following VLAN tag actions can be assigned to an edge or access link interface within the networking domain:

- **Add Tag** - The "add-tag" action adds an outer tag either on all packets on an interface, or only on packets that match a specified VLAN ID or P-bit.
- **Add 2 Tags** - The "add-2-tag" action adds an outer tag and an inner tag with the specified VLAN IDs on packets of untagged traffic.
- **Change Tag** - The "change-tag" action changes the outermost tag to the specified VLAN ID on packets that match a specified VLAN ID.
- **Add and Change Tag** - The "add-and-ch" action changes the tag and then adds an outer tag on packets that either match a specified VLAN ID or are priority-tagged frames.


Note: Tag-actions can only be applied to edge or access link interfaces. To create tag actions for an ONT Ethernet or xDSL port subscriber service, you must use Service Tag Actions.

Configuration guidelines

- Tag actions can only be assigned to edge or access link interfaces within the networking domain and are defined with respect to ingress traffic.
- When you create a tag-action, the interface automatically becomes a member of the VLAN specified in the tag-action.
- When you create an add-tag action, the P-bit of the incoming tag is automatically copied into the P-bit of the added tag.
- You can configure only one of the following attributes on a given VLAN on a given interface:
 - VLAN membership
 - Tag-action (add, add-2-tags, change, or add-and-change)
 - Native VLAN
- To forward untagged traffic, an access interface must have an add-tag action, add-and-change, or add-2-tags action applied to untagged frames, assigning the traffic to a VLAN.
- An edge or access link interface can be associated with a particular VLAN either through a VLAN membership or through a tag-action, but not through both.
- Interfaces can be associated with additional VLANs through memberships or tag-actions.
- An interface role cannot be modified when it is assigned a tag-action.
- The E-Series can perform the following tag actions on a VLAN that has IGMP enabled:
 - Add-tag action to untagged Ethernet frames
 - Change-tag action to change the VLAN ID at network administrative boundaries
- If no matching criterion is assigned to a tag action, then the tag action is performed on all packets entering the specified interface, except for the packet traffic on VLANs that have an associated membership with the interface. When the match rule is set to ignore and p-bit any that is the same as no matching criterion assigned. You can see this if you create a tag action and leave those fields blank. It will be set to ignore.
- For double-tagged packets, the outer VLAN ID must match a provisioned VLAN. The inner tag does not have to match a provisioned VLAN ID.
- A priority tagged frame has a new tag added or changed-to with the specified tag value. Upon egress, the frame still carries the original priority bits onto the newly added or change-to VLAN tag.
- If an Ethernet interface has both a VLAN tag action and a policy map applied, the tag action function is applied to the interface before the policy map function.

Example add-tag-action configurations

Create an action to add a tag (VLAN 400) to all traffic on the associated interface.



Create

Create VLAN Tag Action

Associated Interface * 1-1-GE-1

Action Performed * add-tag


Matching Tag

Matching pbit

S-VLAN (Outer Tag) * 400

CREATE CANCEL

Create an action to add a second tag (VLAN 400) to all VLAN 401 traffic on the associated interface.



Create

Create VLAN Tag Action

Associated Interface * 1-1-GE-1

Action Performed * add-tag


Matching Tag 401

Matching pbit

S-VLAN (Outer Tag) * 400

CREATE CANCEL

Create an action to add a tag (VLAN 400) to all traffic based on the P-bit for the associated interface.



Create

Create VLAN Tag Action

Associated Interface * 1-1-GE-1

Action Performed * add-tag


Matching Tag ignore

Matching pbit pbit-6

S-VLAN (Outer Tag) * 400

CREATE CANCEL

Create an action to add a tag (VLAN 400) to all traffic based on the P-bit and outer tag on the associated interface.



Create

Create VLAN Tag Action

Associated Interface * 1-1-GE-1

Action Performed * add-tag

Matching Tag 401

Matching pbit pbit-6

S-VLAN (Outer Tag) * 400

CREATE CANCEL

Example application configuration for tag-action versus VLAN membership:

For incoming traffic where the voice and video services are single-tagged and the per-port-data VLANs are double-tagged:

- VLAN-per-port for data traffic
- Single voice service VLAN
- Single video service VLAN

Configure the following:

1. Create a voice VLAN and a video VLAN and add the interface to each of the VLAN memberships.
2. Create a no-match add-tag action to the interface for all remaining VLANs to add an outer data service tag.

Before starting

Before starting the VLAN tag action process, check that the following conditions are met:

- The VLANs must already exist in order to create tag actions.
- The interface that is the target for the tag action must have the edge or access role assigned.

Parameters

You can provision the following parameters for VLAN tag actions to an Ethernet interface:

Parameter	Description	Valid Options
Matching Tag	Specifies the match criteria. If you select Enter Value, then you must also specify the VLAN ID value to match. VLANs can be specified by name, or by numeric VLAN ID. (VLAN IDs 1002-1005 are reserved for E7 operation.)	ignore, any, untagged, priority, Enter Value (then enter a VLAN ID 2-4094)
Matching pbit	P-bit value that specifies the VLAN priority value to match.	pbit-0 to pbit-7, pbit-any, or leave blank
Action Performed*	Specifies whether to add, change, or add and change tags to incoming packets.	add-tag add-2-tags ch-tag add-and-ch
S-VLAN (Outer Tag)*	New VLAN ID. (VLAN IDs 1002-1005 are reserved for E7 operation.)	2-4093
C-VLAN (Inner Tag)	New VLAN ID. (VLAN IDs 1002-1005 are reserved for E7 operation.) Note: This parameter only applies if the Action Performed selection is "add-2-tags" or add-and-change.	ignore, any, untagged, Enter Value (then enter a VLAN ID 2-4093)

* Required field

Service-Tag Action to a Provisioned Service

Applying service tag actions performs the following two actions:

- **Classification** – the service match lists are for matching and classifying traffic into services.
- **Marking** – the service tag actions are for marking the traffic with the appropriate VLAN and P-bit required by the service.

Service-tag actions can be used for many subscriber ports that require the same actions to be performed. All configured Ethernet services exist in a VLAN-tagged service flow, requiring an associated service-tag action.

Service-tag actions are applied to subscriber service provisioning as follows:

- Video and data services provisioning on an xDSL or ONT Ethernet port includes a reference to a service-tag action.
- Voice services provisioning on a VDSL2 Voice (POTS) port references an IP Host that sets the output service VLAN and points to a system default service-tag action that specifies the traffic as untagged and assigns the appropriate priority value.
- Voice services provisioning on an ONT Voice (POTS) port implicitly references a Voice Service IP Host that includes a reference to a service-tag action.

The following tag actions can be assigned to a service port.

Add a single tag	Adds a single outer tag to the matched traffic. • Can be applied to tagged and untagged traffic. • Supported for DSCP-to-P-bit translation. • Supported for Full Bridge (FB) interface services. • NOT supported for Residential Gateway (RG) interface services.
Add 2 tags	Adds an inner and outer tag to the matched traffic*. • Can be applied to untagged traffic, only. • Supported for DSCP-to-P-bit translation. • NOT supported for FB or RG interface services.
Add and change tag	Adds an outer tag to the matched traffic and changes the inner tag. • Can be applied to tagged traffic, only. • Supported for RG interface services. • NOT supported for FB interface services. • NOT supported for DSCP-to-P-bit translation.

Change tag**Changes the outer tag on the matched traffic.**

- Can be applied to tagged traffic, only.
- Supported for FB or RG interface services.
- NOT supported for DSCP-to-P-bit translation.

*To properly process ingress double tags, the GE network interface (uplink) must be configured as a Trunk role.

Determining the tag-action values

When a service is provisioned, a match list and service-tag action specify the classifying and marking of packets from the subscriber port into the service VLAN.

You have the following options for determining the added- or changed-tag values:

- **Specify the VLAN value explicitly.**

This option is used for a VLAN-per-service provisioning model (N:1). A service carried on an N:1 VLAN applies to multiple subscriber ports, where a single match list and tag action can indicate the service.

- **Reference the VLAN value from the service provisioning on the subscriber port interface.**

This option is used for a VLAN-per-port provisioning model (1:1). A service carried on a 1:1 VLAN is the same for each subscriber except the customer tag is unique for each single subscriber port, so you can define the match list and tag action pair such that multiple subscriber ports can reference it. This is accomplished by indicating "Specified in Service" for the output tag in the tag action, and then the subscriber-specific value for the output tag is defined when the service is provisioned on the port.

Note: The service provisioning process takes place after the creation of a service-tag action.

You also have the following options for determining the service-tag action P-bit for services provisioned on xDSL ports:

- **Specify the P-bit value explicitly.**

Enter the P-bit value to use for P-bit marking the outer tag.

- **Reference a Layer-3 Priority map.**

Use a DSCP table or an IP Precedence table to map to a P-bit value. From the xDSL port associated interface, you can select to use the system-default profiles ("access") or specify another custom profile.

Configuration rules:

- Only one service per port can use the DSCP-Bit mapping function. However, a second service of untagged traffic on the same port that matches a specific OUI value is supported, where that OUI-matching service takes precedence for any matching frames.
- Layer-3 priority maps work with add-tag and add-2-tags tag actions.

- If you select “Map a layer-3 Priority,” in the service tag action, then when creating the service, you must select a PON CoS parameter of User1-4 or CoS1-4. Do not leave the system-default parameter of “derived” as there is no single unique P-bit specified from which to derive the intended CoS.
- The T-Series ONTs do not mark the P-bit of broadcast and multicast frames based on DSCP. T-Series ONTs handle all broadcast / multicast packets from the subscriber as best-effort and mark them with the default P-bit of the service.

Provisioning sequence:

- Create a DSCP-to-P-Bit map, if you do not want to use the system-default.
- Assign the map to the ONT port, or leave the system-default map "access" assignment.
- Create a service-tag action with an Outer P-Bit Source of **Map a layer-3 Priority**.
- Create a service using the defined service-tag action.

The following table shows the 802.1p to DiffServ DSCP pre-defined, system-default mapping table ("access").

Default P-bits	IP Precedence	DiffServ Code Point**	Service
0	7	CS6, CS7 (48-56)	Network control
0	6	EF (46)	Voice, NSP, CES
5	5	CS5 (40)	Network Management
4	4	CS4 & AF41-43 (32-38)	Video
3	3	CS3 & AF31-33 (24-30)	Data3
2	2	CS2 & AF21-23 (16-22)	Data2
1	1	CS1 & AF11-13 (8-14)	Data1
0	0	0	Best Effort

**The E7 supports alphanumeric DSCP values that denote forwarding classes such as AF21, AF31 in addition to numeric value assignment.

- **Preserve the existing P-bit value upstream.**

Indicate that the E7 preserves the P-bit value set in service provider supplied Residential Gateways (RGs) by specifying a P-bit value of "copy."

- **Promote the existing P-bit value upstream.**

Indicate that the E7 promotes the P-bit value set in service provider supplied Residential Gateways (RGs).

- The P-bit of the inner tag is promoted to the P-bit of the outer tag, allowing the network QoS to handle the priority of the frame consistently.
- For the inner tag P-bit, you can specify the value of "same-as-outer".

Note: The upstream traffic on a service VLAN must have the same P-bit value as is set for downstream traffic. The P-bit value must also be consistent with the class of service type as defined in the PON CoS table.

Note: For VDSL2 and GPON subscribers (if PON CoS = Fixed), downstream traffic with P-Bit values not explicitly defined by the service(s) on the port will be passed through with unmetered bandwidth.

- **(GPON only) Specify the CoS.**

Enter the CoS ID to use for P-bit marking the outer tag to the lowest P-bit value associated with this CoS. This value identifies the PON class of service.

At the GPON and ONT, individual services with priority bit settings are mapped to a particular class of service over the PON, as shown:

GPON CoS	Scheduling Class Type	Default P-bits	Service
CoS 4	Expedited Forwarding	5, 6, 7	Network control, voice, T1
CoS 3	Assured Forwarding	4	Video
CoS 2	Assured Forwarding	3	Application signaling, TLAN
CoS 1	Best Effort Forwarding	0, 1, 2	High-speed internet

Note: The upstream traffic on a service VLAN must have the same priority value as is set for downstream traffic.

VLAN Traffic flow in subscriber ports

For subscriber ports, rather than adding a port interface to a VLAN membership to enable traffic flow, a service-tag action must be created that specifies the VLAN and traffic priority. This service-tag action is then referenced when the service is provisioned. The VLAN must already be created on the E7 for a subscriber port to pass traffic carried on a VLAN.

For xDSL

- To view the VLANs associated with a specific xDSL port, on the Navigation Tree, click the port of interest, and then click **Associated Interface > VLANs**.
- To view the xDSL services associated with a specific VLAN, on the Navigation Tree, click **VLANs**, click the particular VLAN from the list that appears in the Work Area, and then click the **Service Associations** tab.

For GPON

- To view the VLANs associated with specific GPON ports, click **E7 > GPON-4 > GPON**: on the Navigation Tree, and then click the **VLANs** tab.
- To view the VLANs associated with an ONT, click **ONTS** on the Navigation Tree, and then click the **VLANs** tab.
- To view the ONT services associated with a specific VLAN, click **VLANs** on the Navigation Tree, click the particular VLAN from the list that appears in the Work Area, and then click the **Service Associations** tab.

The subscriber ports support the same VLAN services as the E7 Ethernet ports, including IEEE 802.1Q VLAN tagging and IEEE 802.1ad VLAN stacking (Q-in-Q). The following service models are supported.

- **N:1/VLAN-per-service model:** each service is assigned a dedicated VLAN where multiple subscriber ports are assigned to the VLAN for a single service. This model is often referred to as N:1. An additional tag can be added to limit the size of a group for improved reliability. A service carried on an N:1 VLAN applies to multiple subscriber ports, allowing a single match list and tag action to be used to denote the service.
 - Subscriber tag = 110
 - Match tag = 110
 - Change tag = 200 (service VLAN)
- **1:1/VLAN-per-port model:** the traffic frames from each subscriber port are assigned a unique VLAN ID (tag). The tag is applied to a frame before the traffic is aggregated for transport. An additional tag can also be added to group subscriber ports in logical categories.

A service carried on a 1:1 VLAN is the same for each subscriber, except the customer tag is unique per subscriber. You can define the match list and tag action pair such that multiple subscriber ports can reference it. This is accomplished by assigning an attribute for the output tag in the tag action that indicates the value of the output tag is subscriber specific. The customer specific tag is included in the service definition when provisioning the service.

 - Subscriber tag = 110
 - Match tag = 110
 - Change tag = 110
 - 1:1 VLAN-per-Port, Single Tagged – Subscribers on a system are provisioned with individual Customer Tags (C-Tags). Add tag, or change tag action.
 - 1:1 VLAN-per-Port, Double Tagged – Subscribers on a system are provisioned with individual Customer Tags (C-Tags). The S-Tag can be added by the service tag action. Add 2-tags, or change-and-add tag action.

The 1:1 model does not require any new or different traffic handling techniques to isolate and manage subscriber traffic. With the 1:1 service delivery model, each subscriber's traffic (except broadcast video) is sent down a dedicated single VLAN on a per subscriber basis. When ATM framing is used for ADSL/ADSL2+ modes, the single ATM PVC maps to a single VLAN.

Within the GPON OLT and ONT subsystem, MAC learning and switching occurs within a single, outer VLAN ID, referred to as the VLAN C-tag. Packets on the PON always have a C-tag, which consists of a provisioned VLAN ID and p-bit. The ONT is responsible for classifying subscriber traffic and adding/modifying the C-tag.

Example service-tag action configurations for data services

VLAN-per-service

Applies the specified service-tag action to the matched traffic.

Matches all untagged traffic on the xDSL port.

- Adds VLAN ID 300
- Marks with P-bit 0

The upstream traffic on a service VLAN must have the same P-bit value as is set for downstream traffic.

VLAN-per-subscriber

Applies the specified service-tag action to the matched traffic.

Double tags all untagged traffic arriving on the xDSL port.

- Outer tag is explicitly set to 300
- Inner tagged is defined when adding the service to a port
- Upstream traffic is marked with P-bit 0

The upstream traffic on a service VLAN must have the same P-bit value as is set for downstream traffic.

Before starting

Before starting the service tag action creation process, check that the following conditions are met:

- All VLANs that are to be specified in the service tag action have already been created.
- The match list that is to be specified in the tag action has already been created.
- The DSCP or IP Precedence values are defined, if using a map other than the system default.

Parameters

You can provision the following parameters for service tag actions:

Parameter	Description	Valid Options
Name*	Descriptive name of service tag action.	text string
Action*	Specifies the tag action to occur.	Add Tag, Add 2 Tags, Change Tag, Add and Change Tag
Match Criteria List	Name and ID of the service match list to use for this tag action.	Any established service match list
Outer Tag	VLAN ID for the new outer tag. The specified VLAN must already be provisioned in the system. Alternately, "use-svc" indicates that the VLAN ID will be specified when you provision a service on the VDSL port.	Specified in Service, Enter Value (1-4093)
Inner Tag	Name of VLAN for the new inner tag. Alternately, "use-svc" indicates that the VLAN ID will be specified when you provision a service on the VDSL port. Note: This parameter only applies if the selected action is "Add 2 Tags" or "Add and Change Tag."	Specified in Service, Not Used, Enter Value (1-4093)
P-Bit Source	Specifies where to derive the P-bit value for the outer tag, selecting one of the following: - Set an explicit P-bit value in the tags. Note: The upstream traffic on a service VLAN must have the same priority as is set for downstream traffic. - Mapping to either a DSCP map or IP Precedence values. Note: The DSCP map or IP Precedence map is specified in the xDSL port interface.	Specify P-bit ‡, Map a layer-3 priority, From CoS (GPON only)
P-Bit	The P-bit value to use for P-bit marking the outer tag when the P-bit Source is set to "Specify P-bit." Note: When the value of "copy" is selected, the E7 honors the P-bit value of the incoming traffic and passes the existing P-bit value upstream.	0-7, copy 0 ‡

*Required fields

‡ Default

xDSL Configuration Guidelines for Service Tag Actions

- The P-bit value specified in the service tag action must be consistent with the DiffServ Code or IP Precedence as defined in the default profiles, or custom profiles referenced at the xDSL interface.
- If the specified P-Bit source is "copy," the P-bit value cannot be specified because the P-bit value of the incoming traffic is honored.
- Each service tag action references a service-match list (classify) and specifies a tag action (mark and tag).
- Each service tag action is associated with a specific service.
- Multiple services and service tag actions can be applied to a single xDSL port. However, if multiple services on the same xDSL port use the same outer VLAN ID, the priority must be specified for each service in the associated service tag action using the **P-bit Source = Specify P-Bit**. Using the P-Bit Source = Map a layer-3 priority is not supported for this scenario.
- The same service tag action can be used on multiple xDSL ports.

- Each E7 shelf, or E7 Modular Chassis system has a capacity of 256 service tag actions. For configurations that require large numbers of VLANs, Calix recommends selecting the **Specified in Service** value for the "Outer Tag" parameter, and then specifying the VLAN ID when provisioning the service.
- For traffic to flow on the VLAN specified in a service tag action, the VLAN and uplink must already be created on the E7.
- For an xDSL port to pass traffic carried on a VLAN, a service tag action must be created specifying the VLAN and the VLAN must already be created on the E7.
- For installations using multiple PVC assignments, individual PVCs may be mapped to specific VLANs that are applied to the port. For example, installations that have multiple PVC assignments (e.g. 0/33, 8/33, 0/35) may be mapped to individual data services (e.g. "Data 1", "Data 2", "Data 3") that would then be applied to the port.
- If no matching criteria are assigned to a tag action, then the tag action is performed on all packets entering the xDSL port. For example, when a match list has a single untagged rule with the outer tag set to "ignore" and the Outer P-bit set to "Pbit-none."
- Traffic must have a priority assigned in order to be scheduled on the xDSL port.
- The "change-tag" action changes the VLAN ID in the outermost tag.
- For double-tagged packets, the outer VLAN ID must match a provisioned VLAN. The inner tag does not have to match a provisioned VLAN ID.
- Edge interfaces can only use single-tag actions when DHCP Snooping is enabled on the associated service VLAN.
- Within the xDSL subsystem, MAC learning and switching occurs within a single, outer VLAN ID, referred to as the VLAN C-tag. Packets on the xDSL port always have a C-tag, which consists of a provisioned VLAN ID and priority for classifying subscriber traffic and adding/modifying the C-tag.
- The E-Series reserves four VLAN values for system operation. The default values for these VLANs are 1002, 1003, 1004 and 1005. You can change these values to another set of four consecutive VLANs, if required.
- Although the "Native VLAN" (VLAN 1) does not apply for traffic arriving at the xDSL port, the default VLAN 1 is off limits for user provisioning.
- IGMP snooping is provisioned on a VLAN basis.
- IGMP snooping is only enabled on the outer VLAN ID as only single-tag multicast VLANs are supported.
- The following tag actions are supported for VLANs with IGMP enabled:
 - Add-tag (only supported for untagged traffic)
 - Change-tag
- Services provisioned with DSCP or IP Precedence translation to P-bit values only support:
 - add-tag
 - add-2-tags

- change-tag (R2.6.30 and higher)
- One service per port can use the DSCP or IP Precedence to P-bit mapping function
- Either DSCP or IP Precedence can be used for a service, not both
- Once a tag is used as a single-tagged VLAN, it cannot also be used in double-tagged VLANs. For example, if VLAN 300 is single tagged:
 - It cannot be used as the inner tag of a double-tagged VLAN because the double-tagged VLAN would require VLAN 300 to push an outer tag.
 - It cannot be used as the outer tag in a double-tagged VLAN because then it would get popped in the downstream direction.
- In double-tagged VLANs, if the inner tag is applied to multiple xDSL subscriber ports, it is an N:1 VLAN.
- For single-tagged VLANs, if the outer tag is applied to multiple xDSL subscriber ports, it is an N:1 VLAN.
- If a given VLAN, single- or double-tagged is applied to a single xDSL subscriber port, it is 1:1 VLAN.

GPON Configuration Guidelines for Service Tag Actions

- The P-bit value specified in the service-tag action must be consistent with the class of service type as defined in the PON Class of Service (CoS) table.
- The ONT requires that SIP traffic be set to Expedited Forwarding. Therefore, the P-bits set in the service tag action must map to the Expedited Forwarding CoS.
- Each service-tag action references a service-match list (classify) and specifies a tag action (mark and tag).
- Each tag action is associated with a specific service.
- The ONTs do not support tagged and untagged match rules on the same ONT for video services. Only one classification rule is allowed for video services per ONT.
- The logical components within the RSG are functionally similar to a standard GPON ONT with a subtended 3rd party residential gateway (RG): The ONT receives tagged traffic from the downstream RG and sends it upstream single-tagged where the OLT may or may not add a second tag.
- Services provisioned on a Full Bridge interface only support:
 - add-tag
 - change-tag
- Services provisioned on a Residential Gateway interface only support:
 - change-tag
 - add-and-change-tag

- The following tag actions are supported for VLANs with IGMP enabled:
 - Add-tag
 - Change-tag
- Services provisioned with a tag action Outer P-bit source of **Map a layer-3 Priority** (DSCP translation to P-bit values) support the following:
 - add-tag
 - add-2-tags
 - Only one service per port can use the DSCP-Bit mapping function. However, a second service of untagged traffic on the same port that matches a specific OUI value is supported, where that OUI-matching service takes precedence for any matching frames.
 - Each port is independent where it can use a different DSCP map.
 - GPON services only support DSCP to P-bit mapping, (while xDSL services support DSCP or IP Precedence for a service, not both).
 - A PON CoS parameter of User1-4 or CoS1-4 must be used. Do not leave the system-default parameter of “derived” as there is no single unique P-bit specified from which to derive the intended CoS.
 - All of the traffic classified by the DSCP map and P-bit-marked according to the DSCP map table is put into a single queue for upstream scheduling by the OLT.
 - In the event of congestion on the upstream bandwidth allocation, traffic in the queue with the highest P-bit is transmitted first, thus preserving the DSCP set by the downstream router.
- Multiple services and service-tag actions can be applied to a single ONT Ethernet port.
- The all-untagged match rule cannot be applied to multiple service-tag actions on the same port. For example, one of the untagged match rules must point to a specific OUI. The system applies the tag action that has the most specific match rule.
- The same service-tag action can be used on multiple ONT Ethernet ports.
- Each E7 shelf, or E7 Modular Chassis system has a capacity of 256 service tag actions. For configurations that require large numbers of VLANs, Calix recommends selecting the **Specified in Service** value for the "Outer Tag" parameter, and then specifying the VLAN ID when provisioning the service.
- The upstream and downstream values specified in the Ethernet Bandwidth Profile and P-bit values defined in the service-tag action must be consistent with the class of service type as defined in the global CoS table.
- For traffic to flow on the VLAN specified in a service-tag action, the VLAN and uplink must already be created on the E7.
- For a GPON port to pass traffic carried on a VLAN, a service-tag action must be created specifying the VLAN and the VLAN must already be created on the E7.

- If no matching criteria are assigned to a tag action, then the tag action is performed on all packets entering the GPON port. That is, when a match list has a single untagged rule with the outer tag set to "ignore" and the Outer P-bit set to "Pbit-none."
- At least one tag must be specified in the match criteria because the system cannot change the VLAN ID for untagged traffic.
- Traffic must have P-bits assigned in order to be scheduled (prioritized) on the GPON.
- The "change-tag" action changes the VLAN ID in the outermost tag.
- For double-tagged packets, the outer VLAN ID must match a provisioned VLAN. The inner tag does not have to match a provisioned VLAN ID.
- Edge interfaces can only use single-tag actions when DHCP Snooping is enabled on the associated service VLAN.
- To properly process ingress double tags, the GE network interface (uplink) must be configured as a Trunk role.
- The E7 reserves four VLAN values for system operation. The default values for these VLANs are 1002, 1003, 1004 and 1005. You can change these values to another set of four consecutive VLANs, if required.
- Although the "Native VLAN" (VLAN 1) does not apply for traffic arriving at the E7 GPON port, the default VLAN 1 is off limits for user provisioning.
- IGMP snooping is provisioned on a VLAN basis.
- IGMP snooping is only enabled on the outer VLAN ID as multicast traffic is typically transported throughout the network with a single VLAN ID tag.
- For double-tagged VLANs, the same inner tag cannot be used with different outer tags. Any given (inner) tag received upstream on the PON can only push a single (outer) tag. For example, if VLAN 100 is received on the PON, tag 200 can be pushed, and then tag 200 will be popped in the downstream.
- Once a tag is used as a single-tagged VLAN, it cannot also be used in double-tagged VLANs. For example, if VLAN 300 is single tagged:
 - It cannot be used as the inner tag of a double-tagged VLAN because the double-tagged VLAN would require VLAN 300 to push an outer tag.
 - It cannot not be used as the outer tag in a double-tagged VLAN because then it would get popped in the downstream direction.
- For video, the specified P-bit must match the CoS type of bandwidth profile.
- The ONT does not support heterogeneous tagging in the MDU case for video service.
- In double-tagged VLANs, if the inner tag is applied to multiple ONT subscriber ports, it is an N:1 VLAN.
- For single-tagged VLANs, if the outer tag is applied to multiple ONT subscriber ports, it is an N:1 VLAN.
- If a given VLAN, single- or double-tagged is applied to a single ONT subscriber port, it is 1:1 VLAN.

- VLAN Association to voice services:
 - The IP Host object references a service tag action that indicates the voice service VLAN by one of the following methods:
 - An IP host can directly specify an outer tag, when the referenced service tag action states "Specified in Service" for the Outer Tag parameter.
 - This allows you to apply a different VLAN tag for the same service type (SIP, TDM gateway, H.248 gateway, MGCP, or T1 PWE3) on each ONT.
 - An IP host can reference a service tag action that specifies the service VLAN ID for the provisioned voice services.
 - This allows you to apply the same VLAN tagging for all ONTs that use the referenced service tag action object.
 - The VLAN assignment must be different between the VoIP IP host and the subscriber port. Sharing VLANs between a subscriber port and a VoIP IP host is not supported.
- For RG specific (Gateway) services (as opposed to half-bridge [HB] or full-bridge[FB]), only "Change Tag" and "Add and Change Tag" actions are supported. The concept of provisioning this way was derived from the same topology that is observed today in deployments that utilize external non-Calix RGs. This is especially used in scenarios where a third-party RG is downstream of a regular GPON ONT. From a deployment perspective, the HGW ONT then supports two logical components; a regular GPON ONT and an RG. In this scenario, the ONT would receive tagged traffic from the downstream RG and the ONT would either change the received tag and send it upstream single-tagged or either change the received tag, send it upstream, and the OLT would then add the second tag.

IPv4 Security Features

The E-Series provides the following IPv4 security features:

- DHCP Snoop, along with the following related features:
 - MAC Forced Forwarding (MACFF)
 - IP Source Verification
- DHCP Relay Option 82
- DHCPv4 Proxy

DHCP Snoop

When the VLAN parameter "DHCP Snoop" is enabled, the E-Series tracks all DHCP activity on that VLAN and creates/updates a table of DHCP leases granted. Calix recommends DHCP Snoop be enabled for all residential subscriber services all DHCP activity on a VLAN and creating/updating a table of DHCP leases granted. The ability to track IP address-MAC address-interface bindings allows the E-Series to provide intelligent manageability as well as the following security capabilities:

- Host limit per VLAN per port (Lease Limit)
- Blocking multiple ports from using the same IP address
- Blocking MAC spoofing across ports (DHCP Lock)

In addition, MAC Forced Forwarding (MACFF) is only possible in conjunction with DHCP Snoop.

Operation

When DHCP Snoop is enabled, all DHCP traffic is monitored by the E-Series in order to learn the IP address assignment for hosts accessed through it. The Snoop Agent maintains a lease database for each host which boots via DHCP. In this fashion, the E-Series maintains a set of authenticated host entries, with both IP and MAC addresses.

Unless an IP host is statically assigned, if no DHCP lease has been established, all packets arriving from the subscriber port will be discarded. Only DHCP Discover and Request message packets are forwarded.

Configuration considerations

- An operator may manually delete a lease entry from the table. However, if an IP Source Verification is enabled on the VLAN, deleting the DHCP lease entry results in the subscriber traffic being dropped until the subscriber's IP host requests, and is granted, a new IP address using DHCP.

- When DHCP snooping is enabled on a subscriber VLAN, the E7 drops all DHCP server communication originating from subscriber Ethernet interfaces. Calix recommends DHCP snooping be enabled for all residential subscriber services.
- When DHCP snooping is enabled, DHCP requests for VLANs do not appear in an Ethernet port mirror session.
- The craft management ports cannot use IP addresses from the same subnet where DHCP Snooping is enabled.
- DHCP snooping is not supported on management VLANs.
- DHCP Snoop must be enabled on the VLAN to support the AE Discovery Event operation.
- DHCP lease limits allowed on a GPON Ethernet port or xDSL port are specified in an Ethernet Security Profile that is applied to the ONT Ethernet port or xDSL port associated interface. Only VLANs with DHCP Snooping on the port are subject to the DHCP lease limit. In the VDSL2 and GPON subsystems, DHCP leases cannot be learned without also applying a limit to the number of learned leases on the port. Each Ethernet port has an associated Port Security Profile that can limit DHCP leases in a range of 1-16. A security profile used by an xDSL port must have a DHCP lease limit value of 10 or less.
- For VDSL2 and GPON, IP and MAC addresses may be dynamically learned using either DHCP Snooping or manually provisioned with static IP/MAC addresses. Services with static subnets (without MAC address specification) may also be provisioned for IP Source Verification, but are bound to the port only by IP address.
 - DHCP and Static IP host: Binds IP and MAC address to a Port
 - Static IP Subnet: Checks individual host IP addresses to the subnet and binds the subnet to a port
 - The GPON systems have the following limits on the number of IP stations (active DHCP leases, static IP hosts, and dynamically learned hosts within a static subnet):
 - ♦ The ONT supports up to 256 “IP stations” per ONT Ethernet port.
 - ♦ The ONT supports up to 1024 total “IP stations” per ONT.
 - The following capacities apply to Static IP Addresses/Subnets within the E7 VDSL2 subsystem.
 - ♦ Only 1 static IP subnet can be provisioned per xDSL Ethernet service.
 - ♦ Up to 4 static IP hosts can be provisioned per xDSL Ethernet service.
 - ♦ Up to 8 static IP hosts/subnets can be provisioned per xDSL port.
 - ♦ The VDSL2 card supports up to xxx "IP stations" per xDSL port. The number of IP stations includes active DHCP leases, static IP hosts, and dynamically learned hosts within a static subnet.

MAC Forced Forwarding

MAC Forced Forwarding (MACFF), defined by RFC 4562, provides a method for securing end-user traffic on an Ethernet access network, even if subscribers share the same IP sub-network. In VLAN per Service applications, MACFF ensures that subscriber Ethernet frames sent upstream are only forwarded to the MAC address of a known IP gateway (Access Router or AR), and ensures that traffic from one subscriber interface cannot be sent directly to another, because the IP gateway provides IP-layer connectivity between these hosts, reducing the chance that malicious traffic can be transmitted between ports.

With MACFF enabled:

- The access node or ONT learns the IP and MAC addresses of the subscriber device and upstream router.
- The E-Series's Address Resolution Protocol (ARP) proxy intercepts ARP requests from the subscriber device requesting the MAC address of a device at another premises, and responds with the MAC address of the upstream router regardless of the IP destination. This forces subscriber traffic toward the router. The access node and ONT do not respond to ARP requests for a device residing on the same link as the subscriber device.
- The E-Series installs a filter to inspect the destination MAC address of upstream traffic originating from the subscriber, and discards packets to MAC addresses other than the allowed target upstream router.
- MAC FF supports the following:
 - IP hosts learned via DHCP Snooping
 - IP hosts statically provisioned (IP and MAC)
 - IP hosts statically provisioned (IP address only)
 - IP subnets statically provisioned (single IP address/MAC)
 - A single source MAC address with multiple source IP addresses, when the MAC address is not specified

MACFF allows the access network Service Provider to use the AR(s) to perform security filtering, policing, and accounting of all customer traffic. Although subscribers' hosts are connected at Layer 2, the AR ensures IP-layer connectivity between these hosts. To ensure a secure environment in the Service Provider's access network, layer-2 traffic paths should not circumvent an AR.

Layer 2 security is enhanced in the E-Series by blocking the following:

- Unknown unicast traffic going downstream
- Unnecessary broadcast traffic going downstream (any broadcast other than DHCP initiated by subscribers devices)
- All unsolicited multicast traffic (any multicast traffic that has not been requested by an IGMP join from a subscriber)

Configuration guidelines

- MACFF can only be used in conjunction with DHCP Snoop.
- Enable MACFF to limit broadcast traffic on the VLAN domain when using the VLAN Per Service model. It is not required for the VLAN per Port (single and double-tagged, Q-in-Q) data model or PPPoE because this VLAN model by definition limits the broadcast domain to a single access node and the router and makes it easier to manage the broadcast load. However, MACFF should still be enabled on these VLANs when the subscriber edge does not include a residential gateway.
- Limit the number of MAC addresses that the network is exposed to by enabling MACFF on both data and video VLANs.
- If the subscriber edge does not include a residential gateway, enable MACFF when using the VLAN Per Subscriber model.
- Configure static and dynamic subscribers on separate VLANs when enabling MACFF.

IP Source Verify

Enabling the VLAN parameter "IP Source Verify" binds the IP address and MAC address to the physical ONT Ethernet port or xDSL port, preventing subscribers from assigning an IP address to a device and passing traffic on it. IP Source Verification for Static IP hosts requires MAC FF be enabled. Calix recommends IP Source Verification be enabled for all residential and business internet access services.

IP source verify is not supported for a source MAC address with multiple IP addresses.

DHCP Relay Option 82 and LDRA

This topic covers the following information in regard to DHCP Relay Option 82 and LDRA:

- Overview
- Configuration process and guidelines
- Pre-defined Circuit and Remote ID formats
- Custom Circuit and Remote ID formats
- Configuration procedures
 - To enable DHCP Option 82/LDRA at the system level
 - To configure Access Identifier Profiles
 - To enable DHCP Option 82/LDRA at the VLAN level

Option 82/LDRA: Overview

The DHCPv4 L2 Relay Agent (Option 82) and Lightweight DHCPv6 Relay Agent (LDRA) are used to add information to DHCP requests that are relayed to a DHCP server to authenticate the source of the requests. See RFC 3046, RFC 6221, and RFC 3315 for more information.

Note: The E7 OS system implementing LDRA performs a link-layer bridging (i.e., non-routing) function. LDRA resides on the same IPv6 link as the client and a DHCPv6 Relay Agent or server, and is functionally the equivalent of the Layer 2 DHCP Relay Agent for DHCPv4 operation. If configuration requirements are met, both IPv4 and IPv6 are supported for data services simultaneously.

The two sub-options of Option 82 are defined in RFC 3046:

- Agent Circuit-ID (intended for circuits terminated by the system hosting the Relay agent)
- Agent Remote-ID (intended to identify the remote host end of a circuit)

The two sub-options of LDRA are defined in RFC 3315:

- Agent Interface-ID (equivalent to Option 82 Circuit-ID option in DHCPv4)
- Agent Remote-ID (equivalent to Option 82 Remote-ID option in DHCPv4)

The LDRA sub-options are derived from the DHCP Option 82 Circuit-ID and Remote-ID sub-options using the same format.

Trusted and untrusted interfaces

For the operation of Option 82/LDRA, there is a distinction between “Trusted” and “Untrusted” Ethernet interfaces:

- Untrusted interfaces are snooped for client DHCP traffic.
- Trusted interfaces are snooped for DHCP server traffic.

Note: All E7 GPON ONT Ethernet ports are implicitly Untrusted and the “Trusted” attribute cannot be configured.

DHCP message handling upstream and downstream

When Option 82/LDRA is enabled, upstream (client-to-server) DHCP messages are captured at “Untrusted” interfaces, Relay Agent identification information is inserted, and the messages are sent to the DHCP server. For inserting identifying information, the following options and capabilities are provided by the E7 OS system:

- Pre-defined Circuit and Remote ID formats (Calix-format, Calix-format-2, or TR-101-format)
- Custom Circuit and Remote ID formats
- Up to 20 pairs of Circuit and Remote ID formats may be saved as Access Identifier Profiles
- Unique Access Identifier Profiles per VLAN for both GPON and Ethernet/xDSL

Note: For complete details on the Circuit and Remote ID formats that may be used, see *Pre-defined Circuit and Remote ID formats* and *Custom Circuit and Remote ID formats*, below.

When Option 82/LDRA is enabled, downstream (server-to-client) DHCP packets are captured and examined on "Trusted" uplink interfaces:

- If a session match is found for an interface, the Option 82/LDRA string is removed if Option 82/LDRA is enabled, and then the packet is delivered to the Ethernet, xDSL, or ONT Ethernet port interface where the lease request originated.
- If a session match is not found, the DHCP packet will be forwarded unchanged on either the port for which the MAC is learned, or on all "Trusted" interfaces belonging to the VLAN (in case of broadcast DHCP packets).

All packets received on Untrusted interfaces that already have DHCP Relay Agent information will be dropped.

DHCP snooping

An important requirement for the operation of Option 82/LDRA is that DHCP Snoop or Proxy be enabled on the service VLAN. Both modes enable DHCP snooping functionality, which tracks all DHCP activity on a VLAN and maintains a table of DHCPv4 and DHCPv6 leases granted. You can retrieve the table of granted leases via any of the management interfaces containing the following attributes, which may be searched or filtered by VLAN ID, IP address, MAC address, and ONT port.

Neighbor Discovery Protocol (NDP) flood control is enabled automatically with LDRA when DHCP Option 82 insertion is enabled. NDP flood control prunes NDP messages so only IP hosts discover the access router. In order to enable visibility into the NDP flood control processing, the E-Series provides a number of counters such as all packets forwarded and discarded on a per packet type basis, and access to the NDP cached entries. This information can be accessed via CLI, EWI and CMS.

Option 82/LDRA: Configuration process

The following configuration process must be followed for Option 82/LDRA to operate:

- Enable Option 82/LDRA at the system level
- Configure the Ethernet uplink interface as trusted
- If applicable, configure subscriber Ethernet interfaces as untrusted
- Configure Access Identifier Profiles as required, doing any of the following:
 - Use the two system default profiles as-is
 - Edit one or both of the system default profiles
 - Create and save unique profiles (up to 18)

Note: Circuit and Remote ID formats must be selected during the configuration of Access Identifier Profiles. For complete details, see *Pre-defined Circuit and Remote ID formats* and *Custom Circuit and Remote ID formats*, below.

- On the S-VLAN
 - Enable Option 82

Note: If Option 82 is disabled at the system level, enabling it at the VLAN level will not have any effect. The Option 82 enable/disable control at the VLAN level is only effective if Option 82 is enabled at the system level.

- Enable DHCP Snoop/Proxy
- Select Access Identifier Profiles for Ethernet/xDSL and GPON subscribers

Note: A profile must be selected for both Ethernet/xDSL and GPON subscribers, even if only one type of subscriber is served by the VLAN. In this case, any profile may be used for the subscriber type that is not present, or a "null" profile (with Circuit ID and Remote ID = None) may be created and used for this purpose.

- Use the following VLAN service model when adding data subscribers:
 - (For Option 82) Any VLAN service model
 - (For LDRA) Only the N:1 VLAN service model

Note: IPv4 and IPv6 are supported simultaneously on the same VLAN using the N:1 VLAN service model.

Option 82/LDRA: Configuration guidelines

The following rules and guidelines apply to Option 82 and LDRA:

- DHCP Snoop/Proxy and Option 82/LDRA are not supported on the Management VLAN.
- Ethernet interfaces used for LAG or ERPS links cannot be set to "Untrusted."
- Up to 32,000 DHCP leases can be stored in the system database.
- For all E7 OS systems, the DHCP lease database persists during reboot and reset.
- For E7 OS systems with line cards (E7-2 and E7-20), the DHCP lease database also persists during card replacements:
 - DHCP leases are stored on the local card and controller card, and the standby controller syncs databases with the active controller.
 - For leases to persist across card replacements, the cards must be of the exact same type. (For example, leases will be lost if going from VDSLr1 to VDSLr2 cards, or from GPON-4 to GPON-8 cards.)

The following rules guidelines apply to LDRA only:

- LDRA cannot be enabled independently from Option 82.
- LDRA is automatically enabled on any VLAN with DHCP Snoop/Proxy enabled.
- LDRA is supported for data services on VLAN per Service models only.
- IPv6 residential gateways (RGs) are required for the VLAN per Service data model.

- Access interfaces support IPv6 transparency only (no LDRA or NDP flood control) for TLAN point-to-point, TLAN multi-point, and VLAN-per-port topologies.
- E-Series access interfaces support LDRA, but are not expected to be subscriber-facing.
- Edge interfaces do not support LDRA on received RELAY-FORWARD and RELAY-REPLY messages.
- Tag actions can be applied to untagged or single-tagged subscriber traffic processed by LDRA, including:
 - GPON ONT / ETHERNET / VDSL subscriber untagged with tag action: Add Tag
 - GPON ONT / ETHERNET / VDSL subscriber single tagged with tag action: Change Tag
 - GPON ONT / ETHERNET / VDSL subscriber single tagged with VLAN membership
 - GPON ONT subscriber untagged with tag action: Add 2 Tags
- Security features not supported for IPv6 traffic include: MAC FF, IP Source Verification, and static IPv6 host entries.

Option 82/LDRA: Pre-defined Circuit and Remote ID formats

When configuring Access Identifier Profiles, a number of options are available for specifying the format of the Circuit and Remote ID information inserted for Option 82. First of all, you can select the following pre-defined formats for content insertion:

- Ethernet and xDSL ports:
 - Circuit-ID options:
 - Calix-format: <system-ID> eth <shelf>/<slot>/<port>:<Vlan-Id>[-<Vlan-Id>]
 - TR-101-format: <system-ID> <iftype>
<shelf>/<slot>/<tr101port>:<cetag>[-<tag-Id>]
 - ♦ The TR-101 *iftype* should be either “eth” or “atm” (must be all lower case).
 - ♦ The TR-101 *cetag* should be one of 3 formats:
:vpi.vci for DSL lines/groups that are trained in ATM mode (tagged or untagged)
:ce-vlan-id for tagged subscribers that are either PTM DSL lines/groups or ONT
Null for untagged subscribers that are either PTM DSL lines/groups or ONT
 - Calix-format-2: <system-ID>:<shelf>/<slot>/<port>

Note: If the xDSL port is a member of a bonded link group, the port within the xDSL bonded link group with the lowest port value will be selected to fill the <port> field in the Circuit-ID string.

- Remote-ID options:
 - Subscriber ID of the port on which the DHCP lease request is received. The first 64 characters of the Subscriber ID text field are inserted.
 - none (no content is inserted)
- GPON ONT Ethernet ports:
 - Circuit ID options:
 - Calix-format: <system-ID> eth
<shelf>/<slot>/<port>/<OntID>/<Ontport>:<Vlan-Id>[-<Vlan-Id>]
 - TR-101-format: <system-ID> eth
<shelf>/<slot>/<port>/<OntID>/<Ontport>:<etag>[-<tag-Id>]
 - Calix-format-2: <system-ID>:<shelf>/<slot>/<port>/<OntID>/<Ontport>/
 - Remote-ID options:
 - ONT MAC ID. For DOCSIS provisioning, so that the ONT MAC is presented to the DHCP server to validate that the subscriber CPE is connected to a valid ONT virtual Cable Modem (vCM). See the Calix Open Link Cable vCMTS Command-Line Interface (CLI) Reference Guide and Calix Open Link Cable vCMTS SNMP Management Guide for more information.
 - ONT FSAN serial number, which is the default, specified in the "gpon-system-default" profile.
 - Subscriber ID of the port on which the DHCP lease request is received. For ONT VoIP hosts, the subscriber ID of the ONT is used. In both cases, the first 64 characters of the Subscriber ID text field are inserted.
 - none (no content is inserted)

Note: The default Calix format will have a defining letter for the port (x,g,v,etc) followed by the port number. The TR101 format will have a defining letter for the port followed by the port number, except for the VDSL ports which will be only the port number (no leading letter 'v').

Option 82/LDRA: Custom Circuit and Remote ID formats

You can also create custom formats for content insertion by selecting the “user” format option and using the tags in the following table. Per custom format created, these tags may be used in any order and combined with text strings/characters up to a maximum of 63 characters (for the entire custom format string).

Tag	Name	Description
%systemid	SystemID	String of maximum 20 characters defining the global unique ID of the system. It is provisioned by the administrator of the platform. The system default is an empty string "".
%iftype	IfType	Three characters automatically derived from the interface. Its value could be atm for ADSL(2+), eth for others.

Tag	Name	Description
%rack	Node/Rack	A number defined by the administrator identifying the Node/Rack number of the device. The default value is 1. (Always 1 for Calix systems.)
%chassis	Shelf	A number specified by the administrator identifying which shelf this is within a Node/Rack. The default value is 1.
%slot	Slot	Slot number where the interface is located. It is derived by the system.
%port	Port	Port number of this interface.
%ontid	ONT	(GPON only) Number of ONT. The number is locally defined
%ontport	PON	(GPON only) PON number on which the DHCP packet was received.
%cetag		For ATM legacy items, including ":VPI.VCI"
%vlan	VLAN ID	Valid VLAN number, and defines the VLAN sub-interface on which the DHCP packet was received.
%desc	DESC	31-character, free-format string associated with the interface on which the DHCP packet arrived.
%sn	SN	(GPON only) FSAN serial number of the adjacent device.
%mac	DMAC	(GPON only) MAC address of the device sending the DHCP request (most often a CPE)
%ontslot	ONTSlot	The default value is 1. (Always 1 for Calix systems.)
%calixport	Calix-format Port	This field is the port in Calix-defined format. It absorbs the different format per interface type. • For GPON, it is equivalent to "%port/%ontid/%ontport". • For VDSL and Ethernet, it is equivalent to "%port". • For ADSL, it is equivalent to "%port".
%tr101port	TR101-format Port	Port in the TR101-defined format. Since TR101 only defines the format for non-GPON interfaces, for GPON interfaces, the TR101-defined format falls back to Calix-defined format.
%%	%	For the '%' character

For reference, the following custom format strings correspond to the pre-defined formats that are currently available.

Pre-Defined Formats	Equivalent Tag Strings
calix-format	
Ethernet/xDSL	%systemid Eth %chassis/%slot/%port:%vlan
GPON	%systemid Eth %chassis/%slot/%port/%ontid/%ontport:%vlan
Either (using calixport)	%systemid Eth %chassis/%slot/%calixport:%vlan
tr101-format	
Ethernet/xDSL	%systemid %iftype %chassis/%slot/%tr101port%cetag
GPON	%systemid %iftype %chassis/%slot/%tr101port%cetag (Note: If used for GPON, falls back to calix-format.)
Either (using calixport)	%systemid %iftype %chassis/%slot/%calixport%cetag
calix-format-2	
Ethernet/xDSL	%systemid:%chassis/%slot/%port
GPON	%systemid:%chassis/%slot/%port/%ontid/%ontport
Either (using calixport)	%systemid:%chassis/%slot/%calixport
subscriber-id format	%desc
fsan-serial-number format	%sn
mac-addr format	%mac

Option 82/LDRA: Configuration procedures

The following steps of the configuration process are described in detail below

- To enable DHCP Option 82/LDRA at the system level
- To configure Access Identifier Profiles
 - Configure system default profiles
 - Create custom profiles

To enable DHCP Option 82/LDRA at the VLAN level

(Optional) Configuring DHCPv4 Proxy


Note: The configuration of this feature is optional to increase security with DHCPv4 messaging.

Overview

DHCPv4 Proxy, also referred to as Layer 3 DHCP Relay, is a security feature that prevents Denial of Service (DOS) attacks. When this feature is configured, the E-Series system will act as a proxy between DHCP servers and clients.

- When the Proxy detects client DHCP packets, it will capture and send them for DHCP Proxy processing. It will forward client DHCP packets upstream, replacing the source and destination IPs and MACs, as well as the VLAN with values obtained from the dhcp-svc-prof and vlan-ip-host objects.
- Conversely, downstream packets from the server will use the E7 Proxy as the destination IP and MAC. The E7 Proxy will switch packets to the client VLAN, changing the headers to either broadcast or unicast to the client.

- The client will be isolated from the DHCP server, not able to ping the server or the gateway.



Configuration rules and guidelines

- Two separate VLANs are required for this configuration—a server-side VLAN and a client-side VLAN.
- The server-side VLAN may be a dedicated VLAN or the management VLAN.
 - If a dedicated VLAN is used, a VLAN IP Host object is required to configure the DHCP Proxy Agent IP interface.
 - If the management VLAN is used, the In-Band Management Interface (IP Address, Mask, Management VLAN) and System IP Gateway are used to configure the DHCP Proxy Agent IP interface.

Note: If the DHCP server is on the management VLAN, Calix recommends using the management VLAN for the server-side configuration.

- A DHCP Proxy Service Profile is required for association with the client-side data VLAN.
- The VLAN parameter "DHCP Mode" allows provisioning of DHCP proxy (proxy), DHCP snooping (snoop), or no DHCP snooping (none). The server-side uplink VLAN must have DHCP Mode = none, and the client-side VLAN must have DHCP Mode = proxy.

Note: In DHCP Proxy mode, DHCPv4 snooping, DHCPv4 Option 82 operations, DHCPv6 snooping, and DHCPv6 LDRA operations are performed just as in DHCP Snoop mode. The only difference is that DHCPv4 messaging traffic is processed by the Proxy Agent and sent upstream via the server-side S-VLAN. All other IPv4 traffic is sent upstream via the standard data VLAN, and IPv6 traffic is not impacted at all.

- For the server-side VLAN:
 - ID = available ID, different than client-side VLAN
 - DHCP Mode = None
 - (all other parameters as usual)

Note: This VLAN must be configured on all nodes in a ring.

- For the client-side VLAN:
 - ID = available ID, different than server-side VLAN
 - DHCP Mode = Proxy
 - DHCP Proxy Service = <name of profile>
 - (all other parameters as usual)
- You must follow the provisioning steps described in *Provisioning*, below.

Parameters

The relevant objects and VLAN parameters for this feature are described below.

VLAN IP Host

Note: If using a dedicated VLAN as the server-side data VLAN.

- Embedded Web Interface:
 - On the Navigation Tree, click **VLANs**; in the work area, click **VLAN IP Hosts**; on the toolbar, click **Create**.
- CLI:
 - `create vlan-ip-host <name> [vlan|ip|netmask|gateway]`

Parameter	Description	Valid Options
Name*	Name of profile	up to 32 character string
IP Address*	IP address for the DHCP Proxy Agent. This should be in the same subnet as the upstream edge router.	valid IP address
Netmask*	Subnet mask required for the network	valid netmask
VLAN*	Server-side data VLAN ID.	valid VLAN ID
Gateway	Next hop gateway. For example, the IP address of the upstream edge router.	valid IP address

*Required field

Note: If using the management VLAN as the server-side data VLAN, configure the In-Band Management Interface parameters as usual (IP Address, Mask, Management VLAN) and configure the System IP Gateway as usual (System > Provisioning > IP Gateway).

DHCP Proxy Service Profile

- Embedded Web Interface:
 - On the Navigation Tree, click **VLANs**; in the work area, click **DHCP Proxy Service**; on the toolbar, click **Create**.
- CLI:
 - `create dhcp-svc-prof <name> [vlan-ip-host|svr-ip-1|svr-ip-2|svr-ip-3|svr-ip-4|svr-ip-5]`

Parameter	Description	Valid Options
Name*	Name of profile	up to 32 character string
VLAN IP Host*	Two options for configuring the DHCP Proxy Agent IP interface: <name of VLAN IP Host>: Use information from the specified VLAN IP Host object. Use Management Config: Use information from the In-Band Management Interface and System IP Gateway.	valid VLAN IP host name Use Management Config
IP Addr of Server 1*	IP address of DHCP server.	valid IP address
IP Addr of Server 2-5	IP addresses of additional DHCP servers, if required. If more than one DHCP server IP address is entered, the proxy agent will send a unicast DHCP message to each non-zero DHCP IP address. It is up to the server to respond and it is up to the client to select the DHCP reply it will use.	valid IP address

*Required field

VLAN (relevant parameters)

The relevant VLAN parameters should be provisioned as shown below:


- Server-side VLAN:
 - DHCP Mode = None
- Client-side VLAN:
 - DHCP Mode = Proxy
 - DHCP Proxy Service = <name of profile>

Provisioning

Two provisioning models are provided for DHCPv4 Proxy—using a dedicated VLAN or using the management VLAN for the server-side data VLAN.

DHCP proxy provisioning order, using a dedicated server-side VLAN:


- 1.** Create a server-side VLAN with `dhcp-mode = none`.
- 2.** Create a `vlan-ip-host` that references the server-side VLAN.
- 3.** Create a `dhcp-svc-prof` that references the `vlan-ip-host`.
- 4.** Create the client-side VLAN with `dhcp-mode = proxy` and reference to the `dhcp-svc-prof`.
- 5.** Assign the server-side VLAN to the uplink interface for DHCPv4 messaging.
- 6.** Assign the client-side VLAN to the uplink interface for all other data traffic.
- 7.** Assign the client-side VLAN to clients (via the usual subscriber service provisioning steps).



DHCP proxy provisioning order, using the management VLAN:


1. Create the management VLAN, In-Band Management Interface (mgmt-cfg), and System IP Gateway.
2. Create a dhcp-svc-prof that references the mgmt-cfg.
3. Create the client-side VLAN with dhcp-mode = proxy and reference to the dhcp-svc-prof.
4. Assign the server-side VLAN (management VLAN) to the uplink interface for DHCPv4 messaging.
5. Assign the client-side VLAN to the uplink interface for all other data traffic.

6. Assign the client-side VLAN to clients (via the usual subscriber service provisioning steps).




Comparison of DHCP modes


In DHCP Proxy mode, DHCPv4 snooping, DHCPv4 Option 82 operations, DHCPv6 snooping, and DHCPv6 LDRA operations are performed just as in DHCP Snoop mode. The only difference is that DHCPv4 messaging traffic is processed by the Proxy Agent and sent upstream via the server-side S-VLAN. All other IPv4 traffic is sent upstream via the standard data VLAN, and IPv6 traffic is not impacted at all.



The following diagram illustrates that all IPv4 traffic is carried over the standard data VLAN when DHCP Mode = Snoop. Only when DHCP Mode = Proxy does DHCPv4 traffic get handled differently.



The following diagram illustrates that IPv6 traffic is not impacted by DHCP Mode. In either mode (Proxy or Snoop), all IPv6 traffic is carried over the standard data VLAN.



IPv6 Support

For data services (high speed internet), E-Series systems support dual-stack access networks with both IPv4 and native IPv6 subscribers. Going beyond Layer 2 transparency for IPv6 (with TLAN service models), E-Series systems provide the following IPv6 support features:

- DHCPv6 Snooping and Lightweight DHCPv6 Relay Agent (LDRA)
 - DHCPv6 snooping maintains a local lease database for IPv6 hosts
 - LDRA inserts Interface ID / Remote ID DHCPv6 option
 - DHCPv6 snooping with LDRA identifies IP and MAC address for each access port
- Neighbor Discovery Protocol (NDP) Flood Control
 - Access platforms prune NDP messages so IP hosts only discover the access router
- Calix P-Series, GigaCenter, and GigaHub Home Gateways (HGWs) provide IPv6 support for data services.

Lightweight DHCPv6 Relay Agent (LDRA)

Feature Description

Lightweight DHCPv6 Relay Agent (LDRA) is used to insert relay agent options (Interface ID/ Remote ID) in the DHCPv6 message exchanges identifying client-facing interfaces. Lightweight DHCPv6 Relay Agent (LDRA) as defined in (RFC 6221) defines a simple layer-2 relay function which intercepts DHCPv6 messages, inserts the Interface Option, and then relays them to the server as per RFC 3315 (DHCPv6).

The E-Series implementation of LDRA performs a link-layer bridging (i.e., non-routing) function. LDRA resides on the same IPv6 link as the client and a DHCPv6 Relay Agent or server, and is functionally the equivalent of the Layer 2 DHCP Relay Agent for DHCPv4 operation in Layer-2 Relay (L2RA). Both IPv4 and IPv6 are supported simultaneously on Subscriber per Service N:1 VLAN models in all layer-2 topologies including rings, ring-on rings and single node.

By inserting port location into the DHCPv6 request (Interface-ID), LDRA provides a means to assign addresses, nameservers, etc as in IPv4. This provides a similar function as DHCPv4 option 82 for tracking customers. When used in conjunction with a Residential Gateway (RG) at the subscriber premises the user gains the ability to locally provide addresses for premises devices and ensure that traffic destined between local devices remains local. Use of a RG is always recommended to remove the need to have subscriber premises devices requesting addresses from upstream routers via a local NAT function. When a service provider chooses to deploy in a 1:1 VLAN model, a residential gateway is not required for subscriber isolation, as it is isolated via the its own VLAN and all its traffic is sent upstream to the access router.

The table below shows terminology differences between IPv6 LDRA and IPv4 L2 DHCP RA.

LDRAv6 (DHCPv6)	L2 RA (DHCPv4)
DHCPv6 Opt18 Interface ID	DHCPv4 Opt82suboption 1
DHCPv6 Opt37 Remote ID field	DHCPv4 Opt82suboption2

DHCP snooping lease table

When DHCP snooping is enabled on a VLAN, the E-Series will track all DHCP activity on that VLAN within the system and maintain a table of DHCP leases granted. The table of granted leases is retrievable via any of the management interfaces and contains the following attributes, which may be searched or filtered by: VLAN ID, IP address, MAC address, ONT Port.

The DHCPv6 lease info will be visible in the in the CLI, WebGui, and CMS in Release 2.2 in the same DHCP lease table where v4 addresses are contained.

In Release 2.2, configuring DHCP-SNOOP with option 82 enabled on a VLAN will automatically enable LDRA.

Application and Benefits

Lightweight DHCPv6 Relay Agent (LDRA) is simply used to insert relay agent options (Interface ID / Remote ID) in the DHCPv6 message exchanges identifying client-facing interfaces. This allows subscriber identification in Layer-3 aware networks and enables data services over a VLAN per Service N:1 model.

Configuration Rules

Supported on all E7-20 1GE, 2.5GE, 10GE interfaces	Yes
Supported for E7 GPON ONTs	Yes
Supported on all E7-2 1GE, 2.5GE, 10GE interfaces	Yes
Supported on E7-2 VDSL2 interfaces	Yes
LDRA and NDP snoop are automatically enabled on any VLAN with DHCP enabled	Yes

Please note the following system behavior and limitations for IPv6 support on E7 in Release 2.2:

- IPv6 functionality is supported for high speed internet services only
 - Video, VoIP, or System Management services are not supported for IPv6 in Release 2.2 .
- Trunk interfaces
 - Support only IPv6 transparency (no LDRA encapsulation/decapsulation or NDP Snoop support)

- Edge interfaces
 - No LDRA encapsulation/decapsulation performed on received RELAY-FORWARD / RELAY-REPLY messages
 - NDP Flood control enabled when in untrusted mode
- Access interfaces
 - Considered customer facing and may only be set to untrusted
 - Includes GPON ONT Ethernet interfaces, VDSL2Release 2.2 Access interfaces, and Access Ethernet interfaces
 - 10GE Access interfaces are not precluded from support of LDRA but is not expected to be client facing
 - IPv6 Transparency support for the following topologies: TLAN pt-2-pt (EVP-LINE), TLAN multipoint (EVP-LAN), and 1:1 services
- VLAN per Service data model
 - In this service model, IPv6 residential gateways (RGs) are required.
- VLAN per Port (single and double-tagged) data model
 - IPv6 residential gateways (RGs) are recommended, but not required.
- Tag actions
 - Can be applied to untagged or single tagged subscriber traffic processed by LDRA. Tag criteria:
 - GPON ONT / ENET / VDSL Subscriber untagged with Tag Action : add-tag
 - GPON ONT / ENET / VDSL Subscriber single tagged with Tag Action: change-tag
 - GPON ONT / ENET / VDSL Subscriber single tagged with VLAN membership
 - GPON ONT Subscriber untagged with Tag action Add-2-tags
- Security features not supported for IPv6 traffic
 - MAC Forced Forwarding, IP Source Verify
 - Static IPv6 Host and Subnet Snoop table entries not supported
 - IPv4 Relay and LDRA are not independently enabled / disabled
 - IPv6 lease snooping supported through DHCPv6 only (not NDP)
 - Interface ID and Remote ID values are not user defined. These values are derived from DHCP Opt82 sub option 1 Circuit ID and sub-option 2 Remote ID using the same format

Neighbor Discovery Protocol (NDP) Flood Control

Release 2.2 includes support for NDP flood control. The primary purpose of NDP flood control is to prune NDP messages so IP hosts only discover the access router.

NDP flood control is enabled automatically with LDRA when DHCP Option 82 insertion is enabled. In order to enable visibility into the NDP flood control processing, the E7 provides a number of counters such as all packets forwarded and discarded on a per packet type basis, and access to the NDP cached entries. This information can be accessed via CLI, Web GUI and CMS.

Home Gateway IPv6 Support

Calix P-Series, GigaCenter, and GigaHub Home Gateways (HGWs) provide the following IPv6 support for data services:

- Single-Stack IPv6
- Native Dual-Stack IPv4/IPv6
- Dual-Stack Lite (DS-Lite): IPv4 traffic encapsulated in IPv6 packets
- IPv6 Rapid Deployment (6rd): IPv6 traffic encapsulated in IPv4 packets

Note: In 6rd mode, the maximum throughput is limited to 100 Mbps.

Note: IPv6 modes are only configurable with the HGW RG interface in the external management mode.

Note: For detailed information on how to configure these IPv6 support modes, see the P-Series and GigaCenter documentation.

Calix T-Series HGU's provide the following IPv6 support:

- Single-Stack IPv6
- Native Dual-Stack IPv4/IPv6

Note: For detailed information on how to configure these IPv6 support modes, see the T-Series documentation.

MEF Compliant Services


Metro Ethernet Forum (MEF) Private Line (EPL) and Transport LAN (TLAN) business service models can transparently trunk business traffic across a network to other locations, typically a remote office or secondary business location. The traffic received from the business may be untagged, single-tagged, or double-tagged. The E-Series adds an outer tag to all frames to create a private switched LAN with two or more end points.

Transparent LAN (TLS) Services

The E7 support flexible Transparent LAN Service (TLS) architecture to deliver business services. Metro Ethernet Forum (MEF) compliant services (MEF 9 and 14) are supported over point-to-point Ethernet and E7 GPON. All E7 cards support up to 4,092 Ethernet Virtual Connections (EVCs) and can be configured for E-Line (point-to-point) and E-LAN (multipoint-to-multipoint) transparent LAN services over both GPON and GE networks.

The TLS services can transparently trunk business traffic across a network to other locations, typically a remote office or secondary business location. The traffic received from the business may be untagged, single tagged, or double tagged. The E7 adds an outer tag to all frames, creating a private switched LAN with two or more end points.

The diagrams below show examples of an E7 TLAN business service delivery configuration.



TLAN Business Service Delivery

When the E7 is utilized for TLS services, it supports the following VLAN capabilities:

- Single VLAN tag addition to untagged, single tagged or double tagged incoming frames
- VLAN/MAC switching of TLS traffic from any designated TLS interface to any other TLS interface
- Transparent transport of all TLS traffic
- Peer, discard, or tunnel of RSTP BPDU and Link Aggregation control plane traffic

Ethernet Private Line (EPL) Services

The E-Series platform provides flexible, high-capacity Ethernet business services in addition to traditional residential triple play services. Metro Ethernet Forum (MEF) compliant services (MEF 9 and 14) are supported over GPON and point-to-point Ethernet, where Ethernet Private Line (EPL), Ethernet Virtual Private Line (EVPL), and Ethernet-LAN (ELAN) services are supported.

The EPL and EVPL are both point-to-point services within the E-Series transport network, differing only in whether the entire subscriber interface is transported intact (EPL) or selectively based on subscriber VLANs (EVPL).


Because both EPL and EVPL services are point-to-point services, the E-Series does not need to learn MAC addresses for any of the subscriber host devices. If the subscriber has a simple Layer 2 Ethernet switch at one of their office locations, many 10's, 100's, or 1000's of MAC addresses may be needlessly learned by the E-Series during the transport of the EPL or EVPL service.



E7-2 with Ethernet Private Line (E-Line) Service

The operator can disable MAC learning on a per-VLAN basis. For a point-to-point EPL or EVPL services, this results in service delivery without the E-Series adding any subscriber MAC addresses to the E-Series forwarding table. With a MAC table of 32K addresses, learning subscriber MAC addresses on an EPL or EVPL service is not onerous in most cases.

In the diagram below, the EPL VLAN has MAC address learning disabled. Traffic on the subscriber's interface (UNI) is flooded around the ERPS or RSTP ring in both directions. The inherent interface block of the ERPS or RSTP protocol that prevents a forwarding loop also ensures that only a single copy of the subscriber traffic arrives at the far end of the EPL service UNI.



EPL with MAC Address Learning Disabled

Disable MAC Address Learning Planning Guidelines

- MAC Address Learning is disabled/enabled on a per VLAN basis.
- Disabling MAC Address Learning on a VLAN causes the traffic within the VLAN to be flooded to all egress interfaces with membership in the VLAN.
- Disabling MAC addresses must only be used with point-to-point services and is NOT appropriate for a multipoint Ethernet-LAN (ELAN) service or other VLAN-per-Service applications such as internet access.

Data Services

This section describes E-Series data services.

About Data Services

The Calix E-Series supports data service applications over copper-based (xDSL), and for E7, fiber-based (GPON) access networks. Data services can be deployed from any E-series node. That is, the nodes providing data subscriber access (xDSL or OLT) can reside in either the Central Office or in Remote Terminals, with the exception of E7-20 which is only supported in Central Office environments.

The E-Series supports Ethernet trunk interfaces to uplink data traffic to the WAN. Layer 3 Ethernet services such as static routing are also supported. The following data service models are supported over IP on GPON ONTs and VDSL2 ports.

- **VLAN-per-Service:** where each service is assigned a dedicated VLAN where multiple subscriber ports are assigned to the VLAN for a single service. This model is often referred to as N:1. An additional tag can be added to limit the size of a group for improved reliability. A service carried on an N:1 VLAN applies to multiple subscriber ports, allowing a single match list and tag action to be used to denote the service.

Note: If using a N:1 VLAN model for services, Calix recommends enabling MAC Forced Forwarding (MAC FF) and IP Source Verify when creating the VLAN, where appropriate.

- **VLAN-per-Port:** (Calix recommended model) where each subscriber port is assigned a dedicated VLAN. This model is often referred to as 1:1. An additional tag can also be added to group subscriber ports in logical categories. A service carried on a 1:1 VLAN is the same for each subscriber, except the customer tag is unique per subscriber. You can define the match list and tag action pair such that multiple subscriber ports can reference it. This is accomplished by assigning an attribute for the output tag in the tag action that indicates the value of the output tag is subscriber specific. The customer specific tag is included in the service definition when provisioning the service.
- **VLAN Stacking, Q-in-Q:** where a transparent LAN service on each port is configured using VLAN double tagging.

High-speed internet access through GPON

The E7 GPON and point-to-point optical Ethernet access solutions to deliver up to 1 Gbps data connection between Internet service providers and subscribers. Service profiles in the E7 and the ONTs can provide tiered service offerings in 1 Mbps increments up to 1 Gbps. The E7 provides secure internet access using either of the Broadband Forum's TR-101/TR-156 recommended 1:1 VLAN-per-subscriber and N:1 VLAN-per-service models.

High-Speed Internet Access through xDSL

The E7-2/E3-48C and 3rd party modems provide data interconnection between Internet service providers and subscribers. The method of transport over the copper loop is Packet (Ethernet) or Asynchronous Transfer Modes (ATM). PTM can be supported by VDSL2, ADSL2+ and ADSL2 modes, while ATM can only be supported in ADSL2+ / ADSL2 / ADSL / G.dmt modes. Each mode must also be supported by the specific 3rd party modem used. Service profiles provide tiered service offerings. Different models of 3rd party modems provide different types and numbers of Fast Ethernet (100 Mbps) and/or Gigabit Ethernet (1000 Mbps) ports, USB ports, WIFI interfaces, etc.

For information on third party modem interoperability, see *xDSL Modems* (on page [23](#)).

Configuration Process

The high-level process for turning up data service is described below.

Starting point

Before starting the configuration process, check that the following conditions are met:

- Transport is established to all nodes providing broadband data service.
- The uplink card is installed and wired.
- The xDSL and OLT service cards are installed and wired.
- (GPON applications only) The ONTs are installed, wired, turned up and ranged.

Configuration process overview

The data service configuration process follows:


1. Verify that all required equipment is installed and wired.
2. Configure profiles and templates.
 - Rules to classify traffic with service match list and service tag action
 - Ethernet bandwidth profiles to define your different data service offerings
 - Quality of Service (QoS) actions to perform on ingress traffic packets when the packets match selection criteria with policy maps
 - Optional:
 - Enable and/or specify security attributes of the port using an Ethernet security profile
 - Create PPPoE profile that is referenced by the data service VLAN
 - (xDSL only) Create xDSL templates to define port settings.
 - (xDSL only, if required) Configure bonding groups.

3. Configure data subscriber services.
 - a. Configure the xDSL or GPON port.
 - b. (xDSL only) Configure the xDSL port associated interface.
 - c. Add the data service to the subscriber port.

See *Calix E7 GPON Applications Guide* and *Calix E7-2/E3-48C xDSL Applications Guide* for configuration details.

PPPoE Subscriber Authentication and Service Activation

The software subscriber authentication and service activation using Point-to-Point Protocol over Ethernet (PPPoE) can be applied to any end point regardless of access media, xDSL or GPON.



A data service VLAN is set with a PPPoE profile in order to allow this method of subscriber authentication from a device connected to a subscriber. Setting a VLAN PPPoE profile to “none” passes through all PPPoE traffic, transparently. If a PPPoE profile is used with PPPoE snoop, a list of all the active sessions and statistics are available, and the PPPoE stack is enabled, which passes through PPPoE traffic transparently as long as the Clients/BRAS are operating normally (illegal packets will be dropped).

The uplink toward the router, BRAS or PPPoE Server must be on an node interface set to the mode of "Trunk."

PPPoE Relay and Intermediate Agent

The VDSL2 line cards and GPON ONT Ethernet ports support Point-to-Point Over ATM to Point-to-Point Over Ethernet (PPPoA-to-PPPoE) conversion, and PPPoE intermediate agent functionality.

The service can use either single tagged or double tagged (Q-in-Q) frames, to and from the uplink ports, depending on the customer configuration.

When a single VLAN tag is required, a PVID value is assigned to what is considered the Service Provider Tag (S-Tag). When the service model requires Double Tagged traffic (Q-in-Q), an optional inner or Customer Tag (C-Tag) can be provisioned for traffic received for the customer or CPE initiated PPP session. Prioritization bit values can be assigned for both VLAN tags.

When PPPoX intermediate agent function is enabled in the VLAN, the system automatically detects all PPPoE Active Discovery packets, including PADI/PADO/PADR/PADS/PADT, and builds a correspondence table between user provisioned S-Tag / C-Tag PVIDs, and PPP Session ID, for the given DSL port/group. PPPoA LCP (Link Control Protocol) messages are automatically detected so system can initiate a corresponding PPPoE session.

The intermediate agent functionality allows for a PPPoE session access method where the access node inserts a Access Loop Identification (ALID) tag to the PPPoE frames to be sent upstream toward the Broadband Network Gateway (BNG), and then tracks the PPPoE sessions. The inserted PPPoE tag contains the identification of the access loop on which the PADI or PADR packet was received in the access node where the intermediate agent resides. The system complies to TR-101 R-124 requirements for the Agent Circuit ID inserted by the Access Node PPPoE Intermediate Agent (IA) Vendor specific Tag by applying a system-defined access-identifier profile.

Note: The default Calix format will have a defining letter for the port (x,g,v,etc) followed by the port number. The TR101 format will have a defining letter for the port followed by the port number, except for the VDSL ports which will be only the port number (no leading letter 'v').

- The uplink toward the router or PPPoE Server must be on an interface set to the mode of "Trunk."
- At the aggregating Trunk interface northbound of the E-Series, the PPPoE traffic must adhere to a specific VLAN membership on the Trunk interface.
- The PPPoE must be configured so that tag actions are on the Access ports, facing the subscriber modem.
- A modem with a common preprogrammed VLAN can be changed at the xDSL or ONT port interface to either a common service VLAN, or to a 1:1 VLAN associated with that customer.
- Sessions may be manually terminated.
- The VDSL2 card supports PPPoA on the client side, in ATM mode.
- PPPoA client side frames are converted to PPPoE and a circuit to session map is created within the system to keep track of the session to port assignment; downstream frames are converted back to PPPoA.
- Only a single PPPoA to PPPoE conversion can be maintained per xDSL or ONT port at once.
- The PPPoE profile must be set to Auto for the InterWorking function (PPPoA-to-PPPoE conversion).

- Only one PPPoE session is supported per xDSL or ONT port. If a second new session is initiated, the original session tears down and the new session is allowed.
- The PPPoE profile is applied to a service VLAN where simultaneous operation of DHCP Snooping and PPPoE are not supported.
- Setting a VLAN PPPoE profile to “none” passes through all PPPoE traffic, transparently. If a PPPoE profile is used with PPPoE snoop, a list of all the active sessions and statistics are available, and the PPPoE stack is enabled, which passes through PPPoE traffic transparently as long as the Clients/BRAS are operating normally (illegal packets will be dropped).
- PPPoE Profiles are not needed for service, but, they are recommended due to the benefits they provide:
 - A PPPoA-to-PPPoE Interworking Function that gets enabled if you have any legacy ADSL lines using PPPoA. This is legacy support for DSL deployments where the CPE does not provide the Ethernet framing and writes the PPP headers directly in the ATM payload. There can be conversion of ATM VP/VC to Ethernet tags as needed. With VDSL or ADSL PTM mode, everything is PPPoE already. Also, with ATM-based ADSL, most CPE use PPPoE with RFC 1483 bridging natively.
 - PPPoE Intermediate Agent (IA) which is defined in the TR-101 Specification of the Broadband Forum. With PPPoE IA, the E7 actively snoops the PPPoE session handshake. This allows the E7 to keep a table of session statistics that can be useful for troubleshooting. Depending on the mode (“relay” or “snoop”) the E7 can also insert a Circuit and Remote ID in the PPP packets similar to DHCP Option 82.
 - Security agent with awareness of which side has the BRAS and which side has the clients, dropping unexpected packets types in case there is an attempt to spoof the clients. DHCP filtering is also enabled on the VLAN for added security. You can optionally program the actual MACs of the BRAS devices for more security as well. Up to four BNGs are supported.
- Each xDSL (port or bonded group) or GPON subscriber can only have up to 1 Ethernet service with PPPoE enabled.

Parameters

You can provision the following parameters for PPPoE profiles:

Parameter	Description	Valid Options
Name*	Descriptive name for the PPPoE profile.	text string
Mode	Auto – Automatically detects the mode between native PPPoE clients and PPPoA clients, and runs relay (PPPoE IA) for PPPoE clients, and runs IWF for PPPoA clients. Relay – (Works with PPPoE clients only.) Inserts circuit-id and remote-id on the upstream direction, and then removes the circuit-id and remote-id on the downstream direction. Snoop – (Works with PPPoE clients only.) Snoops PPPoE packets without modifying the packets (i.e. does not insert/remove circuit-id and remote-id).	auto ‡, relay, snoop
Stale Timeout(s)	Inactivity time-out for session teardown.	10-300 60 ‡
Discovery Timeout (s)	Discovery timeout (seconds).	1-30 3 ‡
Allowed BNG MAC 1	Allowed (trusted) Broadband Network Gateway (BNG defined in TR-101) MAC address or OUI. Alternately, the keyword "none" indicates that all BNG MACs are allowed.	six hexadecimal digits in the range 0-FF, optionally separated by colons none ‡
Allowed BNG MAC 2	Allowed (trusted) Broadband Network Gateway (BNG defined in TR-101) MAC address or OUI. Alternately, the keyword "none" indicates that all BNG MACs are allowed.	six hexadecimal digits in the range 0-FF, optionally separated by colons none ‡
Allowed BNG MAC 3	Allowed (trusted) Broadband Network Gateway (BNG defined in TR-101) MAC address or OUI. Alternately, the keyword "none" indicates that all BNG MACs are allowed.	six hexadecimal digits in the range 0-FF, optionally separated by colons none ‡
Allowed BNG MAC 4	Allowed (trusted) Broadband Network Gateway (BNG defined in TR-101) MAC address or OUI. Alternately, the keyword "none" indicates that all BNG MACs are allowed.	six hexadecimal digits in the range 0-FF, optionally separated by colons none ‡

* Required

‡ Default

•

The VC encapsulation type used by CPE (LLC vs. VC Mux) is also automatically detected by the E7. There are three protocol encapsulations expected; LLC / PPPoE, LLC / PPPoA, VC-MUX / PPPoA:

- RFC2684 encapsulated Ethernet frames (LLC=0xAA-AA-03)
- LLC Encapsulated PPPoA frame (LLC=0xFE-FE-03)
- VC-Mux mode. No encapsulation present. There can be one and only one PPPoA connection per VC-Mux VC.

Transparent LAN Services Provisioning Model

E7 ports can be configured to support transparent LAN packet forwarding rules. In such a configuration, all traffic is forwarded based on the TLAN tag and all Edge ports on a TLAN are considered 'trusted' interfaces that can be flooded. In addition, the E7 allows you to specify tag actions to define how customer VLANs map to E7 service VLANs.

Consider how traffic is sent and received in this service model: The traffic received from the TLS subscriber (1, below) may be untagged, single tagged, double tagged, or greater as long as the frame size (including the TLAN tag) does not exceed 1550 bytes. Upon receiving the traffic, the E7 pushes a TLAN Tag with associated ID onto the frame and this TLAN Tag ID is used along with the Destination MAC address for all forwarding decisions. When the E7 sends the traffic to a destination interface, the E7 strips the TLAN Tag from the frame and sends the frame un-modified (as it was received) to the destination.

The one exception where a frame is not sent out the E7 as it was received is for Priority-Tagged traffic. In this case, the Priority bit is stripped and the frame is sent to the destination Untagged.

Note: If the destination is unknown, the frame will be broadcast to all interfaces in the TLAN.

Video Service

This section describes E-Series video services.

About Video Services

The E-Series supports both broadcast IPTV and Video on Demand (VOD) video services. On Demand video streams are simply unicast traffic that is no different than any other data service being delivered by the E-Series. VOD streams are merged with multicast IPTV at the GPON or xDSL interface, and transported to the residential gateway (RGW) or modem.

The E-Series supports broadcast IPTV service using IGMPv2 / v3 snooping (RFC 4541). The E-Series sits between the Layer 3 IGMP-enabled multicast router and the modem / RGW and optimizes the channel bandwidth and IGMP join / leave requests to provide a high quality video experience with minimal network resources. IGMP snooping occurs internally to the E-Series and it is leverage by it for efficient networking within the ERPS transport protocol.

The following describes the flow of IPTV video through the E-Series Ethernet network:

- Within the E-Series, only a single copy of a broadcast channel watched by multiple subscribers is placed on the transport. The subscriber ports of users watching the channel will pick up the channel from the E-Series.
- If multiple users on the same GPON ONT or xDSL line card are watching the same stream, the OLT or xDSL line card provides the channel replication function.


The E-Series supports up to 800 IGMP multicast groups, enabling up to 800 independent broadcast video and music channels.

At each video replication point in the system (E7-20 SCP, GPON OLT service access card, and ONT, or xDSL service access card) the hardware-based forwarding / replication engine intercepts multicast control traffic and processes it locally. In this manner, for example, the E7-20 SCP only replicates traffic to the 20 line cards, not the 1,000s of ONTs associated with the system. Likewise, the GPON or xDSL service access cards only replicate multicast traffic to the ports on the card and in the case of the ONT, there is only a need to replicate traffic to the 1-8 Ethernet ports, depending on the models of ONTs deployed. Together, the control processing and replication are distributed within the system and do not suffer degradation even when a full complement of service access cards are placed in the E7 chassis. A similar, yet smaller scale distribution architecture exists in the E7-2 modular chassis.

The E7-20 system has been tested to 350 messages/second on each E7-20 card. The system is rate limited to 350 messages to protect the card CPU and ultimately the multicast router upstream of the E7-20. When a message limit is reached, the subscriber may experience a failure to change the channel and need to press the channel change button a second time.

Note: The E7 also supports RF video overlay with certain Calix ONTs. See the *Calix E7 GPON Applications Guide* for more information.

Video services related objects:



VLAN with IGMP Snooping and multicast IP addresses

The E-Series supports industry standard IETF RFC 4541 IGMP snooping of multicast video leave and join requests sent between the set-top box and the video distribution network.

IGMP snooping enables the E-Series to determine which ports should be a member of a multicast group. It can provide a local response if the channel already exists on the E-Series or it will forward those requests upstream to a multicast router or another IGMP snooping Ethernet switch. When multicast channels are received from the network, the E-Series forwards the channels to all ports that are currently members of the requested IGMP membership group. The ability to perform snooping can be enabled on a per VLAN basis. If multicast traffic is received on a VLAN where IGMP snooping is not enabled (flood), the traffic is handled as broadcast traffic and sent to all ports on the VLAN.

- Passive IGMP snooping (snoop-suppress), augmented with report suppression reduces the number of general query and group specific reports sent to the multicast router (querier).

Note: The suppress-snoop parameter only applies to standalone E7-2 and Modular Chassis systems, as the E7-20 only supports proxy.

- Active IGMP Snooping (IGMP Proxy), becomes capable of acting as a general querier and thus takes an active role in maintaining the IGMP multicast network. When the IGMP mode is set to proxy, an IGMP profile must be referenced from the VLAN. In addition to general querier timing parameters, the following IGMP attributes per multicast VLAN are supported:
 - Robustness parameter (number of times E-Series sends out a general query)
 - Last member query interval
 - Immediate leave option
 - Static multicast router interface
 - Router learn mode

IGMP Proxy provides additional scalability to the IGMP network, reducing the IGMP signaling load on the multicast router.

Video Network

The video network for IPTV can be simplified by thinking of the network as divided into three parts:

- **Video headend** – The Video Headend (VHE) provides the source video content and consists of video encoders, middleware/application servers, routers, and other equipment.
- **Transport network** – The transport portion of the network distributes broadcast video traffic from the headend shelf to all E7 nodes in the network that provide video service.
- **Access network** – The access portion of the network delivers video and data services to subscribers. Access shelves contain the line interfaces (xDSL or OLT cards) that provide the subscriber service drops. E7 shelves providing access may also comprise the transport network, and be located in the central office or distributed in outdoor cabinets or Controlled Environmental Vaults (CEVs).

Video VLAN configuration guidelines

- IGMP snooping is provisioned on a VLAN basis and is recommended for VLANs carrying video services and should be avoided for VLANs in a Transparent LAN Service (TLS).
- IGMP snooping is only enabled on the outer VLAN ID as multicast traffic is typically transported throughout the network with a single VLAN ID tag.
- All nodes in an ERPS or G.8032v2 ring must have the same IGMP Snooping provisioning on the video VLAN for traffic to flow--either all with snooping (snoop-suppress, proxy) or all without snooping (flood).

- The IGMP mode setting affects VLANs as follows:
 - Snoop-suppress mode only uses IGMPv2.
 - Proxy mode references an IGMP profile that specifies the IGMP version setting. You can edit the system-default IGMP profile or create a custom IGMP profile. Other VLANs that use the same IGMP profile are not affected by the version setting unless they are in proxy mode.
 - Snoop-Suppress or Proxy mode can reference a P-bit value for IGMP traffic that passes through the system, allowing the traffic type to be treated differently as it passes through the network. The P-bit value set in a service-tag action, such as change-tag, can override the provisioned IGMP P-bit value.
 - The E7 supports broadcast IPTV service using IGMPv2 Snooping (RFC 4541), in both “Snooping with Report Suppression” and “Proxy” modes. The E7 configured for "Proxy" mode also supports IGMPv3.
 - When a VLAN IGMP mode is set to proxy, you can specify the name of the IGMP profile to associate with the VLAN. You can also leave the system default IGMP profile assigned to the VLAN.

GPON Configuration guidelines

- The GPON card supports up to 32 IGMP-enabled multicast VLANs across the system, which can all be on a PON.
- The video VLAN must be the same on the network side (IGMP router side) and access side (IGMP host side).
- The ONTs do not support tagged and untagged match rules on the same ONT for video services. Only one classification rule is allowed for video services per ONT.
- All nodes in an ERPS or G.8032v2 ring must have the same IGMP Snooping provisioning on the video VLAN for traffic to flow--either all with snooping (snoop-suppress, proxy) or all without snooping (flood).
- The E7 supports up to 2000 IGMP multicast groups, enabling up to 2000 independent broadcast video and music channels.
- In the multicast profile, the Max Streams value should include the boot channels, the programming guide, and any other channels that may be learned, in addition to the number of STBs on the port.
- The upstream router **MUST** be an IGMP querier.
- The Query Interval parameter in the IGMP profile must match the upstream router configuration.
- A multicast whitelist can be applied at the time of creating the service, or later to an existing service as follows:
 - a. In the Navigation Tree, click **ONTs**, and then in the Work Area click **Provisioned ONTs** and then select the ONT from the displayed table.
 - b. Click **Ports > Multicast White List > Create**, and then select the service and multicast white list(s).

- The E7 passes the following multicast IP addresses through the system on a VLAN that has IGMP Snooping enabled. (For VLANs with IGMP Snooping disabled, the E7 passes all multicast traffic transparently.)

HSRP HELLO	224.0.0.2
DVMRP	224.0.0.4
OSPF ALL RTR	224.0.0.5
OSPF DES RTR	224.0.0.6
RIP	224.0.0.9
EIGRP	224.0.0.10
PIM	224.0.0.13
VRRP	224.0.0.18
HSRP	224.0.0.102
MLS ALL SNOOPERS	224.0.0.106

- The E7 can perform the following tag actions on a VLAN that has IGMP enabled:
 - Add-tag action to untagged Ethernet frames.
 - Change-tag action to change the VLAN ID at network administrative boundaries.
- IGMP snooping is only enabled on the outer VLAN ID.
- If a connected C7 is using IGMP proxy, then snooping must be enabled on the video VLAN. Otherwise, snooping should be disabled (set to N).
- With Router Learning Mode configured for 'static-only', IGMP Proxy will not allow a static multicast router ('mrouter') interface to be a multicast destination. If the interface could become a multicast destination in the event of a network topology change, the interface should not be configured as a mrouter interface.
- The network switch providing the video must have IGMP enabled.
- IGMP control messages are forwarded in both directions around an ERPS ring, enabling the population of the multicast forwarding tables for all connected E7 platforms on the ring. In the event of a ring switch, video will be able to quickly flow in either direction without waiting for the IGMP protocol to resolve the video paths.
- IGMP and Multicast traffic are not allowed to traverse ports blocked by the ring protection protocol.
- The E7 monitors all interfaces associated with the IGMP multicast VLAN for the presence of an IGMP General Query message to automatically detect the multicast router.
- The E7 drops joins and leaves received on router (uplink) side ports. Ports connected to a router are assumed to be facing the network and should not receive joins and leaves.
- IGMP is enabled for all ports designated as ring ports.
- IGMP snooping occurs internally to the E7 and Calix ONT and is not a networking protocol.

- P-Series Home Gateways support multicast and unicast IPTV services on either the Routed WAN interface or L2 bridge ports provisioned as either Half-Bridge or Full-Bridge mode.
 - When the multicast profile is assigned to the L2 bridge, a single IGMP Snoop function manages the multicast flows across all Half-Bridge ports, or across the Full-Bridge group.
 - When the multicast profile is assigned to the Routed WAN interface using Home Gateway, it supports both IGMP Snoop and Proxy.
 - All modes support Multicast VLAN Replication (MVR) as part of the Multicast Profile.
- Multicast VLAN Registration (MVR) is used at the ONT to merge the common multicast VLAN content into the subscriber's untagged service on the ONT User-Network Interface (UNI). The E7 MVR implementation allows up to 4 independent multicast VLANs to be merged into the subscriber's unicast service. These multicast VLANs can be used to segregate various content (e.g. HD, SD, Audio, local channels) or facilitate additional video provider content that a service provide may seek to distribute to end users.
- The E7 currently supports RFC 4541 IGMPv2 and RFC 3376 IGMPv3. Most IPTV set top boxes support IGMPv2.

Best Practice: To change parameters of an existing service, first remove the service and then re-add the service with new parameters, especially if using 3rd party provisioning equipment.

Configuration Process

The high-level process for turning up video service on the E-Series is described below.

Starting point

Before starting the configuration process, check that the following conditions are met:

- Transport is established to all nodes providing video service.
- The uplink card is installed and wired.
- The xDSL and OLT service cards are installed and wired.
- (GPON applications only) The ONTs are installed, wired, turned up and ranged.

Configuration process overview

The data service configuration process follows:

1. Verify that all required equipment is installed and wired.

2. Configure profiles and templates.

- Rules to classify traffic with service match list and service tag action
- Create Ethernet bandwidth profiles to define your different video service offerings
- Create a multicast profile (if MVR video is being provisioned, the multicast profile must reference an MVR profile that indicates the MVR VLAN and IP address ranges)
- Create a multicast white list map and ranges (if bundles of video channels are being defined)
- Quality of Service (QoS) actions to perform on ingress traffic packets when the packets match selection criteria with policy maps

3. Configure video subscriber services.

- a. Configure the xDSL or GPON port.
- b. (xDSL only) Configure the xDSL port associated interface.
- c. Add the data service to the subscriber port.

See *Calix E7 GPON Applications Guide* and *Calix E7-2/E3-48C xDSL Applications Guide* for configuration details.

Multicast Profile

Multicast service can be added to any xDSL or ONT Ethernet service by adding a Multicast Service Profiles to the subscriber service. Video service is not available to a subscriber unless the subscriber's service has been explicitly provisioned to include a multicast profile. Rouge IGMP messages sent upstream on a subscriber data service that does not have a multicast profile will be dropped. CLI and Web GUI tools are available to trace all multicast profiles to the associated subscriber Ethernet ports.

Profiles are named and contain the following attributes:

- **Stream Count** – Number of streams allowed at the subscriber port
- **Query Interval** – Time between sending group specific query and termination of the video stream
- **Multicast to unicast** – When enabled, converts multicast to unicast traffic toward the subscriber
- **Multicast Address Map** – Identifies the optional global allowable multicast IP ranges
- **MVR Profile** – Identifies the optional MVR profile to be used by this multicast profile

Global Multicast Address Map

The E-Series provides up to eight Multicast Address Maps to specify the allowable IP multicast addresses that are to be used for a particular multicast profile. Each Multicast Address Map contains up to eight multicast IP address ranges.

The Multicast Address Map is an optional parameter of the Multicast Profile. With no Multicast Address Map identified, multicast group destination addresses are unconstrained.

Multicast VLAN Registration (MVR)

Multicast VLAN Registration (MVR) is designed for distribution of multicast traffic on a dedicated multicast VLAN across disparate access networks, while allowing subscribers who are on different VLANs to join and leave the multicast groups carried in the Multicast VLAN. MVR isolates the network multicast VLAN used for efficient transport of multicast data from the subscriber. Typical applications include:

- Distribution of multicast VLAN and merging multicast traffic into subscribers “single service” associated with a residential gateway
- Distribution of multicast VLAN and converting to a common “Video VLAN” pre-configured on all STB
- Distribution of multiple multicast VLANs (e.g. high definition IPTV, standard definition IPTV, and digital audio) and merge into a single subscriber service

Note: The ONTs do not support tagged and untagged match rules on the same ONT for video services. Only one classification rule is allowed for video services per ONT.

MVR works in conjunction with IGMP. The subscriber continues join and leave multicast groups via IGMP. However, both the IGMP messages and multicast content are mapped from the subscriber’s service into the isolated network multicast VLAN. The E7 supports the ability to map the subscriber’s IGMP and multicast traffic into multiple network multicast VLANs.

MVR Profile Guidelines:

- Each E-Series supports up to 5 MVR profiles
- Each MVR profile can contain up to 4 MVR VLANs
- Each MVR VLAN can have up to 4 multicast IP address ranges

One MVR VLAN per MVR profile is allowed to have no IP range specified. All IGMP and multicast traffic are allowed on this VLAN. All IGMP and multicast traffic not matching any other MVR VLAN with defined IP address ranges will be mapped to this unspecified VLAN

The most common MVR profile will contain a single MVR VLAN with unspecified IP address ranges. Because MVR isolate subscriber video traffic from the network multicast VLAN, Calix recommends that MVR be used for all multicast IPTV applications.

MVR Profile and MVR VLAN Addresses

The Multicast VLAN Registration (MVR) feature gives you the ability to "register" multicast VLANs with a subscriber unicast VLAN.

MVR works in conjunction with IGMP where subscribers join and leave multicast groups via IGMP. However, both the IGMP messages and multicast content are mapped from the subscriber's service into the isolated network-side multicast VLANs. At the subscriber port level, MVR functionality switches multicast video from network-wide multicast VLANs onto the unicast VLAN, delivered tagged or untagged to the subscriber.


This topic describes how to create an MVR profile where a list of ranges specifies how an IP address range maps to a particular multicast VLAN.

Typical applications include the following:

- Distribution of multicast VLAN and merging multicast traffic into subscribers "single service" associated with a residential gateway
- Distribution of multiple multicast VLANs (for example, high definition IPTV, standard definition IPTV, and digital audio) and merge into a single subscriber service

Note: Calix recommends that an MVR profile be used for all multicast IPTV applications to move the multicast traffic out of the VLAN that has all of the middleware traffic, STB DHCP traffic, and so on.

The following diagram shows two variations of this model: One high-speed-data VLAN carrying both HSI and unicast video (left node) and separate VLANs for data and unicast video (right node).



* Note: One xDSL port shown for simplicity.

In this model, multicast video is isolated from unicast video via separate VLANs. One or more dedicated multicast VLANs are used across the entire access network, while unicast video is isolated into smaller broadcast domains on separate VLANs (N:1 or 1:1).

Configuration guidelines

- Each system supports up to 16 MVR profiles.
- Each MVR profile can contain 4 ranges comprised of an upstream MVR VLAN and multicast ranges
 - The full range of Multicast addresses:
 - Start 224.0.0.1
 - Stop is 239.255.255.255
- The address ranges defined in an MVR profile must be distinct from each other, not overlapping.
- The address ranges defined in different MVR profiles, are allowed to overlap.

- One MVR VLAN per MVR profile can have no IP range specified so that all IGMP and multicast traffic not matching any other MVR VLAN with defined IP address ranges will be mapped to this VLAN with no assigned IP range.
- An MVR VLAN cannot also be used as an Ethernet service VLAN.
- A subscriber port is allowed to be associated with only one MVR profile.
- Multicast video traffic (broadcast content) and unicast video traffic (STB DHCP activity, VOD, etc.) must arrive at the E7 on separate VLANs.
- For MVR configuration on a subscriber port, create a minimum of two VLANs:
 - One VLAN for Unicast traffic (DHCP and VOD packets) with IGMP Mode = flood where it correlates to the matched traffic rule and associated service-tag action when provisioning the service on the subscriber port.
 - One VLAN for Multicast traffic (IGMP and video packets) with IGMP Mode = proxy where it correlates to the MVR VLAN profile definition that is associated to the Multicast Profile used when provisioning the service on the subscriber port.
- The ONTs do not support tagged and untagged match rules on the same ONT for video services. Only one classification rule is allowed for video services per ONT.

Multiple Video Providers

- In order to support multiple video service providers over the same access network, the system supports 16 MVR profiles and one video provider per MVR profile. Each MVR profile specifies a multicast VLAN used for delivery of IPTV streams for that provider. Within a given service provider the multicast ranges must not overlap.

Overlapping Multicast Addresses

- Given that each video service provider is assigned a unique MVR profile and associated multicast VLAN, each service provider may use an arbitrary range of multicast addresses without concern for other providers that are sharing the access network. A given subscriber port, can only be associated with a single video service provider. The range of multicast addresses that are available on one subscriber port can overlap with the range of multicast addresses used on another subscriber port without conflict provided that the multicast addresses are sourced from different multicast VLANs (different providers).

Parameters

You can provision the following parameters for an MVR profile and address ranges:

Parameter	Description	Valid Options
Name*	Name of the MVR profile.	32 character string
MVR VLAN ID*	Index value designating the VLAN to be assigned address ranges.	any existing video VLAN
Address Range# Start*	First address in the MVR VLAN range.	Multicast IP address in "dotted quad" format: "224.0.0.1".

Parameter	Description	Valid Options
Address Range# End*	Last address in the MVR VLAN range. The full range of Multicast addresses: - Start 224.0.0.1 - Stop is 239.255.255.255	Multicast IP address in "dotted quad" format: "224.0.0.1".

*Required fields


Multicast White List Map and Ranges

Multicast Whitelists allow a service provider to offer IPTV packages of channel bundles to their subscribers without using a middleware server. The feature can also be described as Ala-Carte video, enabling the service provider to easily create and enable complex channel line-ups on a per subscriber basis. This allows the service provider to define bundles of video channels in a flexible and reusable manner. For example:

- One Multicast Whitelist defines all the local television channels at standard definition and another Multicast Whitelist defines local television channels at high definition.
- Multicast Whitelists define channel bundles like HBO, Cinemax, ESPN, one for each individual bundle the service provider has to offer.


The total number of Multicast Whitelists that can be defined on an E7 system is 128. From the 128 defined Multicast Whitelist, this feature allows up to 16 Multicast Whitelists to be associated with an individual subscriber. This enables the service provider to offer a specific package of channel bundles to each of their subscribers.

The packages of channel bundles that are defined by Multicast Whitelists are available to subscribers served by VDSL or GPON interfaces. Given that any UNI on a subtended modem or ONT device will offer the same channel lineup, there are some limitations of the feature described in the below diagrams.



Multicast Whitelists - SFU Configurations

Multicast Whitelists work with Calix ONTs when services are delivered over GPON fiber interfaces. For an SFU deployment, every set top box in the home has access to the entire channel lineup for the subscriber being delivered from the ONT device. For MDU deployments, each port of the MDU device is associated with a single video package. Only Calix MDU devices are supported. Subtending a switch or router from an SFU ONT to serve an MDU environment is not supported.



Multicast Whitelists - MDU Configurations

The feature offers the service provider flexibility to accommodate a variable number of multicast channel sequences in a channel bundle. A simple channel bundle can be a single multicast address or a single range of multicast addresses. But the service provider may also need to define bundles that consist of a larger number of multicast channel ranges. For this reason, Multicast Whitelists support up to 32 multicast channel ranges, although it is expected that on average only a few ranges will be required in this type of channel filter definition.

Configuration guidelines

- A multicast white list is a list of multicast group ranges which are allowed for the video service. Subscriber channel changes for multicast group IDs that are not included in the configured white list, will be ignored.
- The multicast white list is an optional parameter of the video service, and a multicast address map is an optional parameter of the multicast profile.
- With no multicast white list or multicast address map identified, multicast group destination addresses are unconstrained.
- When both multicast white list and multicast address map are applied to a service, only the white list ranges are passed on to the ONT or xDSL port. The multicast map is ignored.

- Maximum of 16 Multicast white lists are supported per service, where each Multicast white list supports a maximum of 32 multicast channel ranges, for a maximum total of 512 ranges.
- A multicast whitelist can be applied to an existing video service, or at the time of creating the service.
- Each of the 32 ranges of channels within a multicast white list must be contiguous.
- A Multicast Whitelist can NOT have overlapping multicast ranges within a Video Service.
- In MDU environments, only Calix MDU devices are supported.
- Subtending a switch or router from an SFU ONT to serve an MDU environment is not supported.
- For an SFU deployment, every set top box in the home has access to the entire channel lineup for the subscriber service being delivered from the ONT device. For MDU deployments, each port of the MDU device is associated with a single video package.
- The maximum number of multicast filtering entries supported per Calix ONT device depends on the Calix ONT family:

ONT Model Family Maximum Channel Sequences	
T-series SFU (T07x) models	64
T-series MDU (T7x) models	128
P-series 700G, GE, GX models	128
836GE RSGs	128
GigaFamily	128

IGMP Snooping Operation

In an IPTV network, broadcast television channels are delivered using IP multicasting. IGMP is the control mechanism used to control the delivery of multicast traffic. IGMP joins and leaves are sent from the IGMP host (the Set Top Box) to the upstream network to indicate which multicast group (which channel) is being watched. Access networks may use IGMPv2 or IGMPv3 to deliver IPTV services. IGMPv3 improves the interoperability between the access network and the set top boxes used to deliver video services. Many set tops come configured for IGMPv3 by default so having the option of using IGMPv2 or IGMPv3 allows the network operator to pre-provision the interfaces used to deliver IPTV video services to match the preferred operation of the set top manufacturer. In release 2.4, the Calix implementation of IGMPv3 does not support Source Specific Multicast (SSM).

IGMPv2 requires two separate messages when changing a channel, one to leave the current channel and another to join the new channel. IGMPv3 sends one membership report that includes both the leave and join information which reduces the amount of IGMP traffic required to change a channel, improving the overall performance of the channel selection processing.

- The E7 system supports IPv4 IGMPv3, as specified in RFC3376.
- The E7 system supports sending IGMP General Query messages to the subscriber ports.
- The E7 Video Services support IGMPv3 Snooping on ONTs as per RFC 3376.
- The E7 Video Services supports IGMP Snoop for both IGMPv2 and IGMPv3 protocols simultaneously coming from different ports
- The E7 Video Services allows IGMP Snoop for both IGMPv2 and IGMPv3 clients on the same vlan, providing services to both.
- The E7 system supports proxy reporting on the upstream interface of the OLTs / DSLAMs

The E7 supports broadcast IPTV service using IGMPv2 Snooping (RFC 4541), in both “Snooping with Report Suppression” and “Proxy” modes. The E7 configured for "Proxy" mode also supports IGMPv3.

When a VLAN IGMP mode is set to proxy, you can specify the name of the IGMP profile to associate with the VLAN. You can also leave the system default IGMP profile assigned to the VLAN.

The E-Series supports industry standard IGMP snooping of multicast video leave and join requests sent between the set-top box and the video distribution network. IGMP snooping enables the system to determine which ports should be a member of a multicast group. It can provide a local response if the channel already exists on the E-Series or it will forward those requests upstream to a multicast router or another IGMP snooping Ethernet switch. When multicast channels are received from the network, the E-Series will forward the channels to all ports that are currently members of the requested IGMP membership group. The ability to perform snooping can be enabled on a per VLAN basis. If multicast traffic is received on a VLAN where IGMP snooping is not enabled, the traffic will be handled as broadcast traffic and sent to all ports on the VLAN.

Note: For the IGMP profile, Calix recommends setting (or leaving) the IGMP Query Interval value for the video VLAN to 60 seconds (default).

A 60-second query interval allows a good balance between maintaining bandwidth efficiency (channel pruning) versus limiting potential IGMP control message processing overflow by the CPU, and it matches the typical default IGMP Query setting on most routers.

Key IGMP concepts:

- IGMP control messages are allowed to traverse all ports regardless if they are blocked by the ring protection protocol to enable population of the bridging tables for all connected E-Series platforms.
- Multicast traffic will not be allowed to traverse ports blocked by the ring protection protocol.
- The E-Series monitors all interfaces associated with the IGMP multicast VLAN for the presence of OSPF messages to automatically detect the multicast router.
- The E-Series will drop joins and leaves received on router (uplink) side ports. Ports connected to a router are assumed to be facing towards the network and should not receive joins and leaves.
- IGMP will be enabled for all ports designated as ring ports.

IGMP Scalability

The E-Series is a Layer 2 device and supports Layer 3 IGMP snooping and line-rate multicast with non-blocking capacity up to the limit of the physical interface. With proper QoS configured to favor Video over other data flows, channel blocking will occur with an E-Series only if the number of channels being delivered is greater than the capacity of the fiber. The E-Series is optimized for IGMP performance (a considerable portion of the CPU power is reserved for IGMP processing). The E-Series line cards have been tested to handle a sustained rate of 250 channel changes per second. The system will protect itself from an excessive rate of channel changes by dropping the channel change request traffic.

The E-Series IGMP processing delay varies based on the IGMP load on the network. With a typical well-designed network, the contribution of each E-Series on the overall IGMP response time is a few milliseconds. The maximum channel change latency introduced by the E-Series has been measured to be on the order of 10ms to 20ms under load. However, within a network deployment, the E-Series is not the only component in the network that impacts the channel change time as perceived by a subscriber. The subscriber set top box and the multicast router are more than likely the largest contributors to the overall IGMP processing time.

As a transport and aggregation component of the network, the E-Series is highly optimized to provide the least possible impact on the channel change behavior. The E-Series facilitates fast convergence of video services after a topology change in the layer 2 network such that the service can be restored to the effected parts of the network sooner than the General Query time configured on the network. Furthermore, the E-Series reduces the overall network load on multicast services by utilizing IGMP report suppression to reduce the number of message sent upstream to the multicast router.

IGMP Profile

Multicast delivers IP packets to a select group of hosts on the network. The IGMP enables the system to manage the flow of multicast IP traffic, selectively forwarding and blocking flows to ports based on multicast group client join and leave requests.

The system can passively snoop on the following IGMP packets transferred between IP multicast routers, switches, and IP multicast hosts to learn the IP multicast group membership and configure multicasting accordingly:

- Query
- Report
 - Join
 - Leave (IGMP version 2)

The system forwards multicast traffic destined for multicast groups (discovered through IGMP snooping) to ports that are members of that group. The system also discards multicast traffic destined for multicast groups that it does not recognize. IGMP snooping generates no additional network traffic, allowing you to significantly reduce multicast traffic passing through your device.

Configuration guidelines

- The IGMP snooping is enabled on a per-VLAN basis.
- You can edit the system-default IGMP profile or create a custom IGMP profile.
- The IGMP profile applied to a VLAN only takes effect if the VLAN IGMP mode = proxy or snoop-suppress.
- IGMPv3 is only available in Proxy IGMP mode and does not operate in snoop-suppress mode.
- All nodes in an ERPS or G.8032v2 ring must have the same VLAN components configured for video traffic to flow:
 - IGMP mode must be all proxy or all snoop-suppress
 - IGMP version must be all v2, v3 or auto
- For ERPS and G.8032v2 ring configurations:
 - The ring node with a direct connection to the mrouter can be provisioned as IGMPv3 and operate as v3. If at least one node in the ring does not support IGMPv3, all of the other nodes in the ring not connected to the mrouter will operate as IGMPv2 regardless of the provisioned IGMP version (v2, v3, or auto). If or when all ring nodes become IGMPv3 compatible, the nodes operating version will transition to IGMPv3, regardless of the provisioned IGMP version (v2, v3, or auto). That is, the ring node that is connected to the mrouter dictates the operating version of all the ring nodes, assuming they are all IGMPv3 capable.

- For ring nodes with no directly attached routers, set the Router Learning Mode to Static-and-Dynamic. This ensures that only the ring is “learned” or allowed to be the multicast router source so that subtended RSTP rings that might incorrectly flood a general query around the RSTP rings (typically during a topology change) are never learned as the querying router.
- For ring nodes with a direct connection to the multicast router, configure statically-defined router ports to identify the uplink. See Creating VLAN Members.
- Multicast traffic will not be allowed to traverse ports blocked by the ring protection protocol.

Note: Calix recommends that the subtending devices be configured for IGMP proxy to minimize the load on the multicast router and the transport topology.

Note: By default, IGMP snooping is disabled on a VLAN. IGMP snooping should only be enabled for VLANs that carry multicast content. Typically one, but sometimes a few VLANs are utilized to distribute multicast content within the access network.

- The video VLAN must be the same on the network side (IGMP router side) and access side (IGMP host side).
- The system can perform the following tag actions on a VLAN that has IGMP enabled:
 - Add-tag action to untagged Ethernet frames.
 - Change-tag action to change the VLAN ID at network administrative boundaries.
- IGMP snoop (proxy or snoop-suppress) is only enabled on the outer VLAN ID and is expected to be single-tagged through the network.
- If a connected C7 is using IGMP proxy, then snooping must be enabled on the video VLAN. Otherwise, snooping should be disabled (set to N).
- With Router Learning Mode configured for ‘static-only’, IGMP Proxy will not allow a static multicast router (‘mrouter’) interface to be a multicast destination. If the interface could become a multicast destination in the event of a network topology change, the interface should not be configured as a mrouter interface.
- The network switch providing the video must have IGMP enabled.
- For an RSTP network, Calix recommends setting the following parameters as shown:
 - The VLAN IGMP Mode = snoop-suppress or proxy
 - The IGMP profile Router Learning Mode = static-dynamic
 - The IGMP profile Router Solicit On Topology Change = Y (enabled)
 - The interface is NOT designated as a static router port through the VLAN membership.

- For Cross-Card LAG interfaces, set the following parameters as shown to achieve the best re-convergence time:
 - The VLAN IGMP Mode = proxy
 - The Router Learning Mode = static
 - The Router Solicit On Topology Change = N (disabled)
 - The LAG interface is designated as a static router port through the VLAN membership.
- The following scenarios describe the VLAN operation with various IGMP version settings in the applied IGMP profile.
 - **VLAN set to IGMP Auto:**
 - The VLAN defaults to IGMPv3.
 - If the VLAN receives an IGMPv2 General Query from a GE interface, then the VLAN transitions to IGMPv2 and continues operating in v2 mode for a time period defined by two General Query cycles + a max response time ($60 + 60 + 10 = 130\text{sec}$). Each time a v2 General Query is received, the timer is refreshed. When the timer expires, the VLAN transitions to IGMPv3.
 - If the VLAN receives an IGMPv3 General Query from the multicast router when operating in v2 mode, it responds with IGMPv2 messages.
 - For nodes that are configured in a linear chain, where a multicast router connects to E7 Node 1:
 - ♦ If the VLAN in Node 1 transitions from v3 to v2, then a v2 General Query is sent to Node 2, forcing the VLAN in Node 2 to transition to v2 and any IGMPv3 traffic from Node 2 is discarded.
 - ♦ If the VLAN in Node 1 transitions from v2 to v3, then a v3 General Query is sent to Node 2 and any v2 traffic from Node 2 is discarded.
 - For an xDSL interface, the IGMP operation mode (v2/v3) depends on the VLAN configuration mode (v2, v3 mode) and connected devices:
 - ♦ If all of the VLANs on an xDSL interface operating in IGMPv3 mode, then the xDSL interface IGMP mode is controlled by the connected STB/device version. That is, when a IGMPv3 message is received from a subscriber facing interface, then the xDSL interface responds with IGMPv3 query as well.
 - ♦ If at least one of the VLANs on an xDSL interface is configured for IGMPv2 mode or an IGMPv2 message is received from a subscriber facing interface, then the xDSL interface responds with IGMPv2 query to subscriber and sends out IGMPv3 report to the uplink router.

- If all VLANs on an xDSL interface are operating in IGMPv3 mode and a subscriber interface receives an IGMPv2 message, the subscriber interface transitions to v2 mode and remains in v2 mode for two query cycles plus max response time (default: $60 + 60 + 10 = 130\text{sec}$). The system sends IGMPv2 QUERYs towards the subscriber. Each time the interface receives a v2 message from the subscriber, the timer is refreshed.
- **VLAN set to IGMPv2**
 - IGMPv3 messages from the uplink router, are discarded. If the VLAN receives an IGMPv2 REPORT/LEAVE from a GE interface or PON subscriber port, the v2 message is processed but the IGMPv3 QUERYs continue to be transmitted from the interface towards the subscriber.
 - IGMPv3 messages from the subscriber port are discarded and v2 General Queries are sent every query cycle. Only the packets that match the operating IGMP version are processed.
- **VLAN set to IGMPv3**
 - If the VLAN receives an IGMPv2 message from the uplink router or GE or PON subscriber port, it is discarded. Only the packets that match the operating IGMP version are processed.
 - However, if a v2 report is received on an xDSL port, the xDSL interface operation mode for the port changes from IGMPv3 to IGMPv2. In the absence of v2 REPORTs, the interface transitions back to v3.
- The IGMP version of the IGMP Query must match the provisioned VLAN setting in the associated IGMP profile if the version setting is set to either “v2” or “v3”. If there is a chance that a v2 querier may win an election upstream, use auto mode.
- If there are mixed v2 and v3 hosts on GPON or Ethernet ports, use v2 for the IGMP version setting.
- Video Services supports IGMP Snoop for both IGMPv2 and IGMPv3 protocols simultaneously coming from different ports.
- Video Services supports IGMP Snoop for both IGMPv2 and IGMPv3 clients on the same VLAN, providing services to both.
- IGMPv3 provides the ability for a host to selectively request or filter traffic from individual sources within a multicast group.
- For IGMPv3 operation, only 1 source address per multicast group per VLAN is supported.
- Source-specific group records received within a v3 Report are processed if the following conditions are met:
 - The VLAN operating version is IGMPv3.
 - The subscriber facing interface operating version is IGMPv3.
 - Only one source is specified in the group record.
 - The multicast address in the group record is within the provisioned SSM range.

- The following set of counters are supported on a per-VLAN basis:
 - IGMPv2 joins sent and received
 - Leave messages sent and received
 - Query solicits sent and received
 - IGMPv3 reports "to include" sent and received
 - IGMPv3 reports "to exclude" sent and received
 - IGMPv3 reports "is exclude" sent and received
 - Group-specific queries sent and received
 - IGMPv2 General queries sent and received
 - IGMPv3 General queries sent and received
 - Invalid IGMP messages received

IGMP profile parameters

You can provision the following parameters for an IGMP profile:

Parameter	Description	Valid Options
Name*	A descriptive name to identify the profile.	31-character text string
Proxy IP Address*	IP address to use when in proxy mode. This address is used as the source address in IGMP messages sent upstream, and it is also used as the source address for queries downstream towards subscribers. It should be a valid IPv4 address in the same subnet as the upstream router. This address can be the same IP as the Management VLAN for the node.	IP address in "dotted quad" format. For example: "192.168.1.100". Alternately, "none" can be used to reset the value to "0.0.0.0"
Immediate Leave	Whether a multicast stream is dropped as soon as a Leave is received. May be overridden per interface.	selected (enable) cleared (disable) (default)
Robustness	Number of IGMP group-specific queries sent per query interval when proxy is configured.	1-10 1 (default)
Last Member Query Count	Number of group-specific queries sent when a Leave is received.	1-8 2 (default)
Last Member Query Interval	Time to wait in milliseconds for a response to a group-specific query.	100-5000 1000 (default)
Router Learning Mode	Method used for learning the location of upstream routers. Note: With Router Learning Mode configured for 'static-only,' IGMP Proxy will not allow a static multicast router ('mrouter') interface to be a multicast destination. Static-only option requires an interface be added as a multicast router interface to the multicast VLAN with the applied IGMP profile.	static-and-dynamic (default) static-only

Parameter	Description	Valid Options
Router Solicit On Topology Change	Whether to send a Query Solicitation when a topology change occurs. When the topology changes in a network, such as when a link is added or removed from an RSTP domain, a Query Solicitation message is generated on all ports belonging to the VLAN for which IGMP snooping is enabled. If the upstream multicast Querier in the network supports Query Solicitation and has the function enabled, it will respond by sending a General Query out, causing devices to re-adjust to a new multicast source port location, if necessary.	enable (default) disable
Query Response Interval	Time to wait for responses to general queries, in seconds. Note: Setting the value lower than 10 could result in a high traffic rate.	1-20 10 (default)
Query Interval	Time interval in seconds between general queries. Note: Calix recommends setting (or leaving) the IGMP Query Interval value for the video VLAN to 60 seconds (default). A 60-second query interval allows a good balance between maintaining bandwidth efficiency (channel pruning) versus limiting potential IGMP control message processing overflow by the CPU, and it matches the typical default IGMP Query setting on most routers.	10-1000 60 (default)
Startup Query Interval	Time interval between general queries during startup in seconds.	2-250 15 (default)
Startup Query Count	Time to wait for responses to general queries during startup.	1-10 2 (default)
Host Purge Time	Purge time in seconds for host ports.	130-3600 260 (default)
Router Port Purge Time	Purge time in seconds for router ports.	60-600 260 (default)
Proxy VLAN IGMP Version	The VLAN IGMP operating version.	v2 (default), v3, auto
Source Specific Multicast Start Address	IP address that indicates the start of the range used by the system to switch E7 multicast traffic based on a destination IP instead of a destination MAC. Source-Specific Multicast (SSM) is a form of multicast in which the host subscribes to a multicast group from a specific video source. This allows the reuse of the same multicast group address for different video servers to provide different contents. Multicast IP not in the source specific range would be discarded.	232.0.0.0 (default)
Source Specific Multicast End Address	IP address that indicates the end of the range used by the system to switch E7 multicast traffic based on a destination IP instead of a destination MAC. Multicast IP not in the source specific range would be discarded.	232.255.255.255 (default)

Parameter	Description	Valid Options
Upstream Source IP Zero	Allows the option of setting 0.0.0.0 for upstream IGMP messages as the Source IP and the downstream to any valid IP address. This option accommodates the need to have the IGMP proxy address outside the IPTV subnet. This is because using any address other than 0.0.0.0 is rejected by the multicast router due to failing the subnet security check. • enabled (selected) - Use 0.0.0.0 for upstream IGMP messages as the Source IP of the upstream. • disabled (cleared) - Use proxy-ip for upstream IGMP messages as the Source IP of the upstream. Note: For ERPS-ring and G.8032-ring configurations, Calix recommends that the node adjacent to the active multicast router be set to "upstream-source-ip-zero" = enabled for IGMP-Proxy mode. The other ring nodes should have this option disabled.	selected (enable) cleared (disable) (default)

*Required field

VoIP Service

About the E-Series VoIP Services Applications

Voice

The E-Series supports switched voice services from the Calix family of GPON and Active Ethernet ONTs. Integrated Access Device (IAD) function is embedded in the ONT and natively interfaces to the GR-909 metallic POTS ports on the ONT. The ONT performs the call handling and IP packetization required for VOIP. Three ONT VOIP solutions are supported:

- Session Initiation Protocol (SIP) as specified in RFC-3261, with supplemental RFCs related to SIP employed on the ONT for performance enhancements
- ITU-T H.248
- Calix C7 TDM Gateway connection to PSTN switch interfaces (GR-303, TR-08 Mode II) The *Calix C7 VoIP Services Guide* provides a complete service description and provisioning instructions

SIP and H.248 voice services are delivered in conjunction with third party softswitches from BroadSoft, MetaSwitch, Genband, and others. *Calix CAB-10-015* includes softswitches test results and is continuously updated as new switches and ONTs are added to the Calix VoIP portfolio. The E7 provides QoS for voice traffic throughout the Ethernet transport network and uplink interfaces. ONT VoIP traffic is tagged with an 802.1Q VLAN tag and prioritized using 802.1p priority bits.

The ONT SIP solution supports the standard voice service offerings of a Class V switch, so the ONT VOIP implementation is transparent to the subscriber. These features include Caller-ID, Call Waiting, Call Transfer, Three-Way Calling, Distinctive Ring, Ring Splash, and E911 support. The suite of supported CLASS features depends on the type of SIP server (softswitch or C7 TDM gateway) that the ONT operates with because feature support varies by vendor. Consult the above referenced documents for a complete set of ONT voice service attributes.

E-Series based voice services are delivered in conjunction with either or both of the following:

- Calix C7 TDM switch interfaces (GR-303, TR-08 Mode I and Mode II) using a Voice over IP Resource card (VIPR) or Ethernet Gateway (EGW) card, as the VoIP to TDM gateway. The *Calix C7 VoIP Services Guide* provides a complete service description and provisioning instructions.
- Third-party softswitches from MetaSwitch and Genband (C15), for SIP signaling native VoIP services. Calix conducts softswitch interoperability testing through its engineering organization and System Quality Assurance (SQA) process. Calix continuously updates interoperability tests as new switches and protocols are added to the Calix VOIP portfolio.

Lifeline Voice over IP with SIP

The E-Series natively packetizes traditional lifeline, analog voice through an integrated SIP user agent (UA) for delivery across the IP / Ethernet network, ultimately terminating in a VOIP softswitch service environment. In order to ensure maximum voice flexibility, the E7 supports the following voice codecs: G.711 (packetization rates of 10ms or 20ms), G.726-32, and G.726-40. Calix has tested SIP VOIP interoperability with the MetaSwitch Call Agent CA9020 and Genband C15.

Lifeline TDM Voice with C7 GR-303 and TR-08 Gateway

In addition the E-Series has the ability to have VoIP traffic routed to and terminated on a Calix C7 for conversion to TDM, for a GR-303 or TR-08 interface hand off to a traditional Class 5 switch.

CALEA Compliance

P-Series ONTs configured for VoIP service populate the destination MAC address of all RTP packets with the MAC address of the default gateway for that particular VoIP protocol's IP host. This is done in order to ensure that all RTP traffic is able to be captured at the Service Provider's router if necessary for Communications Assistance for Law Enforcement Act (CALEA) purposes. If a direct RTP path is required between VoIP endpoints on the same subnet, it is required that the router must hairpin the RTP packets back to the same interface. This might be the case when media gateways for the VoIP service reside within the same subnet as the ONT VoIP IP host, or when Direct Media is in use. This functionality also provides the benefit that it does not require Split Horizon to be disabled or TLAN Service to be enabled for the VoIP VLAN in order for customers to use Direct Media.

GPON GigaCenter Voice Services

Calix has traditionally supported the transition from TDM to VoIP based services on 700GX and 700GE ONTs by allowing one port to be associated to a TDM interface group and another to a VoIP instance, either SIP, H.248 or MGCP. However, the GigaCenters do not support simultaneous TDM and VoIP instances. Although, GigaCenters do support TDM circuit based voice using C7 TDM Gateway, and SIP and H.248 VoIP protocols.

E7-2 VDSL2-48 Overlay Line Card

This card supports 48 integrated splitter ports, in support for analog POTS services.

Broadband Overlay: With this card, the E7-2 can be used when traditional POTS already exists and is passed through the remote location, originating from the access equipment located in the CO or an existing remote DLC. The DSL overlay is passed through an integrated splitter / combiner, allowing the delivery of voice, video, and data to a customer across a common, twisted pair cable. The Calix E7's GE uplinks may be aggregated on a C7 with GE-2p / FE-4p or GE-4s line cards or on an Ethernet switch or router.

POTS based VoIP Services with VDSL2-48C/E3-48C/E5-48C/E5-48

All 48-subscriber ports on the VDSL2-48C/E3-48C provide broadband DSL access and lifeline POTS service. Analog POTS is carried over the service loop using standard Voice Frequency carriers (300 to 3400 Hertz). A line card integrated voice splitter separates the voice from the data service for each subscriber signal, handing the POTS service to a digital signal processor that digitizes and converts it to native, standards based VOIP streams.

The VDSL2-48C/E3-48C supports different VOIP signaling protocols natively, the first of which used are SIP and a SIP-based Calix proprietary method of supporting TDM based GR-303 and TR-08 uplinks to a Class 5 switch deployed in conjunction with C7.

POTS interfaces of the VDSL2-48C/E3-48C have the following characteristics:

- Loop-Start signaling
- Impedance of 900 ohms + 2.16 μ f
- Subscriber loop lengths of 1830 ohms, including 430 ohms telset
- Ringing support for 15 REN/card ring source (5 REN at 1400 Ohms)
- Ringing Voltage of Balanced 86 Volts AC at 20Hz
- Loop current of up to 25mA (during normal operation), and 18mA (during battery backup)
- User configurable Gain (+2dB for A to D, -2dB for D to A)

Configuration Process

The high-level process for turning up VoIP service is described below.

Starting point

Before starting the configuration process, check that the following conditions are met:

- Transport is established to all nodes providing VoIP service.
- The uplink card is installed and wired.
- The xDSL and OLT service cards are installed and wired.
- (GPON applications only) The ONT's are installed, wired, turned up and ranged.

Configuration process overview

The VoIP service configuration process follows:

1. Verify that all required equipment is installed and wired.

2. Configure profiles and templates.

- Rules to classify traffic with service match list and service tag action
- Voice service profile:
 - xDSL: SIP Gateway, TDM Gateway
 - GPON: SIP Gateway, TDM Gateway, H.248 Gateway, MGCP Gateway
- IP Host for voice services
 - Each VDSL2 card allows one IP host, either TDM or SIP
 - Each GPON ONT allows one IP host per type of voice service: SIP, TDM, H.248, MGCP, and PWE3 if T1 ports exist

3. Configure voice subscriber services.

- a. Enable the VDSL2 voice or GPON port.
- b. Add the voice service to the subscriber port.

See *Calix E7 GPON Applications Guide* and *Calix E7-2/E3-48C xDSL Applications Guide* for configuration details.

TDM Gateway Overview


The Calix C7 TDM Gateway provides a gateway to the circuit-switched PSTN for VoIP lines on VDSL2 cards or the 700 family of ONTs operating in E7 GPON or AE mode. These units interface to the C7 via an IP / Ethernet link, and the C7 in turn provides a GR-303, TR-08 or GR-57 switch interface to the PSTN for the subtended VoIP lines.

Note: The E7 GPON system only supports a single VoIP Interface Group for 700GE, 700GX and 760GX, 800G, although the GX ONTs support multiple VoIP IP Hosts

TDM interfaces to a Class5 switch can be GR-303 or TR-08 (Modes 1 and 2) in combination with the C7's TDM Gateway function.

Operators can use the C7 VIPR / EGW DHCP server to provide IP addresses to VDSL2 or ONT VoIP hosts configured for TDM Gateway services. The supporting C7 must be provisioned to listen for DHCP DISCOVER messages from subscriber ports provisioned for TDM Gateway VoIP ports. The C7 integrated DHCP server simplifies setup and configuration of the end-to-end PSTN-VOIP connection.

For C7 TDM Gateway service, you must configure the SIP voice concentration group (SIP VCG) on the C7 network and add a VoIP connection to the E7. For details, see the Calix C7 VoIP Services Guide. For details, see the *Calix C7 VoIP Services Guide*. If you use CMS to provision a C7 TDM Gateway service, you can specify the CSIP VCG in a TDM Gateway Service Group that is referenced when you activate the service. See the *Calix E7 GPON Applications Guide*.



Ground Start for GPON TDM Gateway and SIP voice

New to the GX ONTs is ground start signaling. Ground start is a signaling technique used to produce dial tone by momentarily grounding one side of a two-wire circuit. The ring side of a wire pair is usually the side that is grounded. Ground start signaling is often used for telephone lines connected to a PBX as well as by many older pay phones. The alternative to ground start is loop start (i.e. POTS).

Both SIP and C7 TDM gateway-based VoIP services support ground start. Operators may choose between loop start and ground start signaling by selecting the desired value in the ONT-port provisioning CLI commands or Web GUI.

Note: Ground start service is only supported on Calix GX ONTs, and is not supported for H.248 VoIP service.

TDM Based VoIP Services on GPON ONTs are Interoperable with B6 Gateway

GPON ONTs now support interoperability with B6-640 TDM Gateway for MGCP VOIP to TDM switch interface translation (GR-303 and TR-08).

Single voice VLAN for networks with multiple trunking gateways

E-Series supports using a single voice VLAN for copper and fiber networks with multiple trunking gateways. Historically, network deployments with multiple trunking gateways required separate VLANs to distinguish voice service associated with specific gateways. Now, a single trunking gateway can be specified for each E-Series or P-Series device. These devices are then associated with the specific trunking gateway by setting the profile to use the server IP address as the DHCP offer filter.

- If the associated IP-host is using DHCP, then the DHCP offer from the specific server will be accepted

- If the local device receives a voice service offering from a trunking gateway that is not associated, it rejects the offer.

SIP Overview

About the Session Initiation Protocol and SIP Network Elements

The Session Initiation Protocol (SIP) is an application-layer control (signaling) protocol used to establish, maintain, and terminate media sessions between two or more endpoints. The SIP standard is defined in Internet Engineering Task Force (IETF) RFC 3261. The SIP signaling and session management functions enable set up and tear down of telephone calls over IP networks.

SIP provides the ability to locate an endpoint, determine its availability and media support capabilities, establish a session between the source and target endpoints, and terminate the session.


SIP works in conjunction with several other protocols, including:

- **Session Description Protocol (SDP):** Describes session parameters to allow the agents to negotiate and define session characteristics (codec choice, packet rate, port to use, etc.) for a given call.

Real-time Transport Protocol (RTP): Defines the format for bearer packets (voice traffic).

- **Real Time Control Protocol (RTCP):** Provides performance statistics and reports on RTP packet flows.

SIP controls only the signaling portion of a communication session, acting as a carrier for SDP packets (to define the session content) and the RTP packets (to carry the actual voice content).



SIP network elements

SIP is a peer-to-peer protocol, with the peers known as "agents." A SIP agent can function as either a client (initiates a session request) or a server (responds to or forwards requests). A SIP agent may be capable of performing either function, but only one or the other per session.

SIP clients typically include:

- **Gateways:** A gateway provides call control and translation functions, including translation between circuit-switched and packet-switched formats.
- **Phones:** IP-capable phones serve as direct SIP clients.

SIP server elements include:

- **Proxy servers:** A proxy server provides several functions, including authentication, authorization, network access control, routing of session requests, security, policy implementation, and call feature support.
- **Registrar servers:** A registrar server processes requests from clients to register their current locations for use by proxy servers.

The distinction between SIP server types is typically more logical than physical, as proxy and registrar servers may be collocated together on one SIP server, or may be separate elements.

E-Series SIP Solution Overview

The E-Series SIP solution supports the standard voice service offerings of a Class V switch, so the E-Series VoIP implementation is transparent to the subscriber. These features include the following:

- Caller-ID
- Call Waiting
- Call Transfer
- Three-Way Calling
- Distinctive Ring
- Ring Splash
- E911 support

The suite of supported CLASS features depends on the type of SIP server (softswitch or C7 TDM gateway) that the E-Series operates with because feature support varies by vendor.

General SIP Operation

The E-Series packetizes the analog, lifeline POTS signals and encodes the voice service payload using standardized voice over IP (VoIP) protocols. Each node includes an individual SIP User Agent (UA). In order to provide voice services, the E-Series needs a dedicated VoIP IP address to enable SIP and RTP communications with other servers. Session registration, authentication, and other required VoIP setup steps are supported as detailed in RFC3261, "SIP: Session Initiation Protocol." The E-Series supports the following voice codecs (coder / decoder):

- G.711
- G.726-32
- G.726-40


By default the E-Series are enabled for SIP services with non-compressed G.711-PCM for analog, lifeline voice. SIP-based signaling of voice services are always used in conjunction with a carrier grade VoIP softswitch. With the E-Series, Calix has tested SIP interoperability with the following softswitches:

- MetaSwitch Call Agent CA9020
- Genband C15.E
- Genband C20 (formerly Nortel CS2000)

Note: This is an ETSI interface (ETSI SIP for the ANSI loss plan).

Lifeline Voice over IP with SIP

The E-Series natively packetizes traditional lifeline, analog voice through an integrated SIP user agent (UA) for delivery across the IP / Ethernet network, ultimately terminating in a VOIP softswitch service environment. In order to ensure maximum voice flexibility, the VDSL2-48C/E3-48C supports the following voice codecs: G.711 (packetization rates of 10ms or 20ms), G.726-32, and G.726-40. Calix has tested SIP VOIP interoperability with the MetaSwitch Call Agent CA9020 and Genband C15.



Calix E-Series Broadband Convergence with Native SIP VOIP Handoff to Softswitch

SIP Server Redundancy - Static Host Provisioning

The Calix E-Series support SIP Server redundancy through the provisioning of Primary and Secondary Static Host IP address of proxy servers.

With SIP Primary and Secondary Static Host provisioning if the primary path or server is disrupted, the system will resolve to a pre-provisioned secondary server without the need for DNS. The ability to explicitly provision both the primary and secondary proxy servers can be used with ESA systems such as the Metaswitch Universal Media Gateway (UMG) to switch to a local gateway server when primary server is not available.

E-Series SIP Service Models Overview

When considering the implementation of SIP VoIP services in an E7 network, you have the choice of three voice service provisioning models:

- SIP Remote (legacy model, for P-Series ONTs only)
- SIP Gateway (recommended model for GigaFamily, P-Series ONTs, T-Series ONTs, and E7 OS family xDSL nodes)
- SIP Gateway Via TR-069 Provisioning (for GigaFamily, P-Series, and T-Series GPON home gateways)

Each model is briefly described below:

SIP Remote (SIP RMT) Provisioning Model

Note: Legacy model, for P-Series ONTs only.

This model uses a remote TFTP VoIP configuration file* along with typical per-port voice attributes selected during E7 subscriber provisioning. Dial plan information is always defined in the configuration file; in addition, the following per-port voice attributes can also be defined in the configuration file, overriding** local E7 provisioning:

- Call Waiting Enable
- Caller ID Enable
- Three way calling Enable
- T.38 Fax Relay
- Message Waiting Indicator Enable
- Direct Connect (Hot Line) and Direct Connect Timer (Warm Line).

***Note:** This configuration file (specified by the SIP Remote Profile) defines a large number of attributes, nearly equivalent to all the attributes defined by the SIP Gateway Profile, Dial Plan Profile/Rules, and the local TFTP VoIP configuration file, described later.

****Note:** To override local E7 provisioning, you must select the "Enable SIP config file (Overrides local provisioning)" checkbox in the Create SIP Service dialog box.

To use the SIP Remote provisioning model, the following components are needed:

- P-Series 700 ONTs
- A fully configured remote TFTP VoIP configuration file (PseriesTFTP_VoipConfig.txt)
- A TFTP server for storing and downloading the file as needed
- A SIP Remote Profile created with the above information
- SIP service created on the ONT voice port with the above SIP Remote Profile selected (via the E7 EWI)

SIP Gateway Provisioning Model

Note: Recommended model for GigaFamily, P-Series ONTs, T-Series ONTs, and E7 OS family xDSL nodes.

For all applications: This model uses the E7 SIP Gateway Profile, E7 Dial Plan Profile and Rules, and typical per-port voice attributes selected during E7 subscriber provisioning.

For GigaFamily, P-Series ONTs, and E7 OS family xDSL nodes: Default values for some additional voice attributes are automatically applied with ONT firmware or E7 OS family xDSL node software, and rarely require any changes. If changes are required, they may be implemented by applying a customized, local FTP VoIP config file.

Note: This file (FTP_VoipConfig.txt) defines a small number of attributes and is completely different than the legacy, remote TFTP VoIP configuration file (PseriesTFTP_VoipConfig.txt). This file can be retrieved and stored locally on the E7, and then applied to ONTs (via OMCI) or E7 OS family xDSL nodes.

For T-Series HGUs: An XML-based VoIP configuration file must always be applied for SIP services to function. This configuration file cannot be created directly; it must be ordered from Calix.

For E7 OS family xDSL nodes: Default values for xDSL VoIP coefficients are included with the operating software, and rarely need to be updated outside of software updates. However, if an update is required, it may be implemented by applying a VoIP Coefficient File.

To use the SIP Gateway Profile provisioning model, the following components are needed:

- GigaFamily, P-Series ONTs, T-Series ONTs, or E7 OS family xDSL nodes
- SIP Gateway Profile*
- Dial Plan Profile* and Rules*
- (Optional, non-T-Series) Local FTP VoIP config file* (FTP_VoipConfig.txt)
- (Required, T-Series only) XML VoIP configuration file
- (Optional/rare, xDSL only) VoIP Coefficient File created and applied
- SIP service created on the voice port with SIP Gateway Profile, Dial Plan Profile, and VoIP Config File instance (if required) selected

***Note:** Together, these provisioning elements define voice attributes that are nearly equivalent to all the attributes defined in the legacy, remote TFTP VoIP configuration file.

SIP Gateway Via TR-069 Provisioning Model

Note: For GigaFamily, P-Series, and T-Series GPON home gateways only.

This model allows SIP VoIP on Calix Home Gateways (HGWs) to be provisioned entirely via TR-069, outside of CMS or the E7. To support this model, a complete solution, including the E7 GPON access network, compatible HGWs, and the Consumer Connect Plus (CC+) ACS, is required. For complete details on the entire solution, see the *E7 GPON Application Guide*.

With the required configuration in place for HGW RGs and the E7 GPON system, this model uses CC+ to provision SIP VoIP on HGWs. For all applications, the following steps are performed in CC+:

- Create a Voice Service Profile (equivalent to the E7 SIP Gateway Profile)
- Create a Dial plan or use the default (equivalent to the E7 Dial Plan Profile)
- Add voice service to subscribers, configuring per-port voice attributes as required (equivalent to the per-port voice attributes selected during E7 subscriber provisioning)

For GigaFamily and P-Series HGWs: Default values for some additional voice attributes are automatically applied with ONT firmware, and rarely require any changes. If changes are required, they may be implemented by applying a customized, local FTP VoIP config file via CC+.

Note: If required, this file (FTP_VoipConfig.txt) is the same one used in the E7 SIP Gateway Profile provisioning model.

For T-Series HGWs: An XML-based VoIP configuration file must always be applied for SIP services to function, and can be applied via CC+. This configuration file cannot be created directly; it must be ordered from Calix.

Note: If required, this XML file is same one used in the E7 SIP Gateway Profile provisioning model.

SIP Provisioning Overview (SIP GW Model)

Note: The high-level process for this voice model is described below. For more details, see the *E-Series (E7 OS) GPON Applications Guide* or *E-Series (E7 OS) xDSL Applications Guide*.

SIP provisioning process for the SIP GW voice model

1. Create a Service VLAN and uplink for SIP voice services
2. Prepare to associate the Service VLAN with SIP voice traffic
 - Create a match list profile with an untagged rule that matches any criteria

Proprietary Information: Not for use or disclosure except by written agreement with Calix.

- Create a Service Tag Action with the following settings:
 - Match List = <the above match list profile>
 - Action = "Add Tag"
 - S-VLAN = <above VLAN ID>
 - Outer P-Bit Source/Value = Specify/5
- 3. Create the following profiles:
 - SIP Gateway Profile
 - Dial Plan Profile and Rules
- 4. Create and retrieve configuration files, if required
 - (Optional, non-T-Series) Local FTP VoIP config file
 - (Required, T-Series only) XML VoIP configuration file
- 5. Create SIP service on voice ports, with the following settings:
 - SIP Profile = <SIP Gateway Profile from Step 3, above>
 - Dial Plan = <Dial Plan Profile from Step 3, above>
 - IP Host settings:
 - Service Tag Action = <from Step 2, above>
 - Host Protocol = DHCP (*recommended*)
 - (Optional/conditional) Config-file-instance = <from Step 4, above>
 - All other settings as required
- 6. (Optional/conditional) For ONTs and xDSL nodes that specify a config file instance, ensure that the config file has been applied.
- 7. (Optional) Edit voice port settings, if required (to change the default signal type, Tx/Rx loss plans, etc.)

SIP Provisioning Overview (SIP GW Via TR-069 Model)

Note: SIP provisioning via TR-069 is only supported with compatible Calix Home Gateways.

This topic shows the sequence of procedures for configuring the SIP Gateway voice service model via TR-069.

Note: For GigaFamily, P-Series, and T-Series GPON home gateways only.

Calix provides a complete solution for provisioning and managing Calix Home Gateway voice lines entirely via TR-069, outside of CMS or the E7. This allows data and IPTV service providers to contract the support of voice services to third-party providers.

This solution includes the E7 GPON access network, compatible Calix Home Gateways, and a compatible ACS such as Consumer Connect (recommended). At the outset, this solution is for SIP voice services only. The attributes of the Calix SIP Gateway model (provisioned via E7/CMS) have been mapped to the TR-104 data model for compatibility with TR-069.

Supported equipment

The following equipment is required for this solution:

- E7 R2.4 or greater
- Compatible Home Gateways:
 - 844G/854G GigaCenters R11.1.30 or greater
 - 844GE GigaCenters R11.2.1 or greater
 - 844E GigaCenters R11.1.10 or greater
 - 812G/813G GigaHubs
 - T07xG T-Series HGU's R4.1.51.032 (R6.3.1.6) or greater
- Compatible ACS:
 - Consumer Connect (recommended)

Process overviews

Separate configuration process overviews are provided for GPON GigaCenters (844G/854G/844GE) and GigaHubs (812G/813G), 844E GigaCenters, and T-Series T07xG HGU's.

Process overview with GPON GigaCenters or GigaHubs

To support this solution, the following configuration must be in place:

- **Starting point configuration:**
 - The E7 GPON is configured with HGW ONT devices connected, discovered, and created with the correct ONT templates
 - The HGW RG interface has been enabled with the external management mode
 - The E7 system uplink for TR-069 management has been configured, providing connectivity between the ACS and the RG WAN Service VLAN for TR-069.

Note: The WAN Service VLAN for TR-069 management should be different than the VLAN for actual voice traffic.

- **HGW RG configuration:**

Note: The following configuration can be accomplished via TR-069, EWI, or RG Configuration File.

- The VoIP Gateway Management Protocol has been set to TR-069 mode

Note: Changing this setting from OMCI to TR-069 mode, or vice versa, requires a reboot.

- A unique WAN Service VLAN for VoIP service only has been configured with the following key settings:
 - IGMP = disabled
 - NAT = disabled
 - P-Bit = <high priority value, such as 5>
- **E7 configuration:**
 - A system uplink for SIP voice services has been created, with connectivity to the upstream SIP softswitch
 - A Gateway service has been added to the HGW RG interface with the following key settings:
 - BW Profile = <as required for 384k per voice line>

Note: For example, for three-way calling on one POTS port, use a 1280k profile (384k x 3).

- PON CoS = Derived
- Service Tag Action that matches on the WAN VoIP VLAN and P-Bit value (5), changes the tag to the E7 VoIP VLAN, and specifies the same P-Bit value (5).

Note: Adding this Gateway service replaces the need for adding voice services to the POTS ports of the HGW.

Once the above configuration is in place, you may use Consumer Connect (via TR-069) to provision the SIP Gateway (including the identification of the VoIP WAN VLAN) and voice features per line per HGW. When complete, the VoIP WAN VLAN will be associated with the POTS ports of the HGW and will be inaccessible to the LAN-side of the RG.

With Consumer Connect, do the following:

1. Create a Voice Service Profile
2. Create a Dial plan
3. Add Voice service to subscribers, configuring the voice features per line (including the dial plan)

The Voice Service Profile parameters and voice features per line are the same as those used with the SIP Gateway model provisioned via E7/CMS.

Note: You may also use the EWI to view or provision the SIP Gateway and voice features per line. However, this is intended for verification of TR-069 provisioning only; it is not a recommended technique for provisioning services throughout a deployment.

Note: If the VoIP Gateway Management Protocol of a HGW has been set to the TR-069 mode, a provisioning alarm will be triggered if traditional OMCI voice provisioning is attempted (for example, if you attempt to add voice services to POTS ports).

Process overview with 844E GigaCenters (with 700GE ONTs as GPON access endpoints)

Note: For carrier grade voice it is assumed that the ONT or access device that is upstream of the 844E will continue to be used for voice services. The 844E as a secondary device subtended from an ONT does not support an in-band management channel, such as OMCI or AE CLI, and does not include UPS telemetry. The 844E is not expected to be the primary voice service, rather the design and management model of the 844E allows it to provide secondary or non-lifeline voice services.

To support this solution, the following configuration must be in place:

- **Starting point configuration:**

- The E7 GPON is configured with 700GE ONT devices connected, discovered, and created with the correct ONT templates
- The 700GE Ethernet port is in the legacy bridged mode
- An untagged data service for TR-069 management is provided from the 700GE ONT Ethernet port to the 844E WAN port
- The E7 system uplink for TR-069 management has been configured, providing connectivity between the ACS and the RG WAN Service VLAN for TR-069.

Note: The WAN Service VLAN for TR-069 management should be different than the VLAN for actual voice traffic.

- **HGW RG configuration:**

Note: The following configuration can be accomplished via TR-069, EWI, or RG Configuration File.

- A unique WAN Service VLAN for VoIP service only has been configured with the following key settings:
 - IGMP = disabled
 - NAT = disabled
 - P-Bit = <high priority value, such as 5>
- **E7 configuration:**
 - A system uplink for SIP voice services has been created, with connectivity to the upstream SIP softswitch
 - A tagged data service for voice traffic is provided from the 700GE ONT Ethernet port to the 844E WAN port, with the following key settings:

- BW Profile = <as required for 384k per voice line>

Note: For example, for three-way calling on one POTS port, use a 1280k profile (384k x 3).

- PON CoS = Derived
- Service Tag Action that matches on the WAN VoIP VLAN and P-Bit value (5), changes the tag to the E7 VoIP VLAN, and specifies the same P-Bit value (5).

Once the above configuration is in place, you may use Consumer Connect (via TR-069) to provision the SIP Gateway (including the identification of the VoIP WAN VLAN) and voice features per line per HGW. When complete, the VoIP WAN VLAN will be associated with the POTS ports of the HGW and will be inaccessible to the LAN-side of the RG.

With Consumer Connect, do the following:

1. Create a Voice Service Profile
2. Create a Dial plan
3. Add Voice service to subscribers, configuring the voice features per line (including the dial plan)

The Voice Service Profile parameters and voice features per line are the same as those used with the SIP Gateway model provisioned via E7/CMS.

Note: You may also use the EWI to view or provision the SIP Gateway and voice features per line. However, this is intended for verification of TR-069 provisioning only; it is not a recommended technique for provisioning services throughout a deployment.

Process overview with T07xG T-Series HGUs

To support this solution, the following configuration must be in place:

- **Starting point configuration:**
 - The E7 GPON is configured with HGW ONT devices connected, discovered, and created with the correct ONT templates
 - The HGW RG interface has been enabled with the external management mode
 - The E7 system uplink for TR-069 management has been configured, providing connectivity between the ACS and the RG WAN Service VLAN for TR-069.

Note: The WAN Service VLAN for TR-069 management should be different than the VLAN for actual voice traffic.

- **HGW RG configuration:**
 - An RG Configuration file has been created and applied over E7 GPON OMCI with the following attributes:

- BBF mode enabled (default).
- WAN Service VLANs for Data, VoIP, Video, and TR-69 specified.
- The following settings specified for the unique VoIP VLAN:
 - ♦ IGMP = disabled
 - ♦ NAT = disabled
 - ♦ P-Bit = <high priority value, such as 5>
- A VoIP configuration file has been created and applied to set up loss parameters and other low-level voice attributes.
- **E7 configuration:**
 - A system uplink for SIP voice services has been created, with connectivity to the upstream SIP softswitch
 - A Gateway service has been added to the HGW RG interface with the following key settings:
 - BW Profile = <as required for 384k per voice line>

Note: For example, for three-way calling on one POTS port, use a 1280k profile (384k x 3).

- PON CoS = Derived
- Service Tag Action that matches on the WAN VoIP VLAN and P-Bit value (5), changes the tag to the E7 VoIP VLAN, and specifies the same P-Bit value (5).

Note: Adding this Gateway service replaces the need for adding voice services to the POTS ports of the HGW.

Once the above configuration is in place, you may use Consumer Connect (via TR-069) to provision the SIP Gateway (including the identification of the VoIP WAN VLAN) and voice features per line per HGW. When complete, the VoIP WAN VLAN will be associated with the POTS ports of the HGW and will be inaccessible to the LAN-side of the RG.

With Consumer Connect, do the following:

1. Create a Voice Service Profile
2. Create a Dial plan
3. Add Voice service to subscribers, configuring the voice features per line (including the dial plan)

The Voice Service Profile parameters and voice features per line are the same as those used with the SIP Gateway model provisioned via E7/CMS.

Note: You may also use the EWI to view or provision the SIP Gateway and voice features per line. However, this is intended for verification of TR-069 provisioning only; it is not a recommended technique for provisioning services throughout a deployment.

For example, to verify provisioning, confirm that BBF Mode Status = True and SIP Line 1 is registered.

H.248 Gateway Overview

H.248 voice services are supported for both GPON and xDSL subscribers, in addition to the supported SIP services. Provisioning H.248 voice service is done by defining H.248 gateway profiles within the E-Series and then assigning the gateway profile to individual voice (POTS) ports. The following parameters are included in the H.248 gateway profile.

- Base port number for RTP packets
- IP address or hostname of primary H.248 gateway
- IP address or hostname of secondary H.248 gateway
- Type of voice soft switch
- Prefix string to use for terminations
- ID to use for ephemeral terminations
- Maximum number of active calls for media gateway
- Number of terminations per call

Service Parameters

You can provision the following parameters for an H.248 Gateway service:

Parameter	Description	Valid Options
Gateway Profile*	The name of the H.248 gateway profile to use for the service.	any existing H.248 profile
Termination ID*	The ID to use for ephemeral terminations.	text string
Administrative Status	Administrative state of the service.	enabled disabled enabled-no-alarms

*Required fields

Note: Voice traffic is carried on expedited Class of Service 4. Since the H.248 Gateway IP host is a trusted entity, it is assumed that a 1:N VLAN model will be used for voice services.

Calix has tested H.248 voice services interoperability with the following softswitches:

- MetaSwitch Call Agent CA9020
- Genband C15.E
- Genband C20 (formerly Nortel CS2000)

Note: Testing is in progress.

MGCP Overview

By default the E7 will be enabled for SIP services with non-compressed G.711-PCM for analog, lifeline voice. SIP-based signaling of voice services are always used in conjunction with a carrier grade VoIP softswitch. The E7 supports VoIP Interoperability for the following softswitches:

- MGCP voice interface with MetaSwitch Call Agent CA9020
- MGCP voice interface with GENBAND C15

For E7 GPON services, you can create a profile that specifies the Media Gateway Control Protocol (MGCP) properties for the ONT's MGCP gateway services. The MGCP Gateway Profile defines the base RTP port, the primary and secondary gateway controller addresses, any termination prefix, and the ephemeral termination ID for the service.

Note: The 'Primary Gateway Controller' and 'Secondary Gateway Controller' parameters must be configured with a fully qualified domain name, when the IP host for MGCP services is provisioned using Fully Qualify Domain Name (FQDN).

Service Parameters

You can provision the following parameters for an MGCP Gateway service on an ONT port:

Parameter	Description	Valid Options
Gateway Profile*	The name of the MGCP gateway profile to use for the service.	any existing MGCP profile
Termination ID*	The ID to use for ephemeral terminations.	text string
Administrative Status	Administrative state of the service.	enabled disabled enabled-no-alarms
GR 303	For GR303 Option, only used for B6 TDM Gateway support. Enables use of the MGCP and RFC 2833 voice control protocol compatible with the B6-640 Trunking Gateway (TGW) in GR-303 mode.	enabled (selected), disabled (unselected)

*Required fields

Profile Parameters

You can provision the following parameters for MGCP gateway profiles:

Parameter	Description	Valid Options
Name*	Name of the MGCP Gateway Profile.	A 31-character text string
Base RTP Port	Base port number for RTP packets.	0-65535 49152 ‡
Primary Gateway Controller	IP address or hostname of the primary MGCP gateway controller, that is, the softswitch.	dot delimited IP address

Parameter	Description	Valid Options
Secondary Gateway Controller	IP address or hostname of the secondary MGCP gateway controller, that is, the softswitch.	dot delimited IP address
Termination Prefix	Prefix string to use for terminations.	text string TP ‡
Retry Timeout	Specifies the number of seconds to wait for a response to a request.	4-50 30 ‡
Restart Delay	Specifies the delay value sent in the Restart in Progress (RSIP) graceful tear down method in seconds.	1-600 1 ‡
Flash Hook Persist	Enables a hook flash signal to be treated as a persistent event. A persistent event is reported to the call agent even if the call agent has not requested to be notified when it occurs.	enabled (selected), disabled (unselected) ‡
On Hook Persist	Enables an on hook signal to be treated as a persistent event.	enabled (selected), disabled (unselected) ‡
Off Hook Persist	Enables an off hook signal to be treated as a persistent event.	enabled (selected), disabled (unselected) ‡

*Required field

‡Default

IP Hosts for VDSL2 Voice Services

Configure an IP host to define how a VDSL2 card obtains an IP address for communication and indicates the voice service VLAN.

For R2.5 and lower, each VDSL2 card supports one voice service type, either SIP or TDM gateway. Therefore, a VDSL2 card supports a maximum of one IP host for voice services.

For R2.6 and higher, the E7-2 VDSL2 combo card supports multiple IP hosts which may be configured using TDM (CSIP) or SIP. When using this configuration a larger line group of 864 lines is supported.

Configuration guidelines

- For TDM Gateway services, the E-Series allows operators to use the C7 internal DHCP server to provide IP addresses for VDSL2 VoIP hosts, or use an external DHCP server option.
- To provision the IP Host for a DHCP host protocol configuration, you must select **dhcp** for the Host Protocol parameter. Any previously assigned Static IP, Static IP Mask, and Static IP Gateway addresses are ignored, yet preserved. The DHCP host configuration occurs when the first POTS port is configured.
- To provision the IP Host for a Static protocol configuration, you must select **static** for the Host Protocol parameter, and then enter the Static IP, Static IP Mask, and Static IP Gateway addresses.
- The static IP gateway and subnet mask attribute are only required when static IP addresses are in use.

- The gateway address and subtending IP addresses must belong to the same subnet, as indicated by the mask.
- The 'Host Name' parameter must be configured with a fully qualified domain name.
- The IP Host object sets the card IP address, indicates the output service VLAN tag, and references a system-default service tag action (LcPotsSvcTagAction) that specifies the traffic as untagged and assigns a P-bit value that maps to the Expedited Forwarding CoS (CoS 4 or P-bit 5, 6, or 7).
- The VLAN indicated in the IP Host must already be created on the E-Series and at the time of the IP Host creation, the VDSL2 card is associated to the VLAN membership. When viewing the VLAN Associations, this association appears as "dsl-svc."

Parameters

You can provision the following parameters for IP host profiles:

Parameter	Description	Valid Options
Name*	Descriptive name for IP Host.	text string
Outer Tag	Indicates the customer-specific tag.	2-4093 (Except for 1002-1005 which are reserved for E7 operation.)
Host Protocol	Host protocol for the SIP client. If you select "static," you must also enter a static IP address, mask, and gateway addresses.	static, dhcp ±
Host Name	Host-name that will be transmitted in DHCP Option 81. The 'Host Name' must be configured with a fully qualified domain name.	text string
Static IP	IP address statically assigned to the VDSL2 card if the host protocol is static. This attribute is ignored, yet preserved, if the host protocol is DHCP.	4-byte IP address
Static IP Mask	IP network mask assigned to the VDSL2 card if the host protocol is static. This attribute is ignored, yet preserved, if the host protocol is DHCP.	4-byte IP address
Static IP Gateway	Static IP gateway 4-byte address for the VDSL2 card to use in routing its traffic to the SIP server, if the host protocol is static. This attribute is ignored, yet preserved, if the host protocol is DHCP.	4-byte IP address

IP Hosts for Voice Services on GPON ONT

IP Host ports on the ONT are pseudo-devices which allow bridge port paths for IP Host Termination at the ONT. These types of bridges are used to define data paths for SIP, TDM Gateway, and T1 over Ethernet (PWE3). Each IP Host interface may be configured with a unique IP address, mask, and gateway address. Optionally, the SIP IP Host may obtain the IP address configuration information via DHCP client protocol. Each IP Host has a dedicated MAC address.

When an ONT has more than one voice port provisioned with the same type of service (for example, SIP or TDM Gateway), all like services provisioned on that ONT are managed by a single IP host.

Configure an IP host to define how an ONT obtains an IP address for communication. An ONT may support the following IP hosts, if equipped with the necessary ports.

- SIP IP host indicated by **SIP** in the ID. (Only voice service T-series supports.)
- TDM Gateway IP host indicated by **TDM** in the ID.
- H.248 Gateway IP host indicated by **H.248** in the ID.
- MGCP Gateway IP host indicated by **MGCP** in the ID.
- T1 pseudowire (PWE3) services IP host indicated by **PWE3** in the ID.

VLAN Association to voice services

The IP Host object references a service tag action that indicates the voice service VLAN by one of the following methods:

- An IP host can directly specify an outer tag, when the referenced service tag action states "Specified in Service" for the Outer Tag parameter.

This allows you to apply a different VLAN tag for the same service type (SIP, TDM gateway, H.248 gateway, or T1 PWE3) on each ONT.

- An IP host can reference a service tag action that specifies the service VLAN ID for the provisioned voice services.

This allows you to apply the same VLAN tagging for all ONTs that use the referenced service tag action object.

Note: The VLAN assignment must be different between the VoIP IP host and the subscriber port. Sharing VLANs between a subscriber port and a VoIP IP host is not supported.

Configuration guidelines

- The ONT requires that voice traffic be set to an Expedited Forwarding Class of Service (CoS 4). Therefore, the associated service tag action must either specify CoS 4 or a P-bit value of 5, 6, or 7.
- For TDM Gateway services, the E7 allows operators to use the C7 internal DHCP server to provide IP addresses for ONT VoIP hosts, or use an external DHCP server option.
- The E7 GPON only supports a single VoIP Interface Group for 700GE, 700GX and 760GX, although the GX ONTs support multiple VoIP IP Hosts.
- The T-series ONTs only support a SIP gateway VoIP service that requires a SIP gateway configuration file which is referenced in the SIP IP Host object.
- T-Series does NOT support static provisioning of a DNS server for the SIP IP-Host. Therefore, you must use DHCP to specify the DNS server to the ONT.

- To provision the IP Host for a DHCP host protocol configuration, you must select **dhcp** for the Host Protocol parameter. Any previously assigned Static IP, Static IP Mask, and Static IP Gateway addresses are ignored, yet preserved.
- To provision the IP Host for a Static protocol configuration, you must select **static** for the Host Protocol parameter, and then enter the Static IP, Static IP Mask, and Static IP Gateway addresses.
- The E7 SIP IP-Host does not support DHCP Option #60.
- If the ip-host is set to **no-ip** (PWE3 only), no IP address is specified.
- For voice and pseudo-wire services the static IP gateway and subnet mask attribute are only required when static IP addresses are in use.
- The gateway address and subtending IP addresses must belong to the same subnet, as indicated by the mask.
- To provision an IP host for H.248 or MGCP services using Fully Qualify Domain Name (FQDN), follow these requirements:
 - The 'Host Protocol' parameter must be set to **dhcp**.
 - The 'Host Name' parameter must be configured with a fully qualified domain name.
 - The 'Primary Gateway Controller' and 'Secondary Gateway Controller' parameters in the H.248 or MGCP Gateway Profile must be configured with a fully qualified domain name.
- The static IP address must not be the same as the gateway address.
- A single IP host (ID indicates TDM) manages all of the TDM Gateway services on the ONT.
- A single IP host (ID indicates H.248) manages all of the H.248 Gateway services parameters on the ONT.
- A single IP host (ID indicates PWE3) manages all of the PWE3 services on the ONT.
- A single IP host (ID indicates MGCP) manages all of the MGCP Gateway services on the ONT.
- A single IP host (ID indicates SIP) manages all of the SIP services on the ONT.
 - The E7 supports 8 SIP configuration files for each combination of vendor/model/product.
 - Configuration file instances 9 through 16 are used exclusively for VoIP configuration files and can be specified either by instance number alone or by "voip-<1 through 8>."
 - The instance specified in the SIP IP Host must match the instance specified in the Retrieve ONT Configuration File operation. See Retrieving and Applying a SIP GW Configuration File.

Before starting

Before starting the configuration process, check that the following conditions are met:

- The service tag action that you want to reference in the IP host is already created.
- The IP address for the ONT, subnet mask, and IP gateway are already determined, if the host protocol is set to "Static."

Parameters

You can provision the following parameters for IP host profiles:

Parameter	Description	Valid Options
Svc Tag Action	Name of service tag action or "none" for no service tag action. IP host traffic is always untagged. Note: The referenced service tag action must specify P-bits that map to the Expedited Forwarding CoS (CoS 4 or P-bit 5, 6, or 7), as the ONT requires that SIP traffic be set to Expedited Forwarding. . For voice, typical deployments set to add a single tag upstream (and strip downstream) and set the upstream TCI to be set to 6 (expedite service)	index-name of tag action
Outer Tag	Indicates the customer-specific tag when the Output Outer Tag attribute of the referenced tag action states "Specified in Service."	2-4093 (Except for 1002-1005 which are reserved for E7 operation.)
Inner Tag	Indicates the customer-specific tag when the Output Inner Tag attribute of the referenced tag action states "Specified in Service."	2-4093 (Except for 1002-1005 which are reserved for E7 operation.)
Host Protocol	Host protocol for the SIP client. If you select "static," you must also enter a static IP address, mask, and gateway addresses. If the ip-host is set to no-ip (PWE3 only), no IP address is specified.	static, dhcp, no ip (PWE3 only)
Static IP	(Static host protocol only) IP address statically assigned to the ONT if the host protocol is static. This attribute is ignored, yet preserved, if the host protocol is DHCP.	4-byte IP address
Static IP Mask	(Static host protocol only) IP network mask assigned to the ONT if the host protocol is static. This attribute is ignored, yet preserved, if the host protocol is DHCP.	4-byte IP address
Static IP Gateway	(Static host protocol only) Static IP gateway 4-byte address for the ONT to use in routing its traffic to the SIP server, if the host protocol is static. This attribute is ignored, yet preserved, if the host protocol is DHCP.	4-byte IP address
Hostname	(MGCP and H.248 only) Host-name that will be transmitted in DHCP Option 81. The 'Host Name' must be configured with a fully qualified domain name.	text string
Config-file-instance	(SIP gateway only) Instances 9 through 16 are used exclusively for VoIP configuration files and can be specified either by instance number alone or by "voip-<1 through 8>." - voip-1 - Instance for VoIP (9) - voip-2 - Instance for VoIP (10) - voip-3 - Instance for VoIP (11) - voip-4 - Instance for VoIP (12) - voip-5 - Instance for VoIP (13) - voip-6 - Instance for VoIP (14) - voip-7 - Instance for VoIP (15) - voip-8 - Instance for VoIP (16) Note: The instance specified in the SIP IP Host must match the instance specified in the Retrieve ONT Configuration File operation.	None ‡, voip1-8, enter value

Voice (POTS) Overview

Voice (POTS) Loss Plan Provisioning

To improve the performance of the ONT Voice (POTS) interfaces the E7/E3-48C includes system transmit (tx) and system receive (rx) signal-level loss plans for POTS (voice) lines that provide attenuation settings, according to various standards, to reduce the perception of noise and echo on the line. Voice signal-level loss can be set to the following:

- The GR-909 loss plan (less attenuation – higher signal level) is more compatible with GR-303 type phone systems.
- The ANSI loss plan (more attenuation – lower signal level) is more compatible with VoIP type phone systems.

Note: The TX and RX loss plan settings must match.

Newer phones based to ANSI standards are designed to receive the voice signal at a lower signal level than older GR-909 based phones. By dropping the GR-909 gain level to match the ANSI loss plan, the signal levels of noise and reflected signal (echo) are lowered and the perceived performance of the ANSI phone is improved.

POTS based VOIP Services with VDSL2-48C/E3-48C/E5-48/E5-48C

All 48-subscriber ports on the VDSL2-48C/E3-48C provide broadband DSL access and lifeline POTS service. Analog POTS is carried over the service loop using standard Voice Frequency carriers (300 to 3400 Hertz). A line card integrated voice splitter separates the voice from the data service for each subscriber signal, handing the POTS service to a digital signal processor that digitizes and converts it to native, standards based VOIP streams.

The VDSL2-48C/E3-48C/E5-48/E5-48C supports different VOIP signaling protocols natively, the first of which used are SIP and a SIP-based Calix proprietary method of supporting TDM based GR-303 and TR-08 uplinks to a Class 5 switch deployed in conjunction with C7.

POTS interfaces have the following characteristics:

- Loop-Start signaling
- Impedance of 900 Ω + 2.16 μ f, supporting ANSI based loop plants
- Subscriber loop lengths of 1830 Ω , including 430 Ω telset
- Ringing support for 15 REN/card ring source (5 REN at 1400 Ohms)
- Ringing Voltage of Balanced 86 Volts AC at 20Hz
- Loop current of up to 25mA (during normal operation), and 18mA (during battery backup)
- User configurable Gain (+2dB for A to D, -2dB for D to A)

Overlay POTS Services with VDSL2-48

All 48-subscriber ports on the VDSL2-48 line card provide broadband DSL access and access to overlay lifeline POTS service. Analog POTS is carried over the service loop using standard Voice Frequency carriers (300 to 3400 Hertz). A line card integrated voice splitter separates the voice from the data service for each subscriber signal, extracting the POTS service, as an analog signal, back out of the line through the POTS overlay connector ports. The analog POTS service can then be handled via a collocated Digital Loop Carrier (DLC), or sent off to a TDM switch analog front end line bay.

POTS interfaces of the VDSL2-48 line card have the following characteristics:

- Impedance of $900\ \Omega + 2.16\ \mu\text{f}$, supporting ANSI based loop plants
- Loop current of up to 25mA (during normal operation), and 18mA (during battery backup)

VoIP coefficient files for international POTS support

Since the operating software on the VDSL2 card has built-in Coefficient values for all countries that are supported when the software release is created, you will only need to load a coefficient file onto the node when either of the following conditions exist:

- The built-in data for the given country needs to be corrected
- A new, previously unsupported country needs to be accommodated

The VoIP coefficient file contains data for all supported countries and can be downloaded from Calix and then transferred to server local to the system.

Note: The countries supported in the current operating software are available from the “Country Code” option in the SIP Gateway Profile.

Using VoIP Coefficient Files Overview


1. **Retrieve** the coefficient file from an external server that is reachable by the system, and place it in the VDSL2 flash memory.
The “cancel” action can be invoked before the “apply” action is performed, either while the “retrieve” action is still in progress, or after the “retrieve” action has completed.
2. **Apply** the coefficient file to the VDSL2 card, causing the card to reset and apply the contents from the file to the provisioning of all xDSL ports on that card. Any previously applied coefficient file is deleted from the VDSL2 card memory.
3. **Remove** the coefficient file from the VDSL2 flash memory, if it is no longer needed.
If no coefficient file is available, the VDSL2 card uses its built-in coefficient settings.

RF Overlay Video Service

The Calix E7 system and GPON card supports Radio Frequency (RF) overlay for Cable TV (CATV) services over the 47 to 870 MHz band. The RF video overlay path requires the following equipment:

- **RF Transmitter:** This device is typically located in the central office. At the headend, video signals are modulated (analog) or encoded (digital) into RF signals. These signals are combined and sent to a transmitter and optically modulated onto a 1550 nm wavelength optical signal.
- **EDFA Amplifier:** This device accepts the optical signal from the transmitter and amplifies the signal so it can be split to support multiple PONs, increase the PON reach, and deliver adequate signal to noise ratio at the consumer TV set.
- **Combiner:** This device combines the 1550 nm overlay video optical signals with the GPON 1490 nm forward path and 1310 nm reverse path, using CWDM passive optics and transported over a single strand of fiber into a PON.
- **ONT:** Each ONT has an optical triplexer that performs electrical-to-optical conversion (and vice versa) for the three bandwidth ranges supported by the GPON standards. The three frequency ranges are used for downstream IP data (including IPTV signals), upstream IP data, and RF video overlay (optional). The RF video signals are demodulated and exit the ONT via a standard coaxial cable fitting (RG-58). The ONT operates as the network interface device (NID) connecting the service provider network with the in-home wiring.

The following diagram shows the typical configuration of an overlay video network:




Note: Decreasing the number of splits increases the reach up to 40 km.

Refer to *CAB-09-037 HPNA 3.1 Home Networking Guide for IPTV and RF Video Applications* for details.

RF Video Overlay with RF Return Path

The E7 can be used with Calix 700 ONTs featuring integrated RF Return supporting a 5 to 42 MHz return path frequency band. This RF Return solution enables bi-directional communications between a set-top box, consumer electronics device, or cable modem.

The design of the system uses a similar topology and identical set-top box hardware for interactive services as current cable operators' HFC technology and practices. The product allows a single ONT for transport of both the forward and reverse RF paths, incorporates power management per FSAN power shedding standards, and can be operated without the GPON data link as a PON RF node for video and interactive set-top box services.



The Calix RF Return solution uses the 1610 nm wavelength as the return path from the ONT to the headend. The ONT generates the optical signal and multiplexes the wavelength with the PON signals. At the headend the 1610 nm wavelength is separated using the 4-wavelength CWDM module. The 1610 nm signal is terminated on a standard high-gain HFC reverse path optical receiver which demodulates it into an electrical RF signal, amplifies and sets the signal level. At this stage the signal may be combined with other return path signals depending on the aggregate number of set-top boxes to be supported. These combined signals are fed into the RF return path demodulators located within the headend and converted back into return path messages.

Refer to *CAB-07-025, RF Return System Features and Planning Guidelines*, for additional detail regarding RF video overlay and RF return network characteristics and planning.

Configuration guidelines

For purposes of configuring the RF Return system, it is assumed that the following conditions exist:

- The planned RF Return network has previously deployed Analog Video Services using the traditional 3-wavelength model.
- That the SA Prisma II chassis has been previously installed and is functioning.
- That an appropriately sized EDFA module has been installed and is functioning in the chassis.
- That the following new equipment has been installed and is functional:
 - Headend 4-wavelength CWDM Module
 - SA Return Optical Receiver Module
- That the 725 ONT has been installed and has been discovered on the network and that the STB is receiving forward path OOB video signals from the headend.
- That Local Craft Interface (LCI) software from Scientific Atlanta has been installed on the network and can communicate with the Prisma II chassis. The software can be downloaded from this site:

www.scientificatlanta.com/products/customers/iframe_EM_PrismaII_LCI.htm
- That the Out of Band (OOB) transmitter and Reverse Path Demodulator are communicating with one or more set top boxes which are currently not on the ONT network.
- To facilitate the connectivity for the subtending STBs off of the return device, all security must be disabled on the VLAN (no ip-source-verify). A Match List/Rule should be in place with the src-mac of the return device and OUI as the src-mac-mask. See Creating a Service Match List and Rules. If ip-source-verify is enabled, the subtending STBs will not work because their IP addresses are never learned by the system.

Note: For more comprehensive information on the Scientific-Atlanta RF Return solution, refer to the *Calix 725 ONT RF Return Configuration Guide*.

Configuration Process

The E7 system and GPON cards allow third-party overlay components to add RF video functionality to the PON. The RF signals do not traverse the E7 line card, and therefore do not consume system resources. For more information, see *Calix GPON RF Overlay Deployment Guidelines*.

Before starting

Before starting the configuration process, check that the following conditions are met:

- The E7 and ONTs are installed and wired, and ONTs are turned up on a GPON port.
- An EDFA amplifier has been set up to insert the analog video overlay signal onto the PON.

Configuration process

The GPON RF overlay video services configuration process is as follows:

- 1.** Verify that the analog video overlay signal is present on the PON.
- 2.** Configure the AVO Profile.
- 3.** Configure the ONT RF video port for service.
- 4.** If applicable, configure the ONT RF return port for service (for ONTs equipped with integrated RF return path only).

T1/E1 PWE3 Service

The E7 supports T1 and E1 services over GPON and point-to-point Ethernet using pseudowire emulation edge-to-edge (PWE3) technology which provides a standards-based approach to T1 or E1 service delivery over a packet switched network.

The E7 also supports T1 or E1 services on Calix ONTs equipped with T1 pseudowire ports (for example, the 766GX ONT), where traditional T1 or E1 services are converted to packet flows at the ONT and carried back to the E7, where they are delivered to a remote pseudowire device that translates the packet stream back to the original T1 or E1 format. Also see *Calix T1 Pseudowire Applications*.

About De-multiplexing PWE3 services

The available transport modes for E7 PWE3 services are UDP-IP and MEF. All UDP demultiplexing options conform to Mode 3 added for clarification RFC 5086 (CESoPSN) and Mode 1 added for clarification RFC 5087 (TDMoIP). MEF options conform to the ECIDs (Emulated Circuit IDs).

- Either the local (source) UDP port or remote (destination) UDP port may be used to multiplex and de-multiplex individual PWs between nodes, resulting in the UDP port acting as the PW Label.
- PWE endpoints **MUST** agree upon use of either the source UDP or destination UDP port as the PW Label.
- The configured source or destination port (one or the other, but not both) together with both the source and destination IP addresses uniquely identify the PW.
- All UDP de-multiplexing options conform to RFC 5086 (CESoPSN) and RFC 5087 (TDMoIP).

Note: Calix recommends using the RFC 5086 method for de-multiplexing. This method requires that the remote (destination) UDP port always be specified, avoiding port 2142.

CESoPSN Mode (RFC 5086):

- The destination UDP port is used to multiplex and de-multiplex individual PWEs between nodes, resulting in the destination UDP port acting as the PW Label.
- The configured destination port, together with both the source and destination IP addresses, uniquely identifies the PW for the receiver.

TDMoIP Mode (RFC 5087):

- When the source UDP port is used as the PW label, the destination UDP port number **MUST** be set to the IANA assigned value of 2142.
- Calix recommends **NOT** using UDP-based PWs, where port-translating middle boxes are present between PW endpoints.
- When using the destination UDP port 2142, the source UDP port must be set to the same value.

Configuration guidelines for operation with B6-440

- E7 supports the ability to provision PWE3 to interoperate with the B6-440 line card and T1 equipped ONTs by adding the ability to provision the RTP Payload Type – which defaults to a value of 96 but the B6 requires a value of 86.
- By default, the 766GX ONT requires header type 96. To operate with the B6-440, the ONT needs to support RTP header type 86.
- Set RTP to Enabled, and RTP Payload Type to 86 in the T1/E1 PWE3 profile.
- Keep in mind the following B6's functionality when configuring the PWE endpoints:
 - The B6 has a fixed Payload Size of 96 bytes.
 - Src UDP is fixed, Dst UDP is configurable.
 - It supports Adaptive and Loop timing modes but not Differential.
 - The Calix B6 only supports CESoPSN de-multiplexing mode, as detailed in RFC 5086.

Configuration guidelines for operation with MiTOP-E1/T1

- The E-Series supports the RAD MiTOP-E1/T1 TDM pseudowire access gateway in an SFP-compatible enclosure that extends TDM-based services over packet-switched networks, allowing for connectivity to existing T1 devices.
 - The MiTOP-E1/T1 can be inserted into any SFP port on the E-Series platform.
 - On the E-Series, create a service VLAN and add VLAN membership to the GE port where you install the MiTOP-E1/T1 module. Also, configure the MiTOP-E1/T1 to use that same VLAN ID.
 - To configure the MiTOP-E1/T1, refer to the *MiTOP-E1/T1 SFP-Format TDM Pseudowire Gateway Installation and Operation Manual*.
 - Also see *Calix T1 Pseudowire Applications* and *E7 QT-15-002 - Using a MiTOP SFP module in an E7 GE port for PWE applications*.

Before starting

Before starting the services configuration process, check the following conditions are met:

1. The network uplinks for GPON services are configured.

- E7 Ethernet ports
- Ethernet port interfaces
- Service VLANs

Within a Layer-2 network, T1/E1 services from an ONT are typically placed in a single VLAN for transport between the remote pseudowire devices and the IP gateway. Multiple remote locations may be split into separate VLANs for ease of network planning and operation.

- VLAN memberships

2. The necessary E7 system profiles that support GPON applications are created.

- ONT profile
- Quality-of-Service (QoS) for GPON traffic management

- Service match list

Accept the Default Match Criteria list: 1 - (all-untagged)

- Service tag action

Transport of T1/E1 services over a packet network requires a CoS priority and sufficient bandwidth allocation within the packet network to reduce latency and packet loss. T1/E1 service VLANs should be given a high IEEE 802.1p priority “p-bit” value (for example, 5).

You must create an "add tag" tag action in order to add a VLAN tag to all pseudowire traffic. Each pseudowire must have its own VLAN ID within the network. Both endpoints of the pseudowire will share the same VLAN ID. Create a separate tag action for each pseudowire.

The service tag is associated with the PWE3 IP host – one IP host, one tag action.

If you want to use a different VLAN for each PWE3 service, set the tag action to be “specified in service”. Then, when you provision the actual PWE3 service, enter the VLAN ID for that PWE3 service in the service provisioning.

- Profiles for T1/E1 services

All ports on an ONT must be either T1 or E1, which is indicated via the ONT PWE3 Profile when the ONT is created and cannot be changed on an existing ONT.

Assign IP addresses for the CO endpoint ONT and the CPE endpoint by defining the IP Host for PWE3 for each ONT at the PWE3 circuit endpoints. Endpoint IP addresses must be located within the same subnet.

3. The PON is configured.

- GPON port
- *ONT provisioning* (on page [514](#))

Step 4. T1/E1 PWE3 services configuration process

Assuming that the conditions stated above are met, the T1/E1 PWE3 services configuration process follows:

1. Provision T1 ports on the ONTs
2. Add pseudowire service to the T1 ports

Information you need

You must have the following information on hand to configure ONT T1/E1 PWE3 service:

- Tag Action name
- T1/E1 PWE3 profile name
- UDP port number
- Destination IP

E1 Bits Timing Support

Voice, video, and data services on the E7 GPON and point-to-point gigabit Ethernet (Active Ethernet) ONTs are all designed to operate without a traceable network timing reference (BITS clock input). In particular, the ONT VoIP subsystem does not require any additional network timing reference to operate. For triple play services, a BITS timing reference for the system is not required.

A traceable network clock synchronization for E7 GPON is only required in rare situations where differential clock recovery is being used for T1/E1 pseudowire applications.

For TDM T1/E1 Pseudowire (PWE) services delivered across a packet network that require a service clock recovery algorithm to recreate the original service clock at the far end pack-to-TDM transition. Two mechanisms are available for T1/E1 clock recovery:

- Adaptive Clock Recovery (ACR)
- Differential Clock Recovery (DCR)

ACR is used in most networks and is sufficient to recover the T1/E1 service clock for most applications. When a network is susceptible to high packet delay variation (PDV), and a traceable clock is available at both PWE3 end point devices, the T1/E1 circuit can be configured to use DCR and will be highly immune to network-induced PDV. Most T1/E1 deployments do not require DCR. A BITS Stratum 1 clock reference is required for T1 services if DCR is used.


Refer to the *Calix T1 Pseudowire Applications Guide*, for additional information on T1/E1 service delivery, clock recovery options, and network timing requirements.

The E7-2 and E7-20 platforms support E1 timing input to the BITS pins on the rear of each shelf. The system can be configured for a accept T1 input (1.544 Mbps) or E1 input (2.048 Mbps). To facilitate conversion from a BNC connector commonly used for E1 services to a 4 wire interface used for T1 services, Calix offers an optional RJ-45 to BNC connector cable.

Only the E7 GPON cards support clock projection to the ONT.

E7-2 GPON-4	projects timing on GE and PON
E7-20 GPON-4x	projects timing on PON
E7-2 GE-24	does NOT project timing
E7-20 GE-24x	does NOT project timing

The E7 shelves have wire wrap pin field for BITS input. To facilitate conversion from a BNC connector commonly used for E1 services to a 4 wire interface used for T1 services, Calix offers an optional RJ-45 to BNC connector cable. (100-02108 MX E1 RJ45 to BNC Female Cable, Un-balanced, 15cm)



PWE3 Timing and Clock Recovery Considerations

Ethernet networks handle system clocking different from SONET networks. SONET networks require that each node recovers its timing off of the previous node. Thus, any node in the system can trace its timing back to the Central Office Terminal (COT) and the network clock is considered "traceable." In the case of T1 service, having a traceable network clock means that the T1 service clock at any given node is the same as any other node.

The BITS clock input is NOT required for most GPON applications, including most T1/E1 applications. A traceable network clock synchronization for E7 GPON is only required in rare situations where differential clock recovery is being used for T1/E1 pseudowire applications.


In an Ethernet network each node may have its own timing source, thus the network clock may not be traceable from one location to the next. Given this lack of traceability, the T1 service clock is no longer guaranteed to be the same at different nodes. The receiver's service clock now needs to be recovered using timing information stored inside the Ethernet frame. The DS1-12 PWE and T1-8 PWE line units support Adaptive Clock Recovery, Differential Clock Recovery, Source and Loop timing.

- Adaptive Clock Recovery – This method uses the average arrival time of the packets to derive the transmitter's clock. This is possible since the source device is producing bits at a constant rate determined by its clock. By averaging the arrival times, the Packet Delay Variance (PDV) can be negated and the original sender's clock recovered.

- **Differential Clock Recovery** – This method employs a residual timestamp at the sender based on the difference between the shared reference clock and the T1/DS1 line. Thus, the receiver utilizes the timestamps in the received frames against its own time reference to correctly adjust its own T1 line timing. When a traceable clock is available between the two devices, this mechanism will produce lower jitter than adaptive clock recovery.
- **Source Timing** - Source timing is act of using a local source clock as the timing reference. This local source may be an injected clock via a BITS interface; derived from the PON, or directly from a different received T1 (GR-303 line, etc.)
- **Loop Timing** - Loop timing entails looping the received T1 clock back, using it to time the T1 transmissions.

Note: Loop timing is not supported on Axerra AXN10 Psuedo-Wire Access Devices for interfaces where the T1 port carries ces-pw traffic.

The figure below shows the relationship between network reference clock and the use of differential and adaptive timing methods for service clock recovery.



The advantage of adaptive clock recovery is that it does not require a traceable clock source at both locations of the T1 circuit. This solution covers deployments that contain any combination of asynchronous and synchronous T1 services where greater delay and jitter can be accommodated.

Unlike most applications transported over a packet switched network, packetized TDM services can be greatly impacted by the timing characteristics of the network. At each end of the T1 pseudowire, the T1 interworking function (IWF) must assemble the packetized T1 payload into its original format, including its original service clock. The ability to reconstruct the service clock is determined by whether a traceable network reference clock exists at both ends of the T1 pseudowire.

The illustration below shows the relationship between network reference clock and the use of differential and adaptive timing methods for service clock recovery. While both adaptive and differential timing algorithms can maintain end-to-end T1 timing stability that conforms to ITU G.823 synchronization requirements, differential timing provides improved wander performance for applications that require greater long term timing stability or that are subject to high PDV.

While the T1 encapsulation and data streams are independent in transmit and receive directions, the two ends are coupled with respect to synchronization and T1 clock recovery. For example, if the remote subscriber end of the T1 pseudowire is configured for differential timing, then the network end must send RTP time stamps in the T1 frames and both ends must have a traceable clock reference.

Timing Modes

Three mechanisms are available with the C7/E7 and MDU ONT for T1 clock recovery:

- Adaptive
- Differential
- Loop

Timing (Clock Recovery) Guidelines

When developing end to end clock recovery strategies, keep the following information in mind:

- If both Network Side and Premises Side components utilize different timing, then Adaptive Clock Recovery must be used.
- GPON ONTs always get their PWE3 clock from the Passive Optical Network.
- By default, AE ONTs always get their PWE3 clock from the local oscillator. However, if the GE uplink is known to be synchronous (SyncE), the command "config sync set on" can be added to the configuration file resulting in PWE3 automatically switching to get its clock from the GE uplink (similar to ONT timing in GPON mode).
- Adaptive Clock Recovery is used in most networks and is sufficient to recover the T1/E1 service clock for most applications. When a network is susceptible to high packet delay variation (PDV), and a traceable clock is available at both PWE3 end point devices, the T1/E1 circuit can be configured to use DCR and will thus be highly immune to network-induced PDV. Most T1/E1 deployments do not require Differential Clock Recovery.

Note: The entire ONT must be configured for either adaptive or differential timing, which applies the same timing algorithm to all eight T1 ports on the ONT. For E7 deployments, the ONT supports only one PWE3 Clock Recovery mechanism (either adaptive or differential) on the same network.


Recommended Timing Combinations

Recommended Timing Combinations			
Central Office Endpoint (C7 T1 PWE Card)	Customer Premises Endpoint (E7 GPON or AE)	Benefits/Limits	Example Deployment Service
Loop	Adaptive	Preferred Timing Method for network when one of the endpoints includes an E7 GPON ONT. Loop/Adaptive timing environments assume both the C7 and the E7 are running their own clock source. In this environment, it is also assumed that the C7 is serving as the CO (timing source) and any attached ONTs are recovering the C7's timing. Used in instances where a more stringent synchronization interface is required (per G.824, Section 6). For example, ISDN networks may find this solution more consistent. Must be book-end solution (Calix MDU ONT at each end of the pseudowire). Includes 74xGE, 76xGX ONT models. Traceable timing to the same source needed end to end (assumed for GPON network, configurable via "config sync set on" command for AE). ONTs must be interconnected via an E7 GPON card (ensures GE and GPON sync) If separate E7 GPON cards connect each endpoint, then cards must be synchronized via BITS clock.	Any Synchronous service needing to meet stringent upstream T1 Tx Jitter specification. See diagram below for a sample architecture.
Differential	Differential*	Eliminates Packet Delay Variation (PDV) from clock recovery. Differential Timing is preferred (necessary) when high PDV exists on the network. Must be book-end solution (Calix MDU ONT at each end of the pseudowire). Traceable timing to the same source needed end to end (assumed for GPON network, configurable via "config sync set on" command in AE). ONTs must be interconnected via an E7 GPON card (ensures GE and GPON synchronization). If separate E7 GPON cards connect each endpoint, then cards must be synchronized via the same SOURCE BITS clock.	Any Synchronous service needing to meet stringent upstream T1 Tx Jitter specifications. Using this combination requires both ends support Differential Timing and both ends use the same reference clock.
Adaptive	Adaptive	<i>Required Timing Method</i> if neither timing source is known. Easy to deploy - recommended for Unstructured (Hi-CAP) services. Reduces upstream under and overrun conditions. Allows the C7 to pass Northbound interface to handle jitter/wander more efficiently. Highly interoperable (compatible with all supported PWE3 endpoints) Optimizes the handling of jitter and delay variations. No special provisioning required at far end	Clear Channel
* - Allowable if both T1 endpoints have traceable timing to the same clock source. The differential mode common clock must be the same at each endpoint for differential clock recovery to operate properly.			

Proprietary Information: Not for use or disclosure except by written agreement with Calix.


The chart above indicates the Calix recommended timing combinations. The combinations shown below have been thoroughly tested and are known to be compatible with the Calix PWE3 architecture.

Note: Although other timing combinations are allowed, Calix recommends using Loop - Adaptive when the timing source is known and an E7 GPON OLT is deployed as part of the T1 solution.



Adaptive Timing

Adaptive Timing (or Adaptive Clock Recovery) takes the average difference in the arrival time of incoming frames and uses this average to maintain the timing. Using this method, frames are sent at regular intervals (from a timed T1 source) but arrive at random intervals after traversing the Ethernet packet switched network. The average arrival time is then calculated and used to time the outgoing T1 signal. This is possible since the source device is producing T1 encapsulated frames at a constant rate. By averaging the arrival times, the Packet Delay Variation (PDV) can be negated and the original sender's clock recovered.




Notes on Adaptive Timing

- Adaptive Clock Recovery must be used if no external reference timing source is available.
- When using Adaptive Clock Recovery, RTP can be enabled or disabled without affect.
- When using Adaptive Clock Recovery, RTP Timestamp mode must be set to Absolute.

Differential Timing

Differential Timing (Differential Clock Recovery) uses an external clock to help maintain timing. Differential timing also employs a residual timestamp at the sender based on the difference between the shared reference (traceable) time and the T1 line. Thus, the receiver utilizes the timestamps in the received frames against its own time reference to correctly adjust its own T1 line timing. When a shared clock is available between the two PWE3 devices (Calix MDU ONT or other interoperable device), this mechanism will produce lower jitter and delay.



To use Differential Timing, both PWE3 endpoints must use the same primary reference source. The ONT endpoints receive RTP timestamps and upon receiving these timestamps, the ONTs can process the difference between the timestamps as well as the actual packet arrival time. The ONT uses this difference to calculate the timing for the outgoing T1.


Notes on Differential Clock Recovery (Differential Timing)

- Differential timing, by design, uses the timestamp field in the RTP header to carry the shared reference "traceable" time across the pseudowire. For this reason, RTP must be enabled when using Differential Timing.
- RTP Timestamp mode must be set to *Differential* when using Differential Clock Recovery.
- If both T1 endpoints are deployed off the same E7 shelf, they can use the internal clock of the E7 as their primary reference source.
- If the endpoints are connected to different E7s, then both shelves must have BITS inputs connected to the same primary reference source.

Loop Timing

Loop timing takes the timing information received from a T1 and uses the same timing information to transmit the T1. Although Loop timing is not really a clock “recovery” mechanism, the received T1 service clock is “looped back” to provide the transmit T1 service clock. Typically, Loop timing is used in a central office or network side of the T1 service and either adaptive or differential timing is used at the remote (subscriber) side of the pseudowire circuit. As a result of the Loop, this mechanism completely eliminates all wander and jitter in the T1 service transmitted back toward the network T1 source, but may require larger jitter buffers, and hence delay, to accommodate this.

Note: Loop Timing should only be used at both endpoints if both T1s are timed from the same source.



ONT T1 Clock Source (E7)

The Calix ONT receives its traceable reference clock for T1 Differential Timing recovery from the synchronous GPON optical network connected to an E7 GPON linecard.

The E7 may be using a local Stratum 3 clock or an external BITS input as its timing reference. Up to ten E7 shelves can be connected together to share a local BITS input. A BITS Chaining Cable available from Calix includes connectors to attach the redundant BITS output pins from one E7 shelf to the redundant BITS input pins on the next E7 shelf, and so on.

Note: If all ONTs use the same clock reference method, then and only then can differential timing be used.

For GPON networks, the traceable timing chain consists of: BITS - > E7 with GPON card - > PON - > 766GX ONT

For AE networks, the traceable timing chain consists of: BITS - > E7 with GPON card - > GE - > 766GX ONT

In either of the above cases, the E7 GPON card synchronizes the OLT and GE ports to the BITS input and projects that timing out to the ONT.

ONT Timing

The ONT derives its PWE3 timing source from one of three places:

- A local oscillator
- The PON uplink for GPON
- The WAN uplink for AE

Since GPON provides synchronization from the OLT to the ONT, the ONT is synchronous to the OLT. By default, the ONT in GPON mode uses the PON uplink as its timing source (assuming the OLT contains a more accurate timing source than the ONT).

In AE, the WAN uplink does not necessarily provide synchronization (asynchronous). In this case, the ONT in AE mode uses the local oscillator as its timing source for PWE3. If the WAN uplink is synchronous (SyncE), the ONT can be switched to utilize the WAN uplink as its timing source for PWE3.

If the timing from the GPON/WAN is traceable to a known good BITS timing source, then it is possible to synchronize separate PWE3 endpoints and provide better end-to-end PWE3 service.

Encapsulation Methods

The Calix PWE3 solution provides two encapsulation methods for Ethernet Layer-2 packet transport:

- The Metro Ethernet Forum (MEF) Ethernet-only encapsulation provides a structured T1 circuit format. The local and remote ends of the circuit are given Emulated Circuit ID (ECID) values. The remote end of the T1 pseudowire can be addressed using either a MAC address or an IP address. Calix recommends using an IP address to allow the two end points to learn each other's MAC addresses dynamically based on ARP exchanges. In this manner, a failed ONT can be replaced, given the same IP address for the T1 service, and the service will restore automatically. When creating the PWE3 service with MEF transport, the Local and Remote ECID must be specified and the Remote MAC or Remote IP must also be specified.
- UDP-IP is used for the pseudowire multiplexing layer while the Calix MDU ONT uses Structure-Agnostic Time Division Multiplexing over Packet (SAToP) for encapsulation and transport of T1 services over the IP/Ethernet packet switched network. The UDP port of the local and remote PWE3 circuits can be set independently, in support for UDP de-multiplexing, as defined in RFC 5086. When creating the PWE3 service with UDP-IP transport, the Local UDP port and the Remote IP must be specified.


Note: The application of port numbering is ONT-wide and is based on the settings of the 1st pseudowire provisioned. All additional pseudowires on that ONT will be configured identically (using either the RFC 5086 or RFC5087 methodology).

Note: MEF and UDP-IP encapsulation can be mixed on the same ONT.

UDP Port Numbering

In a properly configured system, the combination of the IP address and UDP port number uniquely identifies the pseudowire (in other words, a uniquely identified pseudowire "label"). This information is associated with the physical T1 port and directs the pseudowire payload to the proper T1 (TDM) port. The UDP port number must be agreed upon by both ends of the pseudowire (must be the same at each end in order to identify the bi-directional T1 payload flow).

In the aggregation illustration below, notice that Device A and B are transporting the T1 across their internal UDP ports 1-4. This creates problems on the far end (Device C) in that each T1 at the MDU ONT will have its own unique UDP port (any of a possible 8 internal T1 ports).



5860

Device A
IP Address 172.16.0.1

- T1 Port 1 = Local UDP Port 50001 - Maps to Remote IP 172.16.0.3 UDP Port 60001
- T1 Port 2 = Local UDP Port 50002 - Maps to Remote IP 172.16.0.3 UDP Port 60002
- T1 Port 3 = Local UDP Port 50003 - Maps to Remote IP 172.16.0.3 UDP Port 60003
- T1 Port 4 = Local UDP Port 50004 - Maps to Remote IP 172.16.0.3 UDP Port 60004

Device B
IP Address 172.16.0.2

- T1 Port 1 = Local UDP Port 50001 - Maps to Remote IP 172.16.0.3 UDP Port 60005
- T1 Port 2 = Local UDP Port 50002 - Maps to Remote IP 172.16.0.3 UDP Port 60006
- T1 Port 3 = Local UDP Port 50003 - Maps to Remote IP 172.16.0.3 UDP Port 60007
- T1 Port 4 = Local UDP Port 50004 - Maps to Remote IP 172.16.0.3 UDP Port 60008

Device C
IP Address 172.16.0.3

- T1 Port 1 = Local UDP Port 60001 - Maps to Remote IP 172.16.0.1 UDP Port 50001
- T1 Port 2 = Local UDP Port 60002 - Maps to Remote IP 172.16.0.1 UDP Port 50002
- T1 Port 3 = Local UDP Port 60003 - Maps to Remote IP 172.16.0.1 UDP Port 50003
- T1 Port 4 = Local UDP Port 60004 - Maps to Remote IP 172.16.0.1 UDP Port 50004
- T1 Port 5 = Local UDP Port 60005 - Maps to Remote IP 172.16.0.2 UDP Port 50001
- T1 Port 6 = Local UDP Port 60006 - Maps to Remote IP 172.16.0.2 UDP Port 50002
- T1 Port 7 = Local UDP Port 60007 - Maps to Remote IP 172.16.0.2 UDP Port 50003
- T1 Port 8 = Local UDP Port 60008 - Maps to Remote IP 172.16.0.2 UDP Port 50004

In this case, the initial assumption is that UDP ports 1-4 are unique at device A and device B since these are separate devices that already have unique IP addresses. However, UDP port numbers 1-4 cannot be used twice at device C. Therefore, since device C requires unique UDP port numbers 1-8, this requires either device A or device B to utilize UDP port numbers 5-8. Care must be taken when aggregating pseudowires from multiple devices to a single device such that all the pseudowires are uniquely "labeled".

For additional information on configuring UDP port numbers, refer to the topic "Configuring ONT Service Ports" in the *E7 GPON Application Guide*.

T1 Pseudowire Planning Guidelines

Adhere to the guidelines listed below for a successful PWE3 network turn-up:

- Within a Layer-2 network, T1 services from an ONT are typically placed in a single VLAN for transport between the remote pseudowire devices and the IP gateway. Multiple remote locations may be split into separate VLANs for ease of network planning and operation.
- Transport of T1 services over a packet network requires a CoS priority and sufficient bandwidth allocation within the packet network to reduce latency and packet loss. T1 service VLANs should be given a high IEEE 802.1p priority "p-bit" value (for example, 5).
- Each Calix ONT is provisioned with a static IP address for the SAToP interworking function (an IP host). Each T1 pseudowire end point thus has a unique source and destination IP address corresponding to the ONT with a UDP port ID for the individual T1 circuit on the ONT.
- T1 pseudowires can be transported through an IP network or E7 Ethernet network. If two T1 pseudowire ONT end points are within the same IP subnet and Layer 2 broadcast domain (VLAN), they will discover each other using Address Resolution Protocol (ARP). The T1 service will not transit through the IP default gateway and no manual configuration is required for each end point to discover their respective destinations.
- If two T1 pseudowire circuits have IP addresses in different IP subnets, the data path between the end points must transit through a router external to the E7 network even if there is a direct Layer-2 path between the two ONTs confined within the E7 network.

Note: On the E7 platform in GPON mode, T1 services may not be directly connected from one E7 ONT to another E7 ONT on the same GPON OLT port. ONT T1 services on different GPON OLT ports may connect directly only if they are in the same IP subnet and broadcast domain.


- To connect a T1 pseudowire circuit from one ONT to another ONT on the same E7 OLT port, the VLAN broadcast domain for the T1 pseudowire must have a path to the default gateway router. The default router must also be configured to maintain the high CoS priority for the T1 pseudowire.

- Calix recommends keeping the number of pseudowire circuits per VLAN to a maximum of 100. Thus, a fully loaded C7 shelf with 20 DS1-12 cards would require at least 3 VLANs
- When deleting a PWE circuit, ensure that both ends of the connection are deleted to prevent standing LOS-M critical alarms in the network, which could impact other T1 circuits.

NOTE: When traffic shaping is used and bandwidth metering is applied to ETH-0 on a PWE-3 circuit, upon executing the "metershowinfo" command the rate limiting is applied to ETH-2. When this occurs, the added overhead on these PWE-3 packets will cause lost packets at the far end of the PWE-3 circuit.


Configuration Example

The example below details the basic steps required to build a pseudowire circuit between two endpoints.



E7 Synchronous Network Model I - Differential Timing


The diagram below illustrates a synchronous network using the E7 and Calix MDU ONT (766GX, 766GX-R, 767GX-R). This Differential Timing model assumes all E7's are synchronized using the E7's BITS IN/OUT pins located on the back of the unit. By connecting the timing at this location, you create a traceable network clock reference between the E7 and the ONT's. All E7 systems are tied to a local BITS lock or allowed to free run on the E7's Stratum3 local oscillator. The remote subscriber ONT should be configured for differential timing for best performance. The central office network can then use Loop (lowest jitter) or differential (lowest delay) clock recovery.



E7 Asynchronous Network Model - Adaptive Timing

The diagram below illustrates an asynchronous network using the E7 and Calix MDU ONT.

In this scenario, the absence of a traceable network reference clock dictates that adaptive timing be used as the T1 service clock at the remote subscriber end. The CO network must use Loop timing to eliminate timing "wander" in the upstream T1 path. Since the E7 is a hardware based switch with inherent low latency, significant PDV is avoided between the T1 endpoints.



E7 Synchronous Network Model II - Differential Timing

The diagram below illustrates a synchronous network using the E7 and Calix MDU ONT. This Differential Timing model assumes the remote terminal and CO equipment share a traceable network clock reference.


Possible traceable network reference clock sources include:


- Synchronous Ethernet or SyncE

Note: The GE and GPON ports on the E7 GPON card are synchronous, deriving their timing from the same clock source. In addition, the GE and GPON ports can be synchronized to the BITS timing reference. Since the GE and GPON ports are synchronous, all ONTs connected to the E7 GPON card (AE or GPON) are synchronous. For PWE3, this can be used to synchronize the differential common clock between ONTs (a requirement of differential clock recovery). For AE ONTs, the command "config sync set on" instructs the ONT that the GE uplink is synchronous from the OLT.

- Timing over Packet (IEEE 1588 (v2) Precision Clock Synchronization Protocol (future release).
- BITS/GPS clock input (shown below)

In this configuration, differential timing is used to reconstruct the T1 service clocks with high immunity to PDV (a basic construct of a packet switched network over which the pseudowire is transported). The remote subscriber ONT must be configured for differential timing for best performance while the CO network side must use Loop (lowest jitter) or differential (lowest delay) clock recovery.





Chapter 6

GPON Subscriber Access

This chapter provides engineering and planning information regarding E7 GPON subscriber access.

Topics Covered

This chapter covers the following topics:

- Gigabit Passive Optical Network (GPON) Overview
- Supported GPON Equipment
- Configuring an ONT
- ONT Ethernet Port
- Fiber Cabling
- PON Protection

Calix E7 GPON Solution Overview

GPON overview

A Gigabit Passive Optical Network (GPON) is a point-to-multipoint fiber network as defined in the ITU G.984 family of standards. Each subscriber is connected into the optical network via fiber from a passive optical splitter. There are no active electronics in the distribution network, and bandwidth is shared from the feeder to the drop.

A GPON fiber-to-the-premises (FTTP) deployment consists of the following primary components:

- **OLT:** The optical line terminal (OLT) is the access node that provides GPON network terminations. This can be either a Calix E7 shelf equipped with a GPON line card, or an E3-8G that serves as the OLT.

Note: The GPON features of the E3-8G align with the provisioning, performance, and configuration characteristics of the E7-2 GPON-8 line card.

- **ONTs:** The optical network terminal (ONT) is the access endpoint that provides an optical termination at the subscriber premises, and provides subscriber interfaces (Ethernet, POTS, DS1, etc.). Home gateway or residential gateway (RG) functionality may also be integrated with ONTs, and associated with LAN-side Ethernet ports and WiFi SSIDs (if applicable).
- **ODN:** The optical distribution network (ODN) provides the optical connection between the OLT and ONTs. The ODN contains the following:
 - Optical fiber: The E7 supports a single fiber Class C+ (2.4 GPON) ODN with a reach of up to 60 km.
 - Optical splitters: To passively split the PON, located in the central office, remote terminal, outside plant, or any combination of the above. The GPON specification allows up to 64 splits per PON.
 - EDFA/combiner (for RF overlay): Third-party equipment that injects the RF video overlay signal into the PON and amplifies the signal.

E7 GPON solution

With a wire-speed Ethernet architecture supporting delivery of a full 2.5 Gbps to every GPON port, the Calix GPON cards provide capacity to meet the requirements of current and future needs for business and residential services. Multiple E7-2 systems can be linked together using industry standard 10GE SFP+ copper cables resulting in a very high-density, high bandwidth configuration of over 2500 potential subscriber endpoints in a 10RU space (up to 1:64 split). With software version R3.1 and higher, the E7-20 system supports a total of 10,240 subscribers.

Note: For information on configuring a set of Operation, Administration, and Maintenance (OAM) functions that monitor network operation, detect and localize network faults, and provide a measure of network performance within the Calix 760 ONT series, see "Configuring Ethernet OAM" in the *Calix E-Series (E7 OS) Maintenance and Troubleshooting Guide*.

The key attributes of the E-Series GPON solution are listed below:

- Uses the E7 GPON cards as GPON OLTs
- Also uses the E3-8G as GPON OLT
- Based on ITU G.984 GPON family of standards
- GPON: 2.488 Gbps downstream, 1.244 Gbps upstream
- GEM (Ethernet) based GPON
- Class B+ ODN
- Extended reach GPON
- Class C+ ODN (as per G.984.2 Amendment 2)
- Type B PON protection
- Interoperable with the following Calix access endpoints and gateways:
 - Calix 800 GigaCenters, GigaHubs, and GigaPoints
 - Calix P-Series 836GE RSGs, 700GE ONTs, and 700GX ONTs.
 - Calix T-Series T07xG HGU and T7x0G ONTs
- Integrated 10GE and GE/2.5GE aggregation and transport
- Integrated 10GE and GE transport and aggregation
- Ethernet data services support, including High-speed Internet access and MEF-compliant Transparent LAN Service (TLS)
- Voice services support, including ONT SIP and circuit-switched (GR-303, TR-8) services via a Calix C7 TDM voice gateway
- Multicast IPTV video services support
- CATV RF video overlay: 1550 nm video overlay and 1610 nm RF return
- Integrated ONT-based T1 leased line and private line services

Supported GPON Equipment

E7-2 GPON-4 Line Card

Interfaces:

- Four ITU G.984-compliant Gigabit Passive Optical Network (GPON) interfaces
- Eight Gigabit Ethernet (GE) interfaces
- Four ports of integrated 10-Gigabit Ethernet

The Calix E7-2 GPON-4 line card can be installed into one or both of the two universal slots within a Calix E7-2 shelf. Each PON supports up to 64 Optical Network Terminations (ONTs) for a total of 512 GPON and 16 point-to-point Ethernet subscribers per 1RU E7-2 chassis. The E7-2 GPON-4 card supports a full set of Ethernet services and network topology protocols on the Ethernet ports and can be used interchangeably with other E7-2 line cards to create a redundant system configuration.

Note: E7-2 supports 2.5 Gbps pluggable module interfaces on the SFP ports of the GPON cards. The 2.5GE interfaces support equivalent functions and networking protocols as GE and 10GE interfaces.

E7-2 GPON-4 r2 Line Card

Interfaces:

- Four ITU G.984-compliant Gigabit Passive Optical Network (GPON) interfaces
- Eight Gigabit Ethernet (GE) interfaces
- Four ports of integrated 10-Gigabit Ethernet

This card is backward compatible with the original GPON-4 line card. In addition, it provides C+ optics and forward error correction (FEC).

E7-2 GPON-8 Line Card

Interfaces:

- Eight ITU G.984-compliant Gigabit Passive Optical Network (GPON) interfaces
- Four Gigabit Ethernet (1GE/2.5GE) interfaces
- Two ports of integrated 10-Gigabit Ethernet

The Calix E7-2 GPON-8 line card can be installed into one or both of the two universal slots within a Calix E7-2 shelf. Each PON supports up to 64 Optical Network Terminations (ONTs) for a total of 512 subscribers per card. The E7-2 GPON-8 card supports a full set of Ethernet services and network topology protocols on the Ethernet ports and can be used interchangeably with other E7-2 line cards to create a redundant system configuration. The GPON-8 card supports single-fiber Class C+ (32 dB link margin) optical distribution network and forward error correction (FEC).

E3-8G Node-Based GPON OLT

The Calix E3-8G is an environmentally hardened, node-based GPON OLT that is optimized for cable operators. It leverages 60/90VAC cable plant power and is designed for the following OSP installations: strand-mount, pole-mount, wall-mount, and pedestal.

Interfaces and features:

- Same as the E7-2 GPON-8 card, except for form-factor and installation options

E7-20 GPON-4x Line Card

Interfaces:

- Four ITU G.984-compliant Gigabit Passive Optical Network (GPON) interfaces

The Calix E7-20 GPON-4x service access card provides multiservice capability over four GPON OLT ports, used to subtend up to 64 ONTs per OLT port, with a card capacity of 256 GPON ONTs. The GPON-4x card supports single-fiber Class B+ (28 dB link margin) optical distribution network:

- Standard Reach OIM
- Extended Reach OIM

E7-20 GPON-8x Line Card

Interfaces:

- Eight ITU G.984-compliant Gigabit Passive Optical Network (GPON) interfaces

The Calix E7-20 GPON-8x service access card provides multiservice capability over eight GPON OLT ports, used to subtend up to 64 ONTs per OLT port, with a card capacity of 512 GPON ONTs. The GPON-8x card supports single-fiber Class C+ (32 dB link margin) optical distribution network and forward error correction (FEC):

- Maximum 64:1 split per PON

ONTs can be distributed on any PON or card in any number within limits.

E7-20 GPON-16x Line Card

Interfaces:

- Sixteen ITU G.984-compliant Gigabit Passive Optical Network (GPON) interfaces

The Calix E7-20 GPON-16x service access card provides multiservice capability over sixteen GPON OLT ports, used to subtend up to 128 ONTs per OLT port, with a card capacity of 1024 GPON ONTs. The GPON-16x card supports single-fiber Class C+ (31 dB link margin) optical distribution network and forward error correction (FEC):

- Maximum 128:1 split per PON

ONTs can be distributed on any PON or card in any number within the limits.

Pluggable optics

The E7 platform supports industry standard pluggable modules for all service and network interfaces, including ITU G.984 compliant GPON, SFP GE, XFP 10GE ports, and SFP+ 10GE ports.

Note: The 10GE SFP+ port on the E7-20 SCP card does not support GE SFP modules.

See the following Calix Quick Tip bulletins on the Calix website for information on the support for SFP modules and direct-attach cables:

- *Calix E7 Pluggable Transceiver Module Support*
- *CAB-12-023 - Pairing Bidirectional SFPs (to Support Single-Fiber Ethernet Links)*

Module restrictions

The following module restrictions apply:

- Unless otherwise specified, SFP sockets in all Calix products accept any SFP module.
- GE-12 cards (E7-2) support any Calix-keyed 1GE SFP module in the SFP sockets.
- Copper SFP modules should be purchased from Calix to ensure correct operational behavior.
- The 10GE, 2.5GE, GPON OIM, and CSFP modules must always be purchased from Calix.

Calix ONTs

In a GPON system, ONTs provide the access endpoints and may also include home gateway functionality. The E7 GPON system is interoperable with the following Calix access endpoints and gateways:

- Calix 800 GigaCenters*, GigaHubs*, and GigaPoints
- Calix P-Series 836GE RSGs*, 700GE ONTs*, and 700GX ONTs
- Calix T-Series T07xG HGUs* and T7x0G ONTs
- Calix GPON Interface Adaptor with supported Ethernet Service Access Nodes (ESANs)

***Note:** These ONTs support Calix Home Gateway functionality.

Brief overviews of the supported ONT models are included below. For more detailed information, you can consult device-specific user documentation as well as the following matrices on the *Calix Resource Center*:

- *P-Series/T-Series ONT Software Matrix for E7 GPON*
- *Calix P-Series ONT Model/Feature Matrix*
- *Calix P-Series ONTs: Minimum Platform Software*
- *GigaCenter Firmware Matrix*

GigaCenters

The Calix **844G** and **854G** GigaCenters are residential premises service delivery platforms that extend the access network into the home and act as a strategic location for control of the gigabit experience. Supporting broadband connectivity within the home and managing subscriber voice, data and video services, this service platform integrates a 2.5 GPON optical interface with switching and routing functions that manage premises network traffic at speeds up to 1 Gbps. GigaCenters provide advanced home gateway functionality with four Gigabit Ethernet (GE) ports and carrier class 802.11ac Wi-Fi wireless networking for IPTV video and data services. GigaCenters also include two integrated voice lines supporting carrier grade VoIP and network-based TDM voice circuits, a USB port for home networking services, and an option for RF video (**854G**).

The Calix **844GE** GigaCenter provides the same functionality as the 844G, with the addition of an auto-detect optical interface supporting 2.5 GPON or 1.0 Gbps Active Ethernet (AE).

The Calix **844E** GigaCenter is an Ethernet home gateway, optimized for deployments that already contain access endpoints. It differs from the 844G as follows: Provides a 1GE WAN interface, and must be deployed behind an existing access endpoint; provides non-lifeline voice services; does not support an in-band management channel, such as OMCI; and does not include UPS telemetry.

GigaHubs

The Calix **812G** and **813G** GigaHubs are premises service delivery platforms that support broadband connectivity to the home. These devices integrate a 2.5 GPON optical WAN interface with switching and routing functions that manage premises network traffic at speeds up to 1 Gbps. GigaHubs provide advanced home gateway functionality with four Gigabit Ethernet (GE) ports and optional 802.11n Wi-Fi wireless networking (**813G**) for IPTV video and data services. GigaHubs also include two integrated voice lines supporting carrier grade VoIP and network-based TDM voice circuits, and a USB port for home networking services.

GigaPoints

The Calix **801G** and **803G** GigaPoints are indoor, 2.5 Gbps GPON, small form-factor service delivery terminals that deliver broadband connectivity to the subscriber. These GPON terminals provide a single Gigabit Ethernet (GE) port, optimized for providing broadband connectivity to a home gateway device for delivering IPTV video and data service. The **803G** also provides an integrated voice line supporting carrier grade VoIP and network-based TDM voice circuits.

P-Series 836GE RSGs

The Calix **836GE** Residential Services Gateway (RSG) is a service delivery device that supports broadband connections into the home and manages subscriber voice, data and video services. This device integrates an auto-detect 2.5 GPON or 1 Gbps Active Ethernet (AE) optical WAN interface with switching and routing functions that manage premises network traffic at speeds up to 1 Gbps. The 836GE provides home gateway functionality with four Gigabit Ethernet (GE) ports and 802.11n Wi-Fi wireless networking for IPTV video and data services. It also includes two integrated voice lines supporting carrier grade VoIP and network-based TDM voice circuits, and a USB port for home networking services.

P-Series 700GE ONTs

Calix 700GE optical network terminals (ONTs) are 2.5 Gbps GPON and 1.0 Gbps Active Ethernet (AE) ONTs that deliver a wide array of information, communication and entertainment services over fiber. A wide selection of 700GE ONTs are available, including SFU and SBU models. Select SFU models provide home gateway functionality.

P-Series 700GX ONTs

Calix 700GX optical network terminals (ONTs) are 2.5 Gbps GPON and 1.0 Gbps Active Ethernet (AE) ONTs that deliver a wide array of information, communication and entertainment services over fiber. A wide selection of 700GX ONTs are available, including SFU and MDU/SBU models.

T-Series T07xG HGUs

The Calix T07xG Home Gateway Units (HGUs) are a line of single-family service delivery devices that combine ONT functionality with Home Gateway capabilities. These Home Gateway Units integrate a 2.5 GPON WAN interface with Layer 3 switching and routing to efficiently manage premises network traffic. A wide selection of T07xG HGUs are available, with options for 802.11n Wi-Fi, USB port, RF video, and POTS ports.

T-Series T7x0G ONTs

The Calix T710G-24 and T720G-24 GPON Multi-Dwelling Unit (MDU) ONTs are indoor ONTs designed to provide up to 24 subscribers with HD quality IPTV, high-speed Internet services and optional voice services from a single ONT. The T710G-24 ONT has a 2.5 GPON WAN interface that supports 24 Fast Ethernet (FE) data ports. The T720G-24 ONT has a 2.5 GPON WAN interface that supports 24 ETSI Voice over IP (VoIP) lines and 24 Fast Ethernet (FE) data ports.

Calix GPON Interface Adaptor

The Calix GPON Interface Adaptor is a fully functional ONT that has been integrated into an SFP optical transceiver module. When inserted into the SFP port of a supported Ethernet Service Access Node (ESAN), it provides a GPON interface toward the network and a 1 Gbps Ethernet interface for the ESAN. This allows supported ESANs to be GPON-fed, utilizing existing E7 GPON access infrastructure.

Third-party equipment

The Calix E7 GPON solution may interface with the following third-party equipment:

- Passive splitters which passively split the PON
- Switches/routers for LAN/management access, Internet/ISP access, IPTV and Video on Demand (VOD) services
- Middleware and set top box for IPTV services
- Central office voice switch with a GR-303 interface
- EDFA/combiner which injects analog video into the PON
- RF Return premises equipment connected to the ONTs

Equipment Upgrades

Before configuring an ONT for a subscriber service, you must download, activate, and commit the applicable system and ONT software, if not already done. For more information about E7 system upgrades, refer to the following guides:

- This *Calix E-Series (E7 OS) Software Upgrade Guide* describes how to perform software/firmware upgrades as well as database backup and restore operations.
- The *Calix E-Series (E7 OS) GPON ONT Upgrade Guide* describes how to perform upgrades for compatible P-Series, GigaFamily, and T-Series GPON ONTs. This document contains a subset of the information in the *Calix E-Series (E7 OS) Software Upgrade Guide*.

E7 GPON Configuration Guidelines

Extended-reach guidelines

Due to the inherent optical link budget, ONTs deployed in the extended reach region are restricted based on the distance traveled.

Extended Reach PON Guidelines		
Number of Splits	Services	Distance
1 x 8	IPTV only	40 km
	IPTV with RF Video	39 km
	IPTV with RF Video and RF Return	38 km
1 x 16	IPTV only	30 km
	IPTV with RF Video	28 km
	IPTV with RF Video and RF Return	26 km
1 x 32	IPTV only	20 km
	IPTV with RF Video	17 km
	IPTV with RF Video and RF Return	15 km
All values shown above assume the following: • All deployments use (1) Extended Reach SFP connector somewhere in the PON. • All deployments use (1) connector at the ONT. • All deployments utilize (5) fusion splices somewhere in the PON. • For RF Video deployments, (1) CWDM is used. • For RF Return deployments, an additional CWDM is used.		

C+ Optics and FEC

Hardware-capable E7-2 and E7-20 GPON OLT cards will be able to support PON optical distribution networks (ODN) designed to comply with ITU-T G.984.2 Amendment 2 - GPON C+. This amendment provides an enhancement to ITU-T recommendation G.984.2 to support a new OLT optical interface specification (C+) that can enable a loss budget extension of 4 dB, to a maximum optical loss budget of 32 dB on each the 1490nm and 1310nm sides of the PON. This extension to the GPON physical media specification allows for a reach extension to up to 60Km in the ODN.

Loss budgets for the C+ extended GPON system

Items	Unit	Single Fibre
Minimum optical loss at 1490 nm	dB	17
Minimum optical loss at 1310 nm	dB	17
Maximum optical loss at 1490 nm	dB	32
Maximum optical loss at 1310 nm	dB	32
Maximum fibre length	km	60

Downstream & Upstream Forward Error Correction (FEC)

In order to fully comply with the network design framework outlined by GPON C+ (ITU-T G.984.2 Amendment 2) a method for correcting errors over a single-path optical transmission line is used called FEC; Forward Error Correction.

FEC is a software feature that can automatically be enabled on ports of a hardware-capable OLT. GPON ports in C+/FEC capable OLTs can be set to “auto” so that FEC is enabled automatically when a C+ GPON optical interface module (OIM) is inserted into the socket. Typically, FEC defaults to auto which results in FEC being disabled by default on the 20 km SFP and FEC being enabled by default on the 60 km SFP. The following OLTs are capable of supporting C+/FEC:

- E7-20 GPON-8x
- E7-2 GPON-8
- E7-2 GPON-4 r2

C+ optical interface modules can also be used in non-FEC capable E7 OLTs (E7-2 GPON-4 and E7-20 GPON-4x) to provide an extension of the B+ or Reach Extended B+ ODN reach, but FEC is required to comply with the full extent of C+ as defined in ITU-T G.984.2 Amendment 2, combining both the added optical power of the C+ optics, and the error correction algorithm provided by FEC. ODN design reach as a function of the OIM used and FEC:

	Split Ratio				
OIM Type	1 x 64	1 x 32	1 x 16	1 x 8	1 x 4
Standard B+	16 km	20 km			
Reach Extended B+	16 km	20 km	30 km	40 km	
C+ without FEC	19 km	28 km	33 km	43 km	
C+ and FEC	26 km	35 km	40 km	50 km	58 km

Home Gateway Overview

This section describes the Calix Home Gateway (HGW) components, concepts, and application methods.

Topics Covered

This section covers the following topics:

- HGW concepts
- HGW services
- HGW port assignments
- RG management modes
- HGW management interfaces

Cross-References

Overview: HGW Applications - External Mode

Overview: HGW Applications - Native Mode

Overview: HGW Applications - TR-069

Overview: SIP Gateway Service with Configuration File

Calix 800G GigaFamily Service Provider's Guide

Calix 844G/854G GigaCenter User's Guide

Calix 836GE RSG User's Guide

Calix 844E GigaCenter Service Provider's Guide

Calix 844E GigaCenter User's Guide

T-Series ONT Operation and Maintenance Guide

How to Determine HGW ONTs

You can determine which ONT models support home gateway operation by issuing the **show ont-profile** command from the E7 Command Line Interface.

A value of "1" in the "G" column of the CLI's ONT port types list, indicates that the model represents support for "residential-gateway" operation.

You can also find this information via CMS or the Embedded Web Interface (EWI).

In the navigation tree, select ONTs, then view ALL of the displayed ONT profiles to get the complete list of models that support the home gateway/RG function.

You can adjust the "Rows Per Page" count to view all of them at the same time.

Home Gateway Concepts

Calix Home Gateways (HGWs) logically partition an ONT into WAN and LAN IP domains, enabling routing and local bridging of data traffic on the LAN side to concurrently support both Layer 2 and Layer 3 services.

The logical components within an HGW are functionally similar to a standard GPON ONT with a subtended 3rd party residential gateway (RG). That is, the ONT receives tagged traffic from the downstream RG and sends it upstream single-tagged where the OLT may or may not add a second tag.


The following terms are used in relation to HGWs:

- **LAN interface** – The ports which are connected to subscriber devices. These ports are generally unmanaged and support layer-2 switching between them.
- **WAN interface** – The port which is connected to the access provider.
- **Internal interface** – This is another name for the LAN interface when talking in the context of NAT and IP address spaces. It represents an interface which is using a private/locally scoped address.
- **External interface** – This is another name for the WAN interface when talking in the context of NAT. It represents an interface which is using a globally scoped address.

Note: The USB feature is not available in this release.

When an HGW is provisioned, it operates with two partitions:

- **Gateway partition**
 - Local and wireless network
 - Home Gateway services
 - Management using TR-069 client
- **ONT partition**
 - ONT Management
 - Subscriber Services (IPTV, voice, data, RF video)



HGWs can operate using a single WAN IP address with an integrated DHCP component for dynamic IP address allocation to LAN devices. Set-top boxes, DVRs, computers, printers, tablets and other Wi-Fi enabled appliances can access the WAN via the HGW. Having Layer-2 bridging and Layer-3 routing provides support for Microsoft Mediaroom®, Minerva IPTV, and whole-home DVR systems.

With HGWs that support Wi-Fi, implementing a wireless in-home network simplifies networking and increases the number of supported devices by using the single WAN IP address service model. Employing the Home Gateway's embedded Point-to-Point Protocol (PPP) and DHCP for dynamic IP address allocation allows the gateway to support multiple Wi-Fi connected devices.

Management Options

HGW initial provisioning is done using a customer-defined RG configuration file, with the option of on-going management via Embedded Web Interface (EWI) or the TR-069 protocol. Unlike the P-Series support of HGW, the T-Series ONTs always require the download of an RG configuration file to set-up the routed WAN interfaces, as there is no "Native" mode provisioning.

The following options can be used to manage HGWs.

OMCI

HGWs support in-band and management and configuration of the interfaces and services compliant to ITU-T G.984.4. The system EWI and CMS are used to provision, manage and update the HGW services and image files.

TR-069

HGWs can be configured from an Auto Configuration Server (ACS) using the CPE WAN Management Protocol (CWMP) described in BBF TR-069, including Consumer Connect. HGWs comply with TR-098 and TR-104 data models. Residential Gateway (RG) configuration files can be downloaded or uploaded from Consumer Connect or to a third-party ACS using TR-069.

Embedded Web Interface

HGW software includes an Embedded Web Interface (EWI) application that provides access.

Related topics

- For ONTs that support Home Gateway, see *Supported GPON Equipment* (on page [476](#)) or product release notes.
- *HGW ONT Management Interfaces* (on page [500](#))

Home Gateway IPv6 Support

Calix P-Series, GigaCenter, and GigaHub Home Gateways (HGWs) provide the following IPv6 support for data services:

- Single-Stack IPv6
- Native Dual-Stack IPv4/IPv6
- Dual-Stack Lite (DS-Lite): IPv4 traffic encapsulated in IPv6 packets
- IPv6 Rapid Deployment (6rd): IPv6 traffic encapsulated in IPv4 packets

Note: In 6rd mode, the maximum throughput is limited to 100 Mbps.

Note: IPv6 modes are only configurable with the HGW RG interface in the external management mode.

Note: For detailed information on how to configure these IPv6 support modes, see the P-Series and GigaCenter documentation.

Calix T-Series HGU provide the following IPv6 support:

- Single-Stack IPv6
- Native Dual-Stack IPv4/IPv6

Note: For detailed information on how to configure these IPv6 support modes, see the T-Series documentation.

E-Series in conjunction with Premises support of Home Gateway is being enhanced to include direct support of IPv6 connectivity. IPv6 is the next Internet Protocol version to meet the expanding requirements for addressing. It is currently being used to supplement IPv4 but is expected to eventually replace IPv4.

IPv4 addresses are 32 bits, written in dot-decimal notation. IPv6 addresses are 128 bits long, written in colons-hexadecimal notation with eight groups of four digits. Direct connectivity of IPv6 negates the need for Network Address Translation (NAT) with each device having a unique IP address, and includes special addressing features and a significantly larger subnet space.

To help in the transition and implementation of IPv6 from IPv4 there are a number of different strategies to help operators depending on the network infrastructure and environment: Single or Dual-stack IPv4/IPv6, DS-Lite, and 6rd. All GigaCenter, and P-Series supporting Home Gateway Layer 3 services will support Single or Dual-stack IPv4/IPv6. Networks that require DS-Lite for IPv6 carriage (tunneling of IPv4) or 6rd for IPv4 carriage (tunneling of IPv6) must utilize the GigaCenter or P-Series supporting Home Gateway functionality.

The Home Gateway support of IPv6 will only support IPv6 for High Speed Internet (HSI) data services. The release does not support IPv6 for IPTV multicast video, voice services and TR-069 management. Note that only one variant of IPv6 support can be applied to a gateway, and only one service WAN interface can support the IPv6 variant which will be constrained to HSI only. The IPv6 interface can support IPoE Dynamicv6, IPoE Staticv6 and PPPoEv6.


To implement IPv6 support on a GigaCenter or P-Series supporting Home Gateway requires moving to External mode configuration with RG configuration file download, either via OMCI download or Consumer Connect. Support and provisioning of IPv6 will not be supported using Native mode which allowed for OMCI provisioning of routed WAN interfaces.

Note: The support of DS-Lite and dual stack 6rd on P-Series ONTs supporting Home Gateway will limit overall bandwidth throughput. GigaCenters will support these variants with 1 Gbps services but will not be line rate for small packet sizes. Calix will provide engineering guidelines to help Calix customers plan their networks based on their bandwidth and services requirements.

Home Gateway Single or Dual Stack IPv4/IPv6

Single stack IPv6 assumes a WAN interface will only connect using an IPv6 address. Dual Stack IPv4/IPv6 implements both connection types on an interface at the same time, subscriber devices can connect to either the IPv4 or IPv6 address protocol. The process is driven by DNS where a dual stack device will query the name of the destination, and if the response is a IPv6 address the device will send IPv6 packets. It allows the gateway to support simultaneous support of IPv4 and IPv6 content.

The dual stack IPv4/IPv6 implementation is shown in the below figure:



Dual stack IPv4/IPv6 is the most desirable variant of IPv6 support since it facilitates direct connections of both IPv4 and IPv6 devices and avoids complexities of tunneling, security, and timing delays that are introduced when translating between protocols required when using Carrier Grade NAT.


When supporting IPv6, the Home Gateway EWI has separate display of IPv6 statistics and packet performance. It does include support of firewall for IPv6 in same way it supports firewall for IPv4 with a general option of off/low/medium/high and ability to change traffic in/traffic out settings for protocols and ports. The firewall settings for IPv6 are managed separate from IPv4.

Home Gateway DS-Lite

With the depletion of IPv4 public addresses some operators have had to discontinue support of IPv4 in their networks and solely deploy IPv6 network infrastructure. Because not all subscriber devices support IPv6 it requires tunneling and translation of IPv4 addresses to the gateway.

The GigaCenter or P-Series supporting Home Gateway continues to distribute private IPv4 addresses on the LAN and wireless interfaces. DS-Lite encapsulates IPv4 packet inside a IPv6 packet with network termination to an Address Family Translation Router (AFTR) supporting Carrier Grade NAT with global IPv6 connection. At the AFTR the IPv6 packet is decapsulated, restored to IPv4, and routed to the public IPv4 Internet.

The DS-Line implementation is shown in the below figure:



To facilitate the tunneling of IPv4 packets the AFTR uniquely marks each traffic flow using the Gateway IPv6 address, the private IPv4 address and port number. The gateway obtains the URL of the AFTR via DHCPv6 (RFC 6334) or it can be provisioned manually with the AFTR URL via EWI, TR-069 or RG configuration file.

On its WAN side Network Address and Port Translation (NAPT) is disabled and the IPv4 tunnel becomes the default IPv4 route. Via DHCPv4 the gateway can either advertise itself as the DNSv4 server or advertise DNSv4 servers provisioned via EWI or TR69 or RG configuration file. In the former case the gateway will proxy A record queries from IPv4 to a WAN-side DHCPv6 server.


As noted the GigaCenter and P-Series supporting Home Gateway only support a single instance of DS-Lite on a routed WAN interface. The WAN interface is assumed to be supporting HSI services. Support of DS-Lite for HSI is independent of IPTV services, it is not supported for TR-069 management.

Home Gateway 6rd

For service providers with networks that do not have IPv6 infrastructure the GigaCenter and P-Series supporting Home Gateway will support dual stack 6rd. The variant of 6rd allows IPv6 service to be deployed over a pure IPv4 access network. The core network is not aware of IPv6, it does not require IPv6 infrastructure such as core routers, DHCP or DNS servers

The 6rd mechanism encapsulates IPv6 inside IPv4 between the Border Router (BR) and Customer Edge (CE). It follows all of the same IPv4 routing functions. On the GigaCenter or P-Series supporting Home Gateway the LAN interfaces appear as Dual-Stack IPv4/IPv6 to the LAN interfaces and subscriber.

The dual stack 6rd implementation is shown in the below figure:



At the subscriber location the gateway operates in a 'hub-and-spoke' mode with IPv6 tunneled traffic flows between the BR and gateway. The gateway can be provisioned to support 6rd by obtaining network data via DHCPv4 Option 212 or via EWI, TR69 or RG configuration file. The specific 6rd provisioning data consists of:

- IPv4 Mask Length
- 6rd Prefix
- 6rd Prefix Length
- BR IPv4 address

Provisioning of 6rd includes configuring the necessary parameters via EWI, TR-069 and DHCPv4, creation of the prefix, using the created prefix as a "delegated prefix" for purpose of including one of its /64s in RA messages, and modifying the IP header for traffic that goes between the WAN and LAN devices. Once configured for dual stack 6rd the gateway can either advertise itself as the DNSv6 server to LAN-side hosts or advertise DNSv6 servers provisioned via EWI, TR-069 or RG configuration file.

As noted the GigaCenter and P-Series supporting Home Gateway only support a single instance of 6rd on a routed WAN interface. The WAN interface is assumed to be supporting HSI services. Dual stack 6rd is not supported for IPTV multicast over routed interface and TR-069.

HGW Services

Calix Home Gateways (HGWs) can provide service delivery in a number of ways; for example:

- All subscriber services via a single VLAN on the RG interface.
- Data services via one VLAN on the RG interface, video services via another VLAN on the RG interface.
- Data services via a VLAN on the RG interface, video services via a VLAN on the full bridge (FB)#.

High Speed Internet

The first service VLAN (on the RG WAN interface) is required to be assigned to High Speed Internet services. It has additional attributes which define the management mode, WAN protocol and required configuration or credentials. The data service VLAN allows for provisioning the following:

- WAN Protocol as either DHCP, PPPoE, or Static IPoE
- DHCP will boot automatically using the designated IP Host
- PPPoE includes the user credentials Username and Password
- Static IPoE includes the set-up parameters of assigned static IP address, subnet mask and gateway

IPTV Video

HGWs support multicast and unicast IPTV services on either the routed/RG WAN interface or L2 bridge ports provisioned as either Half-Bridge or Full-Bridge# mode. When the video profile is assigned to the L2 bridge, a single IGMP Snoop function manages the multicast flows across all Half-Bridge ports or across the Full-Bridge group. When the video profile is assigned to the routed/RG WAN interface, it supports both IGMP Snoop and Proxy. All modes support Multicast VLAN Registration (MVR) as part of the video profile.

HGWs receive only one copy of a multicast stream on the GPON uplink regardless of the number of subscriber set-top boxes requesting the stream; the HGW replicates the multicast streams on a per subscriber basis. HGWs support multicast traffic switching and routing between Ethernet ports with up to 16 set-top boxes or media clients connected behind the LAN ports with up to 40 service flows per ONT. HGWs manages multicast joins and leaves, with the support of fast leaves.

Note: IPTV requires QoS enforced by the Home Gateway model.

Note: On the E7, T-Series ONTs do not support FB (Full-Bridge) mode. T-Series ONTs are compatible with the E7, but not with the C7.

Voice and Fax

On HGWs, voice services are not resident on the gateway partition, but remain on the ONT partition.

Voice services are supported as stated per HGW device datasheet and depending on the type of OLT (C7, E7, etc.) and software version.

Note: Voice services are managed outside of the gateway partition with no connection to the Home Gateway Profile or TR-069 provisioning.

Wireless LAN

HGWs that supports Wi-Fi can be used as Wi-Fi Access Points (APs) to create an in-home Wireless LAN (WLAN).

Home Gateway Port Assignments

This topic describes how ports are assigned to interfaces on Calix Home Gateways (HGWs). An HGW Ethernet port can be a member of the RG (Residential Gateway) interface group, FB (Full Bridge) interface group, or neither, allowing it to operate as a normal L2 classic port (Half-Bridge mode).

Interface Group Assignments

Interface group	Description
RG (Residential Gateway)	The RG services are named as Gateway Services, where only four services are supported and only one of them can have a multicast profile assignment. Supported service-tag actions: - change-tag - add-and-change-tag
FB (Full Bridge) ⁽¹⁾	With Full-Bridge mode, traffic can pass between LAN ports within the same bridge. Support of Full-Bridge is assumed to be for video multicast services only to enable Whole-Home DVR support, allowing communication between set-top boxes. A single multicast VLAN can be associated to the Full-Bridge group. The Full-Bridge interface only supports 1 service. Supported service-tag actions: - change-tag - add-tag
L2 Bridge ⁽²⁾	Defined by a special VLAN within the RG interface group, the L2 Bridge provides the same functionality as the Full Bridge, but includes the IPTV SSID as a member of the bridge.
<none> (Half Bridge) ⁽³⁾	Ports not associated with an RG or FB are in "Half-Bridge" mode and operate as L2 classic ports. This operation allows traffic to only pass between the WAN and LAN ports and restricts traffic from passing between LAN ports within the same bridge.
Notes: 1: Not supported on 844E GigaCenters (Ethernet Gateways) or T-Series HGUs. 2: Not supported on 700GE ONTs, 836GE RSGs, GigaHubs, or T-Series HGUs. 3: Not supported on 844E GigaCenters (Ethernet Gateways).	

- **When a GigaCenter, GigaHub, or 836 RSG ONT is created**, all ONT Ethernet ports are assigned as members of the RG group. Ports can be removed from the RG group and added to the FB group as needed.
- **When a P-Series or T-Series ONT is created**, the ONT Ethernet ports are not members of any interface group, and can be added to an interface group as needed. 700GE ONTs reset whenever ports transition into or out of RG mode. For example, when the first ONT Ethernet port is added to the RG group on a 700GE ONT, the ONT resets. Similarly, when the last Ethernet port is removed from the RG group, the ONT resets.

HGWs can support the following:

- A single instance of RG/routed mode
- A single instance of Full-Bridge mode for IPTV services
- Up to four instances of Half-Bridge mode for standard L2 bridged services

The use case of mixed High Speed Internet services with bandwidth metering and shaping on the Routed WAN interface and L2 bridge ports is not supported, and Full-Bridge support is only for video without bandwidth metering.

Note: HGWs with Wi-Fi always have the RG interface group enabled. On such HGWs, even if all ONT Ethernet ports are removed from the RG interface group, the Wi-Fi is still a member of the RG group.

Note: Upon turn-up of a HGW, if an IPTV client (STB) is attached to an ONT Ethernet port that is a member of an RG interface group, it would obtain a lease from the embedded DHCP server. Then, if the ONT Ethernet port is subsequently made a member of an FB interface group, the STB retains the IP address.

Examples of port assignment configurations

Port Assignments	Description
1 RG + 3 FB	(Typically deployed for RG HSI/Minerva IPTV services) • HSI is carried on a VLAN via the routed WAN interface • Video service is carried on a separate VLAN on the FB • GE-2–GE-4 ports are associated with FB • WiFi * and GE-1 port are associated with the routed WAN interface • IGMP proxy on the WAN interface is disabled
RG only	• Both video and data services are carried on the same single VLAN via the routed WAN interface • All GE ports and WiFi port * are associated with the routed WAN interface • IGMP-Proxy on the WAN interface is enabled
1 RG + 2 FB + 1 HB	(Typically deployed for TLAN services) • HSI service is carried on a VLAN via the routed WAN interface • WiFi * and GE-1 port are associated with the routed WAN interface • IGMP Proxy on the WAN interface is disabled • Video service is carried on a separate VLAN on the FB • GE-2/GE-3 ports are associated with the FB • GE-4 port is associated with the HB with a 3rd VLAN

* For Wi-Fi enabled HGWs only.

CLI examples of adding & removing ports from an interface

Interface group	CLI	How to add and remove ports
RG	'G1'	add ont-port 1/g1 to-res-gw remove ont-port 1/g1 from-res-gw
FB	'F1'	add ont-port 1/g1 to-full-bridge remove ont-port 1/g1 from-full-bridge
<none> Half-Bridge	'g1'	

ONT template in CMS

Using CMS R12.1 or later, you can create an ONT template that defines all of the key management aspects of a Home Gateway ONT as well as the logical assignments of physical GE ports for the initial setup and whether the RG configuration supports Native Mode versus External Mode. Once the template is created, you can apply it to the applicable ONT. The template can also be shared between AE and GPON ONTs.

There is a system-default template that applies to ONTs implicitly until you specifically change it on an ONT. Once a template is applied to an ONT, any further changes to that template are not reflected on the ONT. Any new ONTs provisioned with the modified template will use the new settings. You can always override the parameter values at the ONT level. The ONT template is a CMS object, only. See the following topics for instructions on using an ONT template in CMS:

- Creating an ONT Template
- Applying an ONT Template

RG Management Modes

On supported Calix Home Gateways, the ONT RG interface can be managed in one of the following modes:

Note: P-Series GPON Calix Home Gateways include the GPON GigaCenters, GigaHubs, 836GE RSGs, and 700GE ONTs.

- **Native** (P-Series only) is the default behavior where the CMS/E7 directly provision up to 4 RG WAN interfaces. However, voice and legacy L2 services are still natively supported for non-RG ports. In this mode, both the RG and ONT partitions are controlled by the CMS/E7.
 - Native mode is the default for all P-series RG ONTs.
 - With the Native management mode, you can provision the following:
 - Create up to 4 RG WAN interfaces.
 - Specify the RG WAN interface protocol: IPoE DHCP, IPoE static, and PPPoE. (WAN protocol is only configurable on the RG WAN interface associated with the Ethernet service Data1. RG WAN interface for the Ethernet services Data2, Data3, and Data4 are DHCP only.)

- Enable IGMP on a single RG WAN interface (video service)
- Enable MVR on a single RG WAN interface (video service)
- In the N:1 VLAN model, all RG WAN interfaces are on the same VLAN.
- In the 1:1 VLAN model, each RG WAN interface is uniquely identified by a VLAN. Typically, 1:1 VLANs are double-tagged, but, single-tagged VLANs are also supported.

Note: T-Series ONTs do not support Native Mode as an option for the Management Mode.

- **External** uses an external RG configuration file that is downloaded (via inband OMCI) to the ONT to fully manage the WAN interfaces on the RG. In this mode, the CMS/E7 do not provision the RG WAN interface on the RG partition. Except that the creation of an out-of-band (OOB) management VLAN is controlled by the E7 OLT.
 - The RG partition is controlled by the configuration file.
 - The ONT partition is controlled by the E7 OLT.

Note: When management mode = External, you must set a config-file-instance value. The T-series Home Gateway (HGW) ONTs only operate in the External mode, and require an RG configuration file for Gateway (Layer 3) services.

Using an RG configuration file, a service provider is able to define up to four Routed WAN interfaces supporting data, IPTV and auxiliary service VLANs as necessary. The RG configuration can explicitly define these service VLANs along with service static routes to bridge traffic on the LAN and wireless interfaces to the correct service VLAN.

Applying a new RG configuration file will reset the RG configuration on the HGW. If a subscriber has modified any of the items below via the EWI, these changes will be lost when the new RG configuration file is applied.

- Admin password
- Primary SSID
- Primary WPA/WPA2 key

Although a configuration file is applied to a Home Gateway system, the E7/CMS is still used to create the ONT Ethernet service for each VLAN, where the service match list used matches the RG WAN interface VLAN as specified in the configuration file.

Service tag actions for RG (Gateway) services

For RG specific services (as opposed to half-bridge [HB] or full-bridge[FB]), only “Change Tag” and “Add and Change Tag” actions are supported. The concept of provisioning this way was derived from the same topology that is observed today in deployments that utilize external non-Calix RGs. This is especially used in scenarios where a third-party RG is downstream of a regular GPON ONT. From a deployment perspective, the HGW ONT then supports two logical components; a regular GPON ONT and an RG. In this scenario, the ONT would receive tagged traffic from the downstream RG and the ONT would either change the received tag and send it upstream single-tagged or either change the received tag, send it upstream, and the OLT would then add the second tag.

T-Series RG configuration file creation

The RG configuration file requires an intake form be submitted to Calix with the defined configuration parameters. This results in the creation of an approved SIP configuration file from Calix that is then transferred to a server local to the E7.

P-Series RG configuration file creation

The RG configuration file must be created with the HGW embedded Web Interface (EWI), saved, and then transferred to an FTP server local to the E7. See Creating an RG Configuration File.

Use Cases for Management Modes

Native Mode (P-Series only)	
CMS, E7, and HGW ONTs	Consumer Connect, CMS, E7, and HGW ONTs
Service WAN VLANs are controlled by the OLT.	
- Default RG attributes are acceptable. - RG attribute manipulation prohibitively increases OPEX due to manually adjusting the RG attributes or loading the RG configuration files via the RG embedded web interface. - No centralized visibility or control of RG features.	- RG attributes controlled through an RG configuration file or RG profile via Consumer Connect. - Zero-touch turn-up. - Centralized visibility and control of RG features.

External Mode	
CMS, E7, and HGW ONTs	Consumer Connect, CMS, E7, and HGW ONTs
- RG VLAN creation is controlled through a configuration file. - ONT partition is controlled by OLT. - Service-tag-actions and OLT transport VLANs. - Zero-touch turn-up.	
- RG configuration file distribution occurs over the GPON OMCI. - No centralized visibility or control of RG features.	- RG attributes controlled through an RG configuration file or RG profile via Consumer Connect. - Centralized visibility and control of RG features.


ONT Template in CMS

Using CMS R12.1 or later, you can create an ONT template that defines all of the key management aspects of a Home Gateway ONT as well as the logical assignments of physical GE ports for the initial setup and whether the RG configuration supports Native Mode versus External Mode. Once the template is created, you can apply it to the applicable ONT. The template can also be shared between AE and GPON ONTs.

There is a system-default template that applies to ONTs implicitly until you specifically change it on an ONT. Once a template is applied to an ONT, any further changes to that template are not reflected on the ONT. Any new ONTs provisioned with the modified template will use the new settings. You can always override the parameter values at the ONT level. The ONT template is a CMS object, only. See the following topics for instructions on using an ONT template in CMS:

- Creating an ONT Template
- Applying an ONT Template

Configuration File Process Flow




HGW Management Interfaces

Calix Home Gateways (HGWs) have a variety of provisioning, activation, and management tools. Each tool offers unique features and access methods that together provide a comprehensive set of management tools.

Note: See the *Calix P-Series Home Gateway User's Guide* and the *Calix T-Series Operation and Maintenance Guide* (if applicable) for complete information on turning up and connecting to HGW interfaces.

The following example graphic shows a configuration where P-Series ONTs are automatically created on the GPON system using SmartActivate.



Management Options

The following methods can be used to manage HGWs.

- RG Configuration File via OMCI
- Home Gateway Embedded Web Interface (EWI), local or remote access
- TR-069 and Consumer Connect Auto Configuration Server

RG Configuration File via OMCI

The HGW management model supports an automatic download of an RG configuration file. This allows an operator to download and deploy common configuration settings across all HGWs. The scope of HGW settings controlled by the RG configuration file depends on the type of OLT (C7, E7, etc.).

Note: The RG configuration file download mechanism is via OMCI file transfer. Applying a new RG configuration file will reset the RG configuration on the HGW. If a subscriber has modified items such as Primary SSID, Primary PSK, or Admin Password via the EWI, these changes will be lost when the new RG configuration file is applied.

Home Gateway Embedded Web Interface (EWI)

HGWs provide an easy-to-use Embedded Web Interface (EWI) that is accessed locally or remotely using a web browser. The EWI provides access to the following functions:

- Retrieving the basic ONT information, such as ONT model, CLEI number, product serial number, and firmware version.
- Configuring services and features on the ONT, such as VoIP, Wi-Fi, CWMP, security and LAN network.
- Managing the ONT, such as rebooting, and account management.

The Home Gateway interface supports two levels of access:

- **EWI via Local Access**

The administer account provides the subscriber local access. For example, a service technician on site with the HGW can logon to its EWI to provision and activate services.

- Attach a PC to any of the RG LAN ports and allow the PC to obtain an IP address, which will be in the 192.168.1.0/24 subnet, by default.
- Launch a web browser and navigate to the following address and login:

http://192.168.1.1

username / password (GigaCenter, GigaHub, 836GE): **support** / **support**

username / password (700GE and T-series): **admin** / **admin**

Note: Hidden service provider pages can be found at the default URL:
http://192.168.1.1/html/support/.

- **EWI via Remote Access**

The remote management account allows the service provider to have remote access over the WAN interface. For example, an administrator at the data center Telnets to the HGW to enable Remote Management access to its EWI. It is important to note that the Remote Management account interface access is disabled by default.

Remote access to the T-series ONT RG interface is currently only supported on the WAN HSI interface, (not the IPTV RG, SIP or TR-069 interface).

1. Enable the remote user interface for a specified time.
 - a. On the Navigation Tree, expand the ONT, and then click **RG 1**.
 - b. In the Set Remote Access Time box, enter the amount of time to enable remote access (up to 1400 minutes/1 day), and then click **Apply**.
 - For CLI: **set ont-port <port-id> remote-access-time <minutes-value>**

At this point the remote interface is accessible through the IP address of the RG WAN interface.

2. Find the WAN interface IP address:
 - In the RG WAN Interface Status area, scroll to the right portion of the window to view the IP address.
 - For CLI: **show ont-port <id> details** (not supported for T-Series)
3. Launch a web browser and enter the WAN interface IP address, followed by the remote access port:
 - **P-Series example:** `http://10.10.200.157:8080`
username / password (GigaCenter, GigaHub, 836GE): **support / support**
username / password (700GE): **admin / admin**
 - **T-Series example:** `https:// 10.45.10.155:309`
username / password: **SuperAdmin / SuperAdmin**

Note: The P-Series default login credentials are stored in the applied RG configuration file, and can be modified through the EWI, Consumer-Connect (TR-069), or by loading another RG configuration file with modifications. To modify these default values for a T-Series RG configuration file, contact Calix Technical Support for assistance. The T-Series ONTs support RG remote access with releases E7 R2.3.20 and T-Series R4.1.38 and above.

For complete information on connecting to HGWs, see:

- *Calix 800G GigaFamily Service Provider's Guide*
- *Calix 844E GigaCenter Service Provider's Guide*
- *Calix 836GE RSG User's Guide*
- *T-Series Operation and Maintenance Guide* (if applicable)

TR-069 and Consumer Connect Auto Configuration Server

HGWs support the TR-069 standard protocol that provides an auto-configuration tool and remote management of CPE devices and CPE Internet access devices such as modems, routers, gateways, set-top boxes, and VoIP-phones.

Without modifying the applied RG configuration file, you can create a different management profile that specifies specific TR-069 management objects with associated parameter values and the ACS will update the appropriate ONT data attribute.

The TR-069 client can be configured to run in the following management modes:

- **In-band** runs on an existing WAN interface. Typically, the path to the ACS server utilizes an existing RG service, such as Gateway Service: 1 (HSI).
- **Out-of-band** runs on its own WAN interface. An independent VLAN outside of the service VLANs is provisioned on HGW ONTs to facilitate communication to the ACS service.

Calix Consumer Connect is an Auto-Configuration Server (ACS) that uses TR-069 standard protocol. All TR-069 compliant devices adhere to a well-defined data model that describes common TR-98 management objects for the CPE, allowing remote configuration, provisioning and device management.

1. CMS provisions the data service to the Routed WAN interface with Consumer Connect ACS URL and credentials.
2. CPE device is the connection initiator with the ACS and must be configured to make contact with the ACS in the following ways:
 - CMS System Settings (ACS Settings)
 - CLI Home Gateway command settings (config acs setup)
 - Home Gateway Embedded Web Interface (EWI) settings (administrative pages access only)
3. After an initial start, the HGW is discovered on the ACS using the bootstrap and INFORM process. The ACS identifies the HGW by device type and instructs it to upload a Service Defined RG configuration profile. The configuration is transferred to the HGW using the HTTP protocol and re-initialized to begin services.

Once a connection between the HGW and the Consumer Connect ACS is established, the CPE device is managed at two distinctly different levels:

- Device level operations – a single HGW
- Group level operations – a group of HGWs

ONT Configuration Files for HGW Services

The Home Gateway (HGW) applications are provisioned in a similar sequence to other GPON services, configuring the network uplinks, creating service-related profiles, and then configuring subscriber services. However, you also need to define all of the key management aspects of a Home Gateway ONT as well as the logical assignments of physical GE ports for the initial setup.

- When configuring Layer-3 HGW services in a GPON system, you use the OMCI download of an RG configuration file.

The P-Series in native mode provisioning allows the E7 to directly provision up to 4 RG WAN interfaces, without the need for an ACS or RG external configuration file.

- When configuring SIP Gateway service in a GPON system, you use the OMCI download of a SIP configuration file.

The P-Series can omit the SIP configuration file and base the service provisioning exclusively on the E7 SIP Gateway Profile.

This topic describes where to obtain the ONT configuration files and the process of applying the files to GPON ONTs.

- **T-series**

The SIP configuration file and RG configuration file is generated by Calix and must be requested via an intake form that you submit to Calix with all configuration parameters pre-defined. Calix then provides an approved configuration file that is then transferred to a server local to the E7. See *Completing Residential Gateway and SIP Configuration File Intake Forms*.

Note: Please allow two weeks for the generation and approval process.

- **P-series**

The RG configuration file must be created with the ONT embedded Web Interface (EWI), saved, and then transferred to a server local to the E7.

The SIP configuration file can be created based on the Calix VoIP template and sample file available with *Calix Application Note: Using the ONT VoIP Configuration File*, and then transferred to server local to the E7.

Note: The P-series and GigaCenters can also omit the SIP configuration file and base the service provisioning exclusively on the E7 SIP Gateway Profile.

Process

The operation is similar for retrieving and applying both VoIP and RG configuration files. However they can be distinguished from each other by the instance number that is specified.

- VoIP files use instances **9** through **16** exclusively and can be specified either by instance number alone or by "voip-<1 through 8>."

The SIP configuration-file instance is also specified in the E7 SIP IP Host applied to the ONT SIP voice services.

- RG files use instances **17** through **24** exclusively and can be specified either by the instance number alone or by "rg-<1 through 8>."

The RG configuration-file instance is also specified in the ONT template applied to the ONT or the ONT RG interface.

The following steps outline the process for retrieving and applying a configuration file:

1. **Retrieve** the configuration file from an external server that is reachable by the E7.
 - The "cancel" action can be invoked before the "apply" action is performed, either while the "retrieve" action is still in progress, or after the "retrieve" action has completed.
2. **Apply** the configuration file to the ONT, causing the ONT to reset and use the new file.
3. **Delete** the old ONT configuration file from the E7 OLT flash memory.

When you load a configuration file onto the E7, all ONT's in Configuration File Management Mode will automatically receive this file.

- You can also manually load the file on a single ONT by specifying a product code to identify a specific ONT when you want to test a new configuration file, or if you want to update the configuration file used by each ONT incrementally. The OLT will not download a new configuration file, unless the version number in the file is different from what is in the file currently on the ONT.
- You can also specify the product code to identify a group of ONTs, including a mix of model types.

Note: The ONT must be reset to factory default values before downloading new configuration files. This can be accomplished by either pressing and holding the reset button for 20 seconds or using the EWI. Applying a new RG configuration file will reset the RG configuration on the HGW. If a subscriber has modified items such as SSIDs, Admin password, Primary SSID, or Primary WPA/WPA2 key via the EWI, these changes will be lost when the new RG configuration file is applied.

The T-Series ONTs will reset to factory default settings during the "retrieve ONT configuration file" step if a prior non-default configuration file is present in the ONT. Perform the "apply ONT configuration file" step as soon as possible after the retrieve step to minimize service disruptions. Once the subscriber RG setting are provisioned, major changes to the L3 services are not allowed via XML Configuration files. (Changes are only allowed via WAN GUI settings on an individual ONT basis or via ACS/TR-69.)

The following is typical of external RG configuration conditions:


Download file over the OMCI from the E7		Download file via ACS	
• Management Mode:	External	• Management Mode:	External
• Config-File Instance:	17 - 24	• Config-File Instance:	none

Ethernet OAM on GPON 760GX ONTs

Calix fiber access solutions offer all the service assurance tools required for operators to sell Carrier Ethernet with service level agreements (SLAs). Integrating Ethernet OAM into fiber access ONTs has several benefits for service providers including:

- Eliminating CAPEX for external OAM network interface devices (NIDs)
- Reduced truck roll OPEX events via integrated test head RFC2544 “reflector”
- Simplified in-service reporting using Y.1731 performance monitoring tools
- Compliance with Ethernet OAM SLA requirements for Metro Ethernet services

Prior to integrating service assurance into fiber access, service providers built OAM “overlay” networks using external test heads and NIDs. Overlay networks complicate operations and increase the cost of commercial service delivery. The figure below shows an E7 GPON and point-to-point Ethernet fiber access network supporting Ethernet OAM on the 760GX family of ONTs.



As Ethernet networks replace the aging TDM infrastructure, integrated service assurance tools are becoming a mandatory requirement to ensure carrier-grade service delivery.

E7 R2.1 and above have the following Ethernet OAM features on the 760GX ONTs:

- ITU Y.1731 Performance Monitoring (PM) provides real-time feedback on SLAs for in-service Ethernet circuits (EVCs), enabling the provider to monitor and report on the service performance.
- IEEE 802.1ag Connectivity Fault Management (CFM) allows the service provider to determine if the circuit trouble is in the network or on the customer side of the Ethernet demarcation interface (UNI).
- RFC-2544 Reflector allows service providers to test a circuit’s performance and compliance prior to turning the circuit live. Integrated “reflector” capabilities on the ONT eliminate truck rolls and placement of a remote test head.

E7 R2.0 and above have the following Ethernet OAM features on the 760GX ONT's:

IEEE 802.3ah Link OAM enables monitoring of individual Ethernet links, link status, and link faults. Link OAM is also used to create local and remote loopbacks on the link between the ONT and the CPE device.


- Discovery – GPON ONT discovers and reports status of the far endpoint
 - Active / passive provisioning
 - Link OAM protocol state (up/down) reporting
 - Advertise local link OAM capabilities
 - Remote capability discovery
- Remote Fault Indication – messages from far endpoint are reported by the E7
 - Remote Failure; a receive link fault has been detected
 - Dying Gasp; loss of power - about to go down
 - Critical Fault; platform specific critical condition
- Remote Loopback
 - Initiate remote loopback - ONT can be configured to start or stop a remote loopback on the device connected to the ONT LAN interface.
 - A start command for an interface will generate a 'remote link in remote loopback' active event. A stop command for an interface that has placed the remote into loopback will generate a 'remote link in remote loopback' clear event.
 - PDUs received at the ONT from the remote endpoint will be discarded. By clearing the ONT Ethernet port statistics prior to the loopback test, the ONT transmit and receive frames can be compared to verify the link. A Remote loopback may also be coordinated with a test device that will inject a fixed number of test frames into the service VLAN.



- ① 802.3ah Loopback Start command sends OAM PDU to remote endpoint
- ② Remote endpoint (CPE) places link in loopback, returning received PDUs to ONT
- ③ E7 reports a "Remote Link in Remote Loopback Active" event


- Respond to remote loopback - ONT can be configured to respond to a loopback command sent by the device connected to the ONT LAN interface.
 - When the ONT interface receives a valid 'start remote loopback' OAM PDU, the ONT will place the interface into loopback and return all non-OAM PDUs to the remote end and discard locally. The E7 will generate a 'local link in remote loopback active' event for the ONT.

- When the ONT interface receives a valid 'stop remote loopback' OAM PDU, the ONT will remove the interface from being in loopback and return to normal operation. The E7 will generate a 'local link in remote loopback clear' event for the ONT.
- The ONT Ethernet port can be configured to enable/disable responding to remote loopback requests from the remote endpoint.




Ethernet OAM Domains

As traffic flows from one CPE to another, it may pass through several different administrative levels, also called domains. The customer domain reaches from one CPE interface to another. An example of this is a business service that connects remote sites. The service provider domain covers all of the network space in between the CPE. A service provider domain may consist of multiple operators. In the example below there are two operator domains: operators A and B. Each operator could be the same organization as the service provider, or it could be an organization that the service provider has contracted with to provide network services. A single device can belong to multiple domains.



Each domain has two endpoints, shown below as a triangle. A domain may also have intermediate points, shown as circles. A device, such as switch C or switch D shown in the diagram, can be an intermediate point for one domain and an endpoint for another. A device could also be an endpoint for more than one domain, as is the case for switch A and switch F.



Open Access Networks


As Calix continues its expansion into global markets, new requirements such as support of Open Access Networks have become an essential requirement imposed by governments or municipalities providing broadband infrastructure. An open access network is managed by the Network Owner / Network Operator and multiple independent service providers have access to the physical network to deploy services to subscribers. Subscribers can choose service packages or individual services from the provider of their choice.

Open access networks deliver world class broadband services at the lowest total Capital Expense and the lowest Operational Expense. Here are some additional advantages of single open optical distribution networks instead of a private optical distribution networks:

- Stimulates competition among service providers
- Competitive access providers can concentrate on marketing and customer service

Types of open access service providers:

- **Traditional carriers** - PTTs, etc.
- **Competitive providers** - Mobile operators can bundle broadband with the cellular service plans
- **Non-traditional providers** - Grocery stores, petrol chains, retailers can bundle broadband with their loyalty plans.



Open Access Example GPON Network Diagram

The above figure shows how a single physical access network can deliver services from multiple providers. Subscribers can choose which provider to use for a given service or take an entire service bundle from the provider of their choice. Notice in the example above that Provider A has one subscriber taking data only while another subscriber is taking voice, video and data. The services offered by Provider A are separate from the services offered by Provider B and Provider C. A given subscriber may choose to take any combination of services from different providers or from the same provider, it is their choice.

Cost Benefits of Open Access Networks

Traditionally, service provider networks are vertically integrated. Each service provider is responsible for their own optical distribution network and the associated construction costs for routing fiber to each potential customer location. In addition, each of the providers must deploy access and aggregation networking equipment to deliver the services offered.

In an Open Access network the physical facilities for delivering services are shared among the various providers that are delivering services. This includes sharing the Optical Distribution Network as well as the Optical Line Termination equipment and aggregation switches and routers.

How does it work?

There are a number of new software features introduced in release 2.4 to facilitate Open Access networks including:

- IGMPv3
- 2K/4K aggregate multicast addresses
- Multiplexed multicast GEM ports
- Overlapping multicast addresses

IGMPv3

In an IPTV network, broadcast television channels are delivered using IP multicasting. IGMP is the control mechanism used to control the delivery of multicast traffic. IGMP joins and leaves are sent from the IGMP host (the Set Top Box) to the upstream network to indicate which multicast group (which channel) is being watched. Access networks may use IGMPv2 or IGMPv3 to deliver IPTV services. IGMPv3 improves the interoperability between the access network and the set top boxes used to deliver video services. Many set tops come configured for IGMPv3 by default so having the option of using IGMPv2 or IGMPv3 allows the network operator to pre-provision the interfaces used to deliver IPTV video services to match the preferred operation of the set top manufacturer. In release 2.4, the Calix implementation of IGMPv3 does not support Source Specific Multicast (SSM).

IGMPv2 requires two separate messages when changing a channel, one to leave the current channel and another to join the new channel. IGMPv3 sends one membership report that includes both the leave and join information which reduces the amount of IGMP traffic required to change a channel, improving the overall performance of the channel selection processing.

The E7 system supports the following:

- IPv4 IGMPv3, as specified in RFC3376
- Sending IGMP General Query messages to the subscriber ports
- Proxy reporting on the upstream interface of the OLTs / DSLAMs

The E7 Video Services support the following:

- IGMPv3 Snooping on ONTs as per RFC 3376
- IGMP Snoop for both IGMPv2 and IGMPv3 protocols simultaneously coming from different ports
- IGMP Snoop for both IGMPv2 and IGMPv3 clients on the same VLAN, providing services to both

2K/4K Concurrent Multicast Addresses

In order to support multiple video service providers, the maximum number of concurrent multicast addresses has been extended in release 2.4. The previous limit was 800 multicast streams per system and the new limit is dependent on switch capabilities of the SCP or line card.

- The E7-20 supports 2000 concurrent multicast streams
- The E7-2 DSL supports 2000 concurrent multicast streams
- The E7-2 GPON supports 4000 concurrent multicast streams

Multiple Video Providers

In order to support multiple video service providers over the same access network, the E7 system supports 16 MVR profiles at the system level and one video provider per MVR profile. Each MVR profile specifies a multicast VLAN used for delivery of IPTV streams for that provider. Within a given service provider the multicast ranges must not overlap. The system allows multiplexing multicast VLANs over a single GEM port on the GPON interface so that multiple providers video streams can be delivered from a single ONT device.

Overlapping Multicast Addresses

Given that each video service provider is assigned a unique MVR profile and associated multicast VLAN, each service provider may use an arbitrary range of multicast addresses without concern for other providers that are sharing the access network. A given subscriber port, whether it be an ONT port or a DSL port, can only be associated with a single video service provider. The range of multicast addresses that are available on one subscriber port can overlap with the range of multicast addresses used on another subscriber port without conflict, provided that the multicast addresses are sourced from different multicast VLANs (different providers).

HPNA Services Deployment

Calix Model 712GE and 722GE ONTs are used in conjunction with HPNA-enabled devices to deliver IPTV and HSI over standard residential coaxial cable or a coaxial to twisted pair adapter.

Note: Twisted pair wiring is not compatible with HPNA deployments if using analog video services as well.

HPNA devices may include set-top boxes, stand-alone HPNA bridges, and Residential Gateways. HPNA services are provisioned identically to Ethernet services.

Calix ONTs with HPNA incorporate a coaxial F-connector port that connects directly to the home network's coaxial cable. These ONTs support both HPNA networking and traditional, one-way RF video.

Before deploying an HPNA network, the service provider must plan and characterize the signal attenuation caused by the in-home coaxial plant. Many options exist for deploying HPNA in a network including:

- **Point-to-point configuration:** a single HPNA device (ONT) connected directly to an HPNA device such as a set-top box or HPNA bridge.
- **Point-to-multipoint with no RF video:** a single HPNA device (ONT) connected to a splitter with multiple HPNA device outputs.
- **Point-to-multipoint with RF video:** a single HPNA device (ONT) connected to a diplexer and splitter with multiple HPNA device outputs as well as RF video signals over the same coax.

Refer to *CAB-13-034 HPNA 3.1 Home Networking Guide for IPTV, RF Video, Whole-home DVR, and Internet Access Applications* for details.

Creating GPON ONTs

ONTs are created on the network in one of two ways:

- Manually where the ONT is pre-provisioned by entering the ONT serial number into the Create ONT dialog box.
- Automatically where the ONT is discovered on the network after successfully completing the Remote ONT Activation Process.

Manually Creating an ONT (Pre-Provisioning)

This topic describes how to configure an ONT using the pre-provisioning and activation approach. Typically, the provisioning of an ONT is done in advance of installation and discovery of the physical ONT.

ONT Pre-provisioning

Create an ONT provisioning record in the E7 database by specifying an ONT Logical ID (or let the system assign the ID), the ONT profile, and then specifying either the ONT serial number or the registration ID and matching Vendor ID.

- **The ONT Global Logical Identifier** assigned by either you or the system. All provisioning related to the ONT (including ports, services, etc.) is indexed using this global identifier as a base.
- **The ONT profile** is named with the model number of the ONT (e.g. "711GX") and describes the numbers and types of ports associated with that ONT. When an ONT is provisioned, the system automatically creates provisioning for the appropriate ONT ports. If you need to create a custom profile for an ONT, see [Creating an ONT Profile](#).
- **The serial number** is the value associated with the specific ONT equipment and that will be reported by the ONT when it arrives and is discovered. The value is determined by the manufacturer and is found on the ONT and its packing.
- **The registration ID** is a user supplied value that is independent of the specific ONT equipment. This value is entered through a RONTA mechanism which requires that the provisioned record contain a registration ID and no serial number. See *Assigning a Registration ID using RONTA* (on page [518](#)). When the ONT arrives and is discovered, linking of a the discovered ONT to provisioning based on Registration ID requires that the Vendor ID match as well. Also upon arrival of the ONT, the serial number will be detected and included in the provisioned record.

ONT Arrival and activation

When an ONT is connected to a PON, the E7 automatically discovers the ONT and detects the serial number. The E7 system attempts to match the ONT to an existing “provisioned but not linked” ONT record in the database. This match may be done by serial number or by registration ID as described above.

When the match between a discovered ONT and a provisioning record occurs, the ONT is automatically linked to the provisioning record.

If there is no match between a discovered ONT and a provisioning record, you can do either of the following actions:

- Edit an existing provisioning record to match the serial number of a discovered ONT.
- Create a new provisioning record with a serial number that matches the discovered ONT.

Once a match is found, the system updates the provisioning record to indicate the ONT is linked to provisioning information, and to note the physical location (Shelf, Slot, GPON Port). If the match was made using the registration ID, the serial number reported by the ONT is added to the provisioning record. All pre-provisioning associated with the ONT (including ports and services) is sent to the ONT and the service is brought up as provisioned.

ONT Provisioning record modification

You can select an ONT listed in the table view to modify the provisioning record attributes, such as adding or changing the Registration ID or changing the profile, if the change is not service affecting.

ONT Removal

To remove an ONT that is present on the system, first unlink the ONT from the provisioning record and then delete it from the system. You can either change the ONT serial number on the provisioning record or directly issue an action to unlink the ONT. Once an ONT is unlinked, it will range on any E7 PON that is discovered. If the E7 system that it connects to has the ONT's serial number in a pre-provisioned or pre-existing ONT record, the ONT automatically links to the new system and services resume. Within a single E7 system, the single provisioning record in the database can be associated with any card in the system. See "Replacing or Installing Equipment" in the *Calix E7 Maintenance and Troubleshooting Guide*.

You can also replace an ONT by issuing an action or effectively remove a faulty ONT from a system by placing it in quarantine. See "Adding an ONT to Quarantine" in the *Calix E7 Maintenance and Troubleshooting Guide*.

Power shedding option

The GPON ONT profile does not include an option for "Disable on battery." Set the power option on each of the individual ONT ports.

Parameters

You can provision the following parameters for an ONT:

Parameter	Description	Valid Options
ID*	Identifies ONT by a global ID. For example, 10001. You must first select the User Assigned ID option to manually assign the ONT ID value. The default setting allows the system to assign the ID.	1-64000000
ONT Profile*	Select one of the default Calix ONT profiles, or select a custom profile that was previously created. See Creating an ONT Profile.	Any established ONT profile
Admin State	Administrative status of ONT. The 'enabled-no-alarms' state enables the ONT but suppresses alarms.	Disabled, enabled enabled-no-alarms
Serial#	ONT serial number that is a numerical portion of the FSAN Serial Number of the anticipated ONT. Select one of the following methods to pre-provision an ONT: - Serial# - Registration ID	1-8 Hex digits
Registration ID	ONT registration ID that is the RONTA identifier. Enter a value that you supply independent of the specific ONT equipment and is unique among all provisioned ONT records in the system. Select one of the following methods to pre-provision an ONT: - Serial# - Registration ID	1-10 alpha-numeric characters
Subscriber ID	ONT subscriber ID that matches back-office ID.	1-63 text string characters
Description	Description that labels the ONT location.	1-48 text string characters
PWE3 Profile	Name of the PWE3 profile to use from the list of available profiles. See Creating an ONT PWE3 Profile.	system-default Any established PWE3 profile
Template	(CMS object only) Select one of the created ONT templates that define all of the key management aspects of a Home Gateway (HGW) ONT as well as the logical assignments of physical GE ports for the initial setup. The template can also be shared between AE and GPON ONTs. There is a system-default template that applies to ONTs implicitly until you specifically change it on an ONT. Once a template is applied to an ONT, any further changes to that template are not reflected on the ONT. Any new ONTs provisioned with the modified template will use the new settings. You can always override the parameter values at the ONT level. See Creating an ONT Template.	Any established ONT template

*Required fields

To pre-provision an ONT before arrival

1. On the Navigation Tree, double-click **ONTS**.
2. Click the **Provisioned ONTS > Provisioning** tabs.
3. From the menu, click **Create**.
4. Reference the table above to configure the parameters.
5. Click **Save & New** if you are creating the provisioning record, or click **OK** if you are modifying an existing ONT provisioning record.
6. To delete an ONT provisioning record, do the following:
 - a. Click the **Provisioned ONTS > Provisioning** tabs.
 - b. In the provisioning record table, click the row of the record that you want to delete.
 - c. In the menu, click **Delete** and then click **Delete** at the dialog box prompt.

For CLI:

- `create ont <ont ID> [ip-host|profile|serial-number|reg-id|subscriber-id|description|ont-pwe3-profile|admin-state]`
- `set ont <ont ID> profile <p-name> [serial-number|reg-id|subscriber-id|description|ont-pwe3-profile|pon-port|admin-state]`
- `show ont [detail|unassigned|unassigned detail]`
- `enable ont`

To link an ONT to a provisioning record upon arrival

1. On the Navigation Tree, double-click **ONTS**.
2. Click the **Discovered ONTS** tab to view the ONTs that the system recognizes.
When a match between a discovered ONT and a provisioning record occurs, the ONT is automatically linked to the provisioning record.
3. In the Search/Filter, click **Not Linked** to view the discovered ONTs that are not linked to a matching provisioning record.
4. Click the row of a discovered ONT, and then from the menu, click **Action** and select the method to link the ONT to a provisioning record:
 - **Link to New Provisioning** allows you to create a new provisioning record with a serial number that matches the discovered ONT.
 - **Link to Existing Provisioning** allows you to edit the serial number in an existing provisioning record to match the serial number of the discovered ONT.
5. In the Link ONT dialog box, do the following:
 - a. In the ONT Profile box, do one of the following:
 - If you selected "Link to New Provisioning," either leave the System Assigned ID check box selected to automatically assign an index to the new provisioning record, or clear the check box and enter an ID.
 - If you selected "Link to Existing Provisioning," chose an index to one of the existing provisioning records that you will edit to match the discovered ONT.
 - b. In the Serial# box, the ONT serial number discovered by the system appears.
6. Reference the table above to configure the parameters.
7. Click **Link**.

For CLI:

- `create ont <ont ID> [ip-host|profile|serial-number|reg-id|subscriber-id|description|ont-pwe3-profile|admin-state]`
- `set ont <ont ID> profile <p-name> [serial-number|reg-id|subscriber-id|description|ont-pwe3-profile|gpon-port|admin-state]`
- `show ont [discovered|discovered detail|discovered on-pon-port|detail]`

Automatic ONT Creation (RONTA)

The Calix Remote ONT Activation (RONTA) feature simplifies the installation and turn-up of any Calix ONT of the specified equipment type. Commands are entered via butt set at the customer premises, which allows configuration of an ONT in a variety of environments, without assistance from the central office, or the need for special equipment.

The RONTA feature is used to reset the ONT to its factory defaults, locally provision service, or assign a registration ID as part of the pre-provisioning process. The RONTA feature is activated when the ONT is powered with the drop fiber disconnected. The ONT leaves the RONTA mode as soon as an optical signal is detected, leaving the Registration ID on the ONT disabled once a successful link has occurred.

Remote ONT Registration Rules

Follow these rules when using the remote ONT (RONTA) registration feature:

- RONTA can only be used on Calix ONTs.
- The ONT must have a physical POTS port in order to connect the butt set. The butt set must be connected to the first POT's port on the ONT.
- The Registration ID must be unique for each ONT on the network.
- The ONT must be powered off when connecting the butt set. Reapplying power initiates the RONTA sequence.
- The fiber pigtail must be disconnected from the ONT when using RONTA. Reconnect the pigtail once RONTA is completed and confirmed.

Note: The Registration ID must be unique per PON. Yet, if you are using the Registration ID to identify the gateway in Consumer CONNECT Plus, then the Registration ID must be unique in that Consumer CONNECT Plus partition.

To turn-up an ONT using RONTA

1. Disconnect power to the ONT.
2. Disconnect the network fiber (pigtail) from the ONT.
3. Connect an RJ-11 terminated butt set to the first (LINE ONE) voice port on the ONT.

Note: For butt set devices using alligator clips, back-out the Tip and Ring screws and clip the butt set leads to the T and R posts (black to Tip, red to Ring). Verify that the network fiber is disconnected from the ONT.

4. Reapply power to the ONT and wait for the ONT to start up.
 - For 700G or 700GX ONTs, listen to the butt set and wait until you hear a click sound (ranges from 10 to 45 seconds after power has been re-applied).
 - For 700GE ONTs, wait for the OFF HOOK LED on the ONT to start blinking (approximately 50 seconds for 2 POTS ONTs, 60 seconds for 4 POTS ONTs).

- For GigaCenter ONTs, wait approximately 150-180 seconds for the GigaCenter to complete its boot-up cycle. Voice Activate is ready to accept input when a clicking sound occurs.

Note: If digits are entered prior to the ONT being ready (click sound or OFF HOOK Blink), those digits are ignored. Also, if digits are entered *after* the ONT ready-time ends (this could be as short as 2 minutes from the time power was applied to the ONT), you must power cycle the ONT in order to get another ONT ready-time window.

5. Press "*", "*", "*" ,"0" (star, star, star, zero) keys on the butt-set key pad. The butt set sounds DTMF tones as the keys are pressed.

Note: The volume of the voice prompt is adjustable by pressing the asterisk (star) key. At each successive depression of the key, the volume increases (four distinct volume levels are available) while the voice prompt is repeated at the current volume level. Replaying of the voice prompt alone is not supported in this release. In other words, to replay the voice prompt (pressing the star key once), a corresponding change in volume will result. You can toggle through the four volume levels as often as necessary to return to the desired volume.

6. A voice prompt occurs: "Please enter [GPON or AE] Registration ID followed by pound".

Note: If you do not hear a message after entering the code, you may need to hang up the phone (butt set) and go off hook again (or perform a flash hook), and then re-enter the code and wait for the next voice prompt.

7. Enter a Registration ID of your choosing (up to 10 alpha-numeric characters in length that do not have to relate directly to the FSAN), followed by the # (pound) key.

Note: When using the RONTA technique for turn-up, make sure all Registration ID numbers are unique.

8. The voice prompts: "You entered 'xxxx'. If correct, enter 1, otherwise enter 0".
9. If you pressed 1, the voice prompt responds: "Registration ID saved".
10. If you pressed 0, the system returns to step 6 above.

Once complete, the following operational behavior is expected:

- The Registration ID is persisted in flash memory.
- The Registration ID is erased only when the ONT is re-set to factory default via a butt set.

With the Registration ID now embedded, pre-provisioning is allowed without knowing the ONT serial number.

Note: Reattach the network fiber (pigtail) to the ONT before securing the ONT enclosure.

Automatically Creating an ONT - Smart Activate

Calix has developed a family of Remote P-series ONT creation tools known as Smart Activate. Depending on the ONT being deployed, installers can activate ONTs in one of two ways:

- Web Activate is an enhanced version of the previously available Remote ONT Activation (RONTA) feature
- Voice Activate simplifies the installation and turn-up of any Calix ONT.

Web Activate

This web browser based ONT activation tool allows the field installer to use a laptop (plugged into the first Ethernet port on the ONT) to activate the ONT. New features available in this tool include a Static IP configuration option for ONT deployments as well as ONT firmware upgrade and factory reset options. Web Activate is currently available for use with the 740GE family of ONTs as well as the 716GE-I ONT.

The Calix ONT Web Activate feature is used to pre-register ONTs as they are deployed in the field. The application and activation process is similar to the butt-set Voice Activate feature.

During the ONT power-up boot process, the ONT enables the first Ethernet port and creates an IP host that accepts a static IP address of 192.168.100.100. The field technician connects the first Ethernet port to a PC and configures the PC Ethernet port to the static address 192.168.100.1 and the subnet mask 255. 255. 255. 0. The technician opens a web browser using the address <http://192.168.100.100> and is prompted for a username and password. The technician then enters the factory default credentials ("admin", "fiber1ttx") and is presented with the Smart Activate page.

Note: For Calix P-Series GPON Home Gateways (HGWs), the Ethernet port used for Smart Activate-Web connection must be a member of the RG-1 interface.

To initiate Smart Activate, enter the 10-character alpha-numeric Registration ID into the empty field. No special characters are allowed.

Upon clicking Apply, the ONT resets in order to apply the RegID.

Then, attach the active fiber pigtail to the bulkhead fitting of the unit.

Note: The GPON ONT profile does not include an option for "Disable on battery." Set the power option on each of the individual ONT ports.

Certain parameter entry fields are automatically shown or hidden depending on the selections made by the user. For example, all static IP configuration entry fields are shown only if the user has selected "Management IP Address Mode" or "Static".

Note: Web Activate is not supported on 700GX ONTs.

Voice Activate

The Calix Voice Activate feature simplifies the installation and turn-up of any Calix ONT. Commands are entered via buttset at the customer premises, which allows configuration of an ONT in a variety of environments, without assistance from the central office, or the need for special equipment. The Voice Activate feature can be used to reset the ONT to its factory defaults, locally provision service, or assign a registration ID as part of the pre-provisioning process.

Note: The Registration ID must be unique per PON. Yet, if you are using the Registration ID to identify the gateway in Consumer CONNECT Plus, then the Registration ID must be unique in that Consumer CONNECT Plus partition.

The Voice Activate feature is activated when the ONT is powered with the drop fiber disconnected. The ONT leaves Voice Activate mode as soon as an optical signal is detected, leaving the Registration ID on the ONT disabled once a successful link has occurred. Voice Activate is available for use with all Calix ONTs. See *Assigning a Registration ID to an ONT* (on page [521](#)).

The following GPON ONT Voice Activate commands are available:

GPON ONT Voice Activate Commands	
Function	Command Keystrokes
Master Reset for an ONT	***#
Assigning a Registration ID (GPON Mode)	***()

Assigning a Registration ID to an ONT

Butt set provisioning via the Voice Activate feature simplifies the installation and turn-up of any Calix ONT of the specified equipment type. Commands are entered via butt set at the customer premises, which allows configuration of an ONT in a variety of environments, without assistance from the central office, or the need for special equipment.

After using the management interface to assign a Registration ID and pre-provision services to the ONT, the Voice Activate feature is used to install and turn-up a Calix ONT at the subscriber site. Using a butt set, the craftsman enters the same Registration ID into the ONT allowing the installation and the service activation to be under the complete control of the installer at the subscriber site.

Note: The Registration ID must be unique per PON. Yet, if you are using the Registration ID to identify the gateway in Consumer CONNECT Plus, then the Registration ID must be unique in that Consumer CONNECT Plus partition.

Keep in mind the following when working with Voice Activate:

ONT Registration Rules

Follow these rules when using the voice activation feature:

- Voice activation can only be used on Calix ONTs
- The ONT must have a physical POTS port in order to connect the butt set. The butt set must be connected to the first POT's port on the ONT.
- The Registration ID must be unique for each ONT on the network.
- The ONT must be powered off when connecting the butt set. Reapplying power initiates the voice activation sequence.
- The fiber pigtail must be disconnected from the ONT when using this feature. Reconnect the pigtail once voice activation is completed and confirmed.

To turn-up an ONT using Voice Activate

Note: Ensure you listen to the entire voice prompt prior to pressing any keys on the butt set.

1. Disconnect power to the ONT.
2. Disconnect the network fiber (pigtail) from the ONT.
3. Connect an RJ-11 terminated butt set to the first voice port on the ONT.

Note: For butt set devices using alligator clips, back-out the Tip and Ring screws and clip the butt set leads to the T and R posts (black to Tip, red to Ring). Verify that the network fiber is disconnected from the ONT.

4. Reapply power to the ONT.
 - For 700G or 700GX ONTs, listen to the butt set and wait until you hear a click sound (ranges from 10 to 45 seconds after power has been re-applied).
 - For 700GE ONTs, wait for the OFF HOOK LED on the ONT to start blinking (approximately 50 seconds for 2 POTS ONTs, 60 seconds for 4 POTS ONTs).
 - For GigaCenter ONTs, wait approximately 150-180 seconds for the GigaCenter to complete its boot-up cycle. Voice Activate is ready to accept input when the PHONE1 LED comes on steadily when going off-hook.

Note: If digits are entered prior to the ONT being ready (click sound or OFF HOOK Blink), those digits are ignored.

5. Press "*", "*", "*", "0" (star, star, star, zero) keys on the butt set key pad. The butt set sounds DTMF tones as the keys are pressed.
 - For GigaCenter ONTs, the PHONE1 LED blinks at a 50% duty cycle indicating voice activate sequence has been initiated.

6. A voice prompt occurs: "Please enter [GPON or AE] Registration ID followed by pound".

Note: If you do not hear a message after entering the code, you may need to hang up the phone (butt set) and go off hook again (or perform a flash hook), and then re-enter the code and wait for the next voice prompt.

Note: The volume of the voice prompt is adjustable by pressing the asterisk (star) key. At each successive depression of the key, the volume increases (four distinct volume levels are available) while the voice prompt is repeated at the current volume level. Replaying of the voice prompt alone is not supported in this release. In other words, to replay the voice prompt (pressing the star key once), a corresponding change in volume will result. You can toggle through the four volume levels as often as necessary to return to the desired volume.

7. Enter a Registration ID of your choosing (up to 10 alpha-numeric characters in length), followed by the # (pound) key.

Note: When using the voice activation technique for turn-up, make sure all Registration ID numbers are unique.

8. The voice prompts: "You entered 'xxxx'. If correct, enter 1, otherwise enter 0".
9. If you pressed 1, the voice prompt responds: "Registration ID saved".
10. If you pressed 0, the system returns to step 6 above.

Once complete, the following operational behavior is expected:

- The Registration ID is persisted in flash memory.
- The Registration ID is erased only when the ONT is re-set to factory default via a butt set.
- The Registration ID is included in Option-61 searches during DHCP discover requests.
- If the DHCP server does not respond with a configuration file name, the ONT requests its configuration file using the following name: "<registrationID>.cfg"

With the Registration ID now embedded, pre-provisioning is allowed without knowing the ONT serial number.

Note: Reattach the network fiber (pigtail) to the ONT before securing the ONT enclosure.

If the volume of the recorded messages is not loud enough you can change the volume by pressing the asterisk key on the phone. This will change the volume and replay the last message at the new volume. You can do this until you get the desired volume. All further messages will play with this volume.

When the maximum volume is reached the volume will be set back to the lowest volume. You can continue pressing the asterisk to increase the volume. There are 4 volume levels.

Previously the asterisk just replayed the last message. You can not just replay the last message. It will always change the volume. If you want to just replay the message you will either have to accept the new volume or press the asterisk key several more times to get to the desired volume.

Web Activate

The Calix ONT Web Activate feature is used to pre-register ONTs as they are deployed in the field. The application and activation process is similar to the butt-set Voice Activate feature. During the ONT power-up boot process, the ONT enables the first Ethernet port and creates an IP host that accepts a static IP address of 192.168.100.100. The field technician connects the first Ethernet port to a PC and configures the PC Ethernet port to the static address 192.168.100.1 and the subnet mask 255. 255. 255. 0. The technician opens a web browser using the address <http://192.168.100.100> and is prompted for a username and password. The technician then enters the factory default credentials ("admin", "fiber1ttx") and is presented with the following page:

The Web Activate section of the page is shown in the lower panel, where the traditional Voice Activate (RONTA) and static IP configuration parameters can be entered by the user.

- Traditional Voice Activate (RONTA) parameters:
 - Registration ID (AE and GPON)
 - Management VLAN (AE)

Note: The GPON ONT profile does not include an option for "Disable on battery." Set the power option on each of the individual ONT ports.

Certain parameter entry fields are automatically shown or hidden depending on the selections made by the user. For example, all static IP configuration entry fields are shown only if the user has selected "Management IP Address Mode" or "Static".

Web Activate - DHCP Discovery



The image shows the Calix Smart Activate web interface. At the top, the Calix logo and 'Smart Activate' title are displayed. Below this, a section shows device information: Model (725GE), FSAN Serial (CXNK00040CD3), S/N (081108003971), Firmware (10.5.255.73), and MAC (00-06-31-42-b8-0e). Three links are provided: Smart Activate, Firmware Upgrade, and Factory Reset. A 'Help' link is also present. The main configuration section is titled 'Smart Activate' and contains a 'Submit' button, a 'Registration ID' input field, a '725GE Mode' section with radio buttons for 'Legacy' (selected) and 'SCTE RFoG', an 'Active Ethernet Only' section with a 'Management VLAN' input field (value 85), and a 'Management IP Address Mode' section with radio buttons for 'DHCP' (selected) and 'Static'. A second 'Submit' button is at the bottom. The number 6474 is visible in the bottom right corner.

Calix
Smart Activate

Model: 725GE
FSAN Serial: CXNK00040CD3
S/N: 081108003971
Firmware: 10.5.255.73
MAC: 00-06-31-42-b8-0e

[Smart Activate](#) [Firmware Upgrade](#) [Factory Reset](#)

[Help](#)

Smart Activate

Registration ID

725GE Mode:

- ☒ Legacy
- ☐ SCTE RFoG

Active Ethernet Only

Management VLAN

Management IP Address Mode

- ☒ DHCP
- ☐ Static

6474

Web Activate - Static IP Discovery

Smart Activate

Registration ID

725GE Mode: ☒ Legacy ☐ SCTE RFoG

Active Ethernet Only

Management VLAN

Management IP Address Mode ☐ DHCP ☒ Static

IP Address *

Network Mask *

Gateway *

Configuration File Server IP Addresses *

Configuration Filename ☐ MIC Mandated

CMS IP Address

NTP Server IP Addresses

NTP Time Offset (seconds) ☐ Daylight saving

Syslog Server IP Addresses

* denotes mandatory fields

6485

Resetting an ONT to the Factory Default

The E7 GPON ONT buttset master reset procedure returns the ONT settings to the factory default.

Another feature of Smart Activate is the ability for the technician or service provider to restore the ONT to the factory state through the “Factory Reset” page.

To reset the ONT to factory default

1. Disconnect power to the ONT.
2. Disconnect the network fiber (pigtail) from the ONT.
3. Connect an RJ-11 terminated buttset to the first (LINE ONE) voice port on the ONT.

Note: For buttset devices using alligator clips, back-out the Tip and Ring screws and clip the buttset leads to the T and R posts (black to Tip, red to Ring). Verify that the network fiber is disconnected from the ONT.

4. Reapply power to the ONT and wait for the ONT to start up.
 - For 700G or 700GX ONTs, listen to the butt set and wait until you hear a click sound (ranges from 10 to 45 seconds after power has been re-applied).

- For 700GE ONTs, wait for the OFF HOOK LED on the ONT to start blinking (approximately 50 seconds for 2 POTS ONTs, 60 seconds for 4 POTS ONTs).
- For GigaCenter ONTs, wait approximately 150-180 seconds for the GigaCenter to complete its boot-up cycle. Voice Activate is ready to accept input when a clicking sound occurs.

Note: If digits are entered prior to the ONT being ready (click sound or OFF HOOK Blink), those digits are ignored. Also, if digits are entered *after* the ONT ready-time ends (this could be as short as 2 minutes from the time power was applied to the ONT), you must power cycle the ONT in order to get another ONT ready-time window.

5. Press "*", "*", "*" and "#" (star, star, star, pound) keys on the buttset key pad. The buttset sounds DTMF tones as the keys are pressed.
6. A voice prompts to acknowledge your selection of requesting an ONT master reset.
7. After the confirmation prompt, press '1' to confirm the reset. A voice prompt replies, "ONT Master Reset is completed".
8. Press '0' to abort. A voice prompt replies, "ONT Master Reset is cancelled".
9. Continue entering new RONTA commands as required.
10. Re-connect the network fiber to the ONT and wait for the ONT to come on line.

To restore the ONT to the factory state through Smart Activate

1. From the Factory Reset" page, click the "Check here to confirm the factory reset" checkbox to confirm the intent to reset the ONT.
2. click **Reset** to return the ONT to a factory default condition.

Upgrading ONTs

Before configuring an ONT for a subscriber service, you must download, activate, and commit the ONT software, if not already done. For more information about E7 system upgrades, refer to the following guides:

- This *Calix E-Series (E7 OS) System Upgrade Guide* describes how to perform software/firmware upgrades as well as database backup and restore operations.
- The *Calix E-Series (E7 OS) GPON ONT Upgrade Guide* describes how to perform upgrades for compatible P-Series, T-Series, and GigaFamily GPON ONTs which does not rely on the OLT to distribute the software image to the ONT. The independent ONT upgrade allows you to upgrade the ONT without upgrading the OLT system software.

This document contains a subset of the information in the *Calix E-Series (E7 OS) System Upgrade Guide*.

Since the E-Series platform supports ONT firmware upgrades independent of OLT software upgrades, do the following before performing the ONT upgrade:

- Refer to the specific ONT platform release notes to determine the latest ONT firmware image supported by the E-Series OLT software currently running in the network. Upgrade the system software first, if necessary.
- Ensure the free flash space on the GPON card is adequate to accommodate additional files, or delete obsolete or unwanted ONT images to clear memory for additional images.

CMS Visibility to Software Images

CMS allows you to identify the following status conditions for ONT software images:

- Downloaded to the system and running in the network
- Downloaded to the system, but not running in the network
- Not present (deleted) on the system, but running in the network (software marooned ONTs)

ONT image release inventory indications

The following example shows the release management page where ONT images are identified in CMS.

PROVISIONED ONTs DISCOVERED ONTs QUARANTINED ONTs SW MAROONED ONTs PROFILES RELEASE MANAGEMENT CONFIG FILE MANAGEMENT							
System Version		APPLY CLEAR					
DETAILS		ACTION	REFRESH	Rows Per Page: 20	< > < Card 1-1 >		
ID	RELEASE NAME	VERSION	VENDOR	MODEL	PRODUCT	RUNNING	DOWN
1	2.4.1.65	10.8.0.15	CXNK	7.	SF	Y	Y
2	2.4.1.65	10.8.0.15	CXNK	7. GE.	S6	Y	Y
3	2.4.1.65	10.8.0.15	CXNK	8. GE.	S6	Y	Y
4	2.4.1.65	10.8.0.15	CXNK	74 GE.	S6	Y	Y
5	2.4.1.65	10.8.0.15	CXNK	74 GE.	B0	Y	Y
6	2.4.1.65	10.8.0.15	CXNK	76.	MD	Y	Y
7	CALIX_P-SERIES_11.0.2...	11.0.20.21	CXNK		P0	Y	Y

- **Running** indicates the OLT image is affiliated with the currently “running” E-Series software.
- **Downloaded** indicates the E-Series OLT was instructed to download the ONT image to the corresponding ONTs. For ONT images that are bundled with the E-Series software, this happens automatically. When performing an ONT-only upgrade, you must initiate this action.
- **Activated** indicates the E-Series OLT was instructed to activate the ONT image that was previously downloaded on the corresponding ONTs. When this is “Y”, then the corresponding ONTs shown are running that version of the image.
- **Committed** indicates the E-Series OLT was instructed to commit the ONT image that was previously activated on the corresponding ONTs. When this is “Y”, the corresponding ONTs have this image as the default version.

Note: When an image is “Activated” but not “Committed”, and the ONT reboots, the ONT initially falls back to its previous image. Committing the release for a single ONT upgrade is not valid, as this is a temporary state.

- **Flash Available** indicates the E-Series OLT flash memory that is still available for storing additional ONT images. You can delete obsolete or unwanted ONT images to clear memory for additional images. See Deleting and ONT Image from the System for details.

The CLI supports filtering the ONT releases per card via the “**show ont-release**” command.

ONTs without a corresponding image indication

The following example shows the software marooned ONTs page where a list of currently connected ONTs are identified for which there is no corresponding image present in E-Series flash.

PROVISIONED ONTS	DISCOVERED ONTS	QUARANTINED ONTS	SW MAROONED ONTS	PROFILES	RELEASE MANAGEMENT	C
System Version APPLY CLEAR						
CREATE DELETE ACTION ▾ REFRESH APPLY						
ID ▲	SYSTEM VERSION	VENDOR	MODEL	PRODUCT	COUNT	
2	24.1.65	ERSN	R1B	T0	1	

To ensure that all ONTs currently connected to the E-Series can be upgraded to run the desired version of ONT software, it is necessary that all ONT images are present in the E-Series flash. If there are a mix of different ONT models connected to the E-Series, you must perform “retrieve ont-release” to retrieve each of the needed ONT images.

The CLI supports filtering the missing ONT image releases per system via the “**show sw-marooned-onts**” command.

ONT Ethernet Port

Note: See *Calix Home Gateway Services* (on page [485](#)) for an overview on the Home Gateway (HG) modes and configuration.

Note: Upon turnup of a GigaCenter, GigaHub, or 836GE RSG ONT, if a CPE (STB) is attached to an ONT Ethernet port that is a member of an RG interface group, it would obtain a lease from the embedded DHCP server. Then, if the ONT Ethernet port is subsequently made a member of an FB interface group, the CPE retains the IP address.

ONT Template in CMS

Using CMS R12.1 or later, you can create an ONT template that defines all of the key management aspects of a Home Gateway ONT as well as the logical assignments of physical GE ports for the initial setup. Once the template is created, you can apply it to the applicable ONT.

Maximum Transmission Unit

The Maximum Transmission Unit (MTU) is a fixed value (read only) for an ONT Ethernet port and depends on the configured ONT:

- 700GX ONT Ethernet interfaces support a fixed 1600 byte MTU
- 76xGX and 700GE ONT Ethernet interfaces support a fixed 2000 byte MTU

Parameters

You can provision the following parameters for an ONT Ethernet port:

Parameter	Description	Valid Options
Admin State	Administrative state of the port. Note: For Ethernet ports that are members of an RG interface, the Admin State parameter is irrelevant, as the RG is considered under the control of the subscriber, not the service provider.	Enabled Enabled-no-alarms (default) Disabled
Subscriber ID	Subscriber ID information, such as name, address, or account number.	String up to 27 characters (blank) (default)
Description	Optional description field for the port or service.	String up to 27 characters (blank) (default)
DSCP/IP Precedence Map	Name of DSCP or IP-precedence to P-bit map to use on ingress. There are system-default profiles named "access" that you can use or you can create custom profiles and assign them to the interface.	Any previously-created map, or the system-default map

Parameter	Description	Valid Options
Interface	(Home Gateway ONTs only) Whether the ONT Ethernet port is assigned to an Residential Gateway (RG) 'G1' in CLI, Full-Bridge (FB) 'F1' in CLI, or not a member of either, allowing it to operate as a normal L2 classic port (Half-Bridge mode). Residential Gateway (RG) Mode The RG services are named as Gateway Services, where only four services are supported with only one of them can have a multicast profile assignment. The ONT RG interface supports up to 4 services. Full-Bridge (FB) Mode (P-series only) With Full-Bridge mode, traffic can pass between LAN ports within the same bridge. Support of Full-Bridge is assumed to be for video multicast services only to enable Whole-Home DVR support, allowing communication between set-top boxes. A single multicast VLAN can be associated to the Full-Bridge group. Applications that require Full-Bridge support of IPTV enabled devices along with media-clients with a single IP Host supporting data and unicast applications requires an auxiliary switch be installed behind the LAN port and continued support of Half-Bridge. The ONT Full-Bridge interface only supports 1 service. Half-Bridge (HB or Classic) Mode In Half-Bridge mode, the port operates as a normal L2 classic port. This operation allows traffic to only pass between the WAN and LAN ports and restricts traffic from passing between LAN ports within the same bridge. When a GigaCenter, GigaHub, or 836GE RSG ONT is created on the E7, all ONT Ethernet ports are assigned as members of the RG port. Ports can be removed from the RG group, and possibly added to the FB interface group. When a 700GE or T-series ONT is created on the E7, the ONT Ethernet ports are not a member of a HGW interface group, and can be added to an interface group. The 700GE resets whenever it transitions into or out of RG mode. For example, when the first ONT Ethernet port is added to the RG group on a 700GE ONT, the ONT resets. Similarly, when the last Ethernet port is removed from the RG group, the ONT resets. Note: For Ethernet ports that are members of an RG interface, the following parameters are irrelevant, as the RG is considered under the control of the subscriber, not the service provider: Admin State, Disable on Battery, Speed, Duplex.	FB, RG, blank selects half-bridge
Speed	Data rate of the Ethernet port in Mb/s. HPNA ports default to 100 Mb/s and cannot be changed. FE ports default to auto but can be changed to 10 and 100 Mb/s. GE ports can be changed to 1000. Note: For Ethernet ports that are members of an RG interface, the Speed parameter is irrelevant, as the RG is considered under the control of the subscriber, not the service provider.	auto (default) 10 100 1000
Duplex	Duplex mode for an Ethernet port. HPNA ports default to half-duplex and cannot be changed. FE ports default to duplex and cannot be changed. Note: For Ethernet ports that are members of an RG interface, the Duplex parameter is irrelevant, as the RG is considered under the control of the subscriber, not the service provider.	Full (default) Half
Link OAM Events	Whether to enable link Operation, Administration and Maintenance (OAM) event monitoring.	unselected (disabled) (default) selected (enabled)
GoS	Grade of Service (GOS) profile number to use on this port.	(select from list) (blank) (default)

Parameter	Description	Valid Options
Disable on Battery	Port operational state when the ONT is operating on battery backup power (during a power outage). You can also disable the ONT Battery-Present monitor from the entire ONT. Note: For Ethernet ports that are members of an RG interface, the Disable on Battery parameter is irrelevant, as the RG is considered under the control of the subscriber, not the service provider.	unselected (No) (default) selected (Yes)
Security Profile	Name of Security Profile to use.	system-default (default) Text string
Accept Link OAM Loopback	Sets the ONT Ethernet port to either accept or reject a received request in the form of an 802.3ah frame that is sent by a host, requesting that the receiving host put its port into loopback.	unselected (No) (default) selected (Yes)
DHCP Limit Override	For DOCSIS provisioning (with the Calix Compass: Open Link Cable vCMTS) to allow an override of the DHCP Lease Limit specified in a security profile applied to a port. This allows the same security profile to be reused and also allows the required DOCSIS provisioning where the lease limit is specified individually on the port on which the new service will be added. See the <i>Calix Open Link Cable vCMTS Command-Line Interface (CLI) Reference Guide</i> and <i>Calix Open Link Cable vCMTS SNMP Management Guide</i> for more information.	None (default) No Limit Enter Value (0-255)
Force 802.1x	An 802.1x supplicant attribute to force the supplicant to be unauthorized or authorized until the force attribute is set to none. Valid values: none, authorized, unauthorized.	None (default) authorized unauthorized
Role	For TR-167 devices where the ONU UNI port is a single UNI port attached to a trusted device. Role change requires that you delete and recreate any services provisioned on the port. uni - the ONT port will be facing untrusted subscriber equipment (for example RG, PC, STB). inni - the ONT port will be facing G.fast nodes where the following behavior occurs: - Provisioning of the MVR VLAN on the INNI port is allowed. IGMP is carried upstream on their own unicast gempport/VLAN used for multicast flows. - When a packet contains option 82, it is passed through intact. Note: The inni role only applies to ONT Gigabit Ethernet ports, ONT Fast Ethernet ports, and ONT HPNA Ethernet ports. This attribute can be changed only when there is no Ethernet service currently provisioned on the port. Only one single video VLAN per INNI port is supported. One PPPoE session per G.fast node port. Four DHCP sessions per G.fast node port.	uni (default) inni
PoE Priority	The ports can be prioritized (low, medium, high) for Power over Ethernet (PoE). If there is not enough power available to source all ports, the ONT drops the lower priority ports first.	low medium (default) high
PoE High Power Mode	The port can be classified to the type of Powered Device (PD) that will be connected to the port. Different classes of PD require different amounts of power.	disabled (default) enabled
Voice Policy Profile	Indicates the voice policy profile applied for this port that defines the VLAN ID, P-bit, and DSCP.	any existing voice policy profile
PoE Admin Status	Enables or disables Power over Ethernet (PoE) on the port. GPON will implicitly enable Link Layer Discovery Protocol (LLDP) on the ONT when either PoE is enabled or when the Voice Policy is defined.	disabled (default) enabled

Ethernet Security Profile

This topic shows you how to create an Ethernet security profile that is applied to ONT Ethernet ports to enable and/or specify security attributes of the ONT Ethernet port.

Note: The E7 implementation of security profiles applies to non-TLAN services only. For TLAN services, the L2CP Filter parameter must be set to all-tunnel.

- DHCP Lease Limit
- Upstream Broadcast/Multicast Limit
- L2CP Filter

Creating an L2CP Filter

This topic shows you how to create an Layer 2 Control Protocol (L2CP) filter that specifies which L2CP ranges to discard or tunnel L2CP protocol frames. The E7 already has two system default filters that either discard (all-discard) or pass (all-tunnel) all L2CP protocol frames for all ranges.

The L2CP filter is referenced from the Ethernet security profile that is applied to ONT Ethernet ports to enable and/or specify security attributes of the ONT Ethernet port.

Note: The E7 implementation of security profiles applies to non-TLAN services only.

Layer 2 Control Protocol (L2CP) handling


The E7 supports provisionable Layer 2 Control Protocol (L2CP) handling (tunnel, discard) on GPON ONT Ethernet interfaces and VLANs on the interface. The L2CP filtering is an MEF-driven requirement for business services that use BPDU tunneling. The L2CP feature supports per subscriber port configuration options to pass or discard the following L2CP protocols:

- **BPDU Range** - indicates L2CP protocol frames with destination MAC addresses 0x01:80:c2:00:00:00 through 0x01:80:c2:00:00:0f.
- **GARP Range** - indicates L2CP protocol frames with destination MAC addresses 0x01:80:c2:00:00:20 through 0x01:80:c2:00:00:2f.
- **All-LANs Range** - (Bridge Management Group) indicates L2CP protocol frames with the destination MAC address 0x01:80:c2:00:00:10.

Configuring Type B PON Protection

Type B PON protection provides redundancy of the optical distribution network between an OLT and a remote 2xN splitter. By enabling this feature, the E7 system can switch from the protected PON interface to the protecting PON interface within milliseconds of the detection of a failure in the protected path. Services are typically restored within 100ms of a fail-over to the protecting path. The solution is compliant with G.984.1 specification for Type B protection.

Applications that can take advantage of Type B protection include business Ethernet services, mobile backhaul services, dense residential MDUs or DPUs, home office SFUs, and any residential subscriber the service provider chooses to offer redundancy to. The protected applications benefit by minimizing outage times and providing end users with a higher quality of service. In most cases, service recovery from an outage is seamless and is not perceptible to the subscriber. In order to enable this feature, the service provider must define PON protection groups that identify a working (protected) PON port and a protect (protecting) PON port for each group. The protected port and protecting port must be on different line cards within the same E7 shelf, same E7 MC node, or same E7-20 chassis.



As the above figure shows, there are redundant fibers between the OLTs and the 2xN splitters. The fiber shown with the green bubble is the protected port. Upon a failure of the protected path, all of the ONTs below the splitter are rehomed to the protecting OLT port. This is usually done without reranging the ONTs on the protecting port. All services automatically move to the protecting path.

Configuration Guidelines

- Supported E7 platforms:
 - E7-2, stand-alone or MC node, with GPON4r2 or GPON8 cards (release R2.4.30 or greater).

Note: The GPON cards used in a protection group may be mixed.

- E7-20 with GPON-8x cards (release R2.4.30 or greater)
 - E7-20 with GPON-16x cards (release R3.1.10 or greater)
- Supported ONTs:
 - 800G GigaCenters
 - P-series 700GE ONTs, 700GX ONTs, and 836GE RSGs
 - T-series HGUs and MDUs
- 2 x N splitter requirements:
 - Must provide two inputs.
 - Must be compliant with Telcordia GR-1209-CORE and GR-1221-CORE specifications.
 - Split ratios up to 2 x 64 are supported.

Note: Typically, 2 x N splitters come in 2 x 8, 2 x 16 or 2 x 32 split ratios.

- Cabling requirements:
 - The protected port and protecting port must be on different line cards, but within the same E7-2 shelf, same E7-2 MC node, or same E7-20 chassis.
 - One input of the splitter must be connected to the protected port and the other input must be connected to the protecting port, routed diversely or non-diversely.

Note: For Type B PON protection, the number of optical fibers between OLTs and optical splitters is doubled.

- Scope of protection/redundancy:
 - The optical distribution network between the OLT and the remote splitter.

Note: This includes the GPON OLT cards and associated GPON ports, the optical modules OIMs, patch cords, connectors, splices and the fiber plant between the Central Office and the remote splitter devices.

- Reversion:
 - No automatic reversion of services once the protected path has been restored.
 - A forced switch from one path to the other is supported.
 - Services will also switch back to the protected port if there is a subsequent failure of the protected path.

- Switch times:
 - For facilities failures (voice and data services): < 100 ms*

*** Note:** If <= 32 ONTs per PON, < 100 ms for up to two FFP groups switched at the same time on the same card. If > 32 ONTs per PON, < 100 ms for one FFP group switched at a time on the same card. Otherwise, the switch times will be greater than 100 ms.

- For facilities failures (video services): < 3 seconds
- For equipment failures, resets, or upgrades: > 100 ms

Note: Switch times for an equipment failure require additional time to restore services due to the need for the ONTs to re-range before coming back into service. During the recovery from an equipment failure, it may take several minutes for all ONTs on a GPON line card to restore services.

Note: Switch times greater than 100 ms may result in ONTs being rediscovered.

- Operational notes:
 - An FFP group contains a maximum of 2 PON ports, and a PON port can only belong to 1 FFP group.
 - Upon creation of an FFP group, the GPON port settings of the working port are duplicated to the protecting port with one exception: Split Horizon settings are not duplicated; instead, they must be manually configured on the protecting port, if applicable.

Note: Split Horizon is not supported for GPON ports on MCE shelves.

- GPON cards with ports in an FFP group cannot be deleted, even if the "forced" option is used. The port must be removed from the FFP group before the GPON card can be deleted.
- Services are not impacted when creating or deleting FFP groups.
- During an FFP switch, existing voice calls will remain up.
- During an FFP switch, ONTs should only have to go through a re-ranging process, not a full rediscovery process. To the network operator, the ONTs should stay ranged/discovered.
- Alarms:
 - FFP Down (Major) – Raised when both the working and protection ports are unavailable.
 - FFP No-Protection (Minor) – Raised when either the working or the protection port is down.

Parameter Values

You can provision the following parameters for an FFP group:

Parameter	Description	Valid Options
Name	Fast Facility Protection (FFP) group name.	String up to 32 characters
Working Port	The default protected (or active) port in the FFP group. This port does not have to be OOS before creating the FFP group. It may be either IS or OOS.	Available GPON port
Protect Port	The default protecting (or standby) port in the FFP group. This port must be on a different GPON OLT card in the same E7 shelf, same E7 MC node, or same E7-20 chassis. This port must be OOS without any linked ONTs before creating the FFP group. Note: The first port to come up active after the creation of an FFP group is considered the ACTIV/IS port, regardless of whether it was identified as the working or protecting port.	Available GPON port

ONT Power Shedding

All Calix ONTs are managed via an ONT Management Control Interface (OMCI). An important element of OMCI is determining what services remain active and which services can be disabled when an ONT loses its primary power source. Known as power shedding, the E7 system allows you to conserve battery backup capacity when a loss of power is detected by provisioning the Disable on BATT field for the port. Power shedding is controlled at the ONT port level and can be set to determine whether services (ports) remain on or are disabled when an ONT goes on battery back-up power. The table below summarizes how power shedding is implemented in the E7 system.

Port Type	Default Behavior	Available Override
POTS	POTS circuits remain active when on battery (<i>Disable on BATT</i> is set to NO).	No override. POTS services remain on when on battery.
FE/GE or HPNA	Disabled after 5 minutes. Allows mission critical applications to be shut down gracefully (<i>Disable on BATT</i> is set to YES).	<i>Disable on BATT</i> can be set to NO - Ethernet Port services remain active indefinitely (or until battery capacity is depleted).
RF	Disabled Immediately. Typically, RF video not a mission critical application (<i>Disable on BATT</i> set to YES).	<i>Disable on BATT</i> can be set to NO - RF Video services remain active indefinitely (or until battery capacity is depleted).
T1	T1 circuits remain active when on battery (<i>Disable on BATT</i> set to NO).	No override. T1 services remain on when on battery.


Fiber Cabling

Standard single-mode fiber cables in buffer tubes must be fusion spliced to pigtails or jumper cables for physical protection whenever the fibers connect to electronics. In the field, the fibers can be fusion spliced directly to splitters to reduce total network insertion loss and total cost. This is not recommended when the splitters are collocated with the E7 system at the remote cabinet, as this central location is the common place for field technicians to run Optical Time Domain Reflectometer (OTDR) traces both at initial deployment (for benchmarking) and for subsequent troubleshooting.

PON Protection

In some service provider applications it may be desired to support protection of the GPON equipment and/or fiber facilities. PON protection may be used to protect a PON serving high-end residential / resort communities, mobile backhaul services, or business services with strict service level agreements.

The FSAN / ITU has defined the protection types: A, B, and C, providing varying degrees of equipment and fiber protection with increasing levels of equipment redundancy / cost. Type B is shown below and is the most popular. Type B protection uses 2xN splitters and provides feeder fiber protection with no additional PON optical loss, while providing 1:1 GPON OLT equipment protection. In this configuration, two GPON cards provide equipment protection in the same E7 system. In an alternative configuration, the two GPON cards could be in separate E7 systems.




Within the E7 GPON operations model, an ONT is created in the E7 database and is not associated with a specific PON until it is discovered (on the PON). Once discovered, the ONT-to-PON association will persist through system restart, software upgrade, LOS, or other transient events. Only the administrative process of “unlinking” the ONT from the PON will allow the ONT database record to go back to an unassigned state where it can be discovered on another PON.

Once an ONT is unlinked, it will range on any E7 PON that is discovered. If the E7 system that it connects to has the ONT's serial number in a pre-provisioned or pre-existing ONT record, the ONT will automatically be linked to the new system and services will resume. Within a single E7 system, the single provisioning record in the database can be associated with any card in the system.

Using the “unlink ONT” command, a simple PON protection scheme can be enabled using the following operational procedure:

1. A “loss of PON” alarm is raised by the E7 management system (typically within 15 seconds of fault).
2. Operator disables the failed OLT port or GPON card to ensure that the failure event is not transient.
3. Using the “unlink ONT” command, the operator unlinks all ONTs associated with the failed OLT port or GPON card.
4. The operator enables the standby OLT port(s). (The standby GPON card should always be enabled.)
5. All newly unlinked ONTs will range on the standby OLT ports. Configurations and services will automatically restore as each ONT re-registers with the E7 system. All ONT services should resume in less than 5 minutes.

The above operational process requires manual monitoring and intervention to enable the PON protection switch. However, the process could be scripted by an external monitoring system.



Chapter 7

xDSL Subscriber Access

This chapter provides engineering and planning information regarding E7-2/E3-48C xDSL subscriber access.

Topics Covered

This chapter covers the following topics:

- Physical Access to Copper Loops
- Multiple Modes of DSL Deployment and Multi-pair Bonding
- xDSL Bonding Groups
- Integrated Voice Services on VDSL2 Combo Line Card
- G.INP Support
- High-Speed Internet Access
- DSL Port Noise Protection
- VDSL2 Vectoring
- Third-Party Modem Interoperability

Physical Access to Copper Loops

Copper loops with subscriber interfaces are terminated on 25-pair twisted pair RJ-21 connectors (Champ Telco-50 connectors) on to the rear of the E7-2 chassis. To connect subscriber interfaces to a protector panel or main distribution frame Calix recommends using CAT5 cables to achieve optimal Signal to Noise (SNR) margins and performance. It is also recommended to use 110° exit, right hand oriented Avaya 525e RJ-21 connectors on the CAT-5 cables, to avoid interference of the cables and adjacent connectors. For more detailed information please refer to the *Calix E7-2 Installation Guide*.

On the rear of the E7-2, the RJ-21 connectors are used differently depending on whether the double wide VDSL2-48 Overlay or single wide VDSL2-48C Combo cards are being deployed.

- VDSL2-48 Overlay:** Connectors labeled P1 and P2 correspond to DSL service loops for subscribers 1-48, and connectors labeled P3 and P4 correspond to POTS service loops for subscribers 1-48. The following graphic shows the rear view of an E7-2 RJ-21 connector layout for a single VDSL2-48 Overlay card.



- VDSL2-48C Combo:** Connectors labeled P1 and P2 correspond to combo subscribers 1-48 for Slot 1, and connectors labeled P3 and P4 correspond to combo subscribers 1-48 for Slot 2. The following graphic shows the rear view of an E7-2 RJ-21 connector layout for two VDSL2-48C Combo cards.



Multiple Modes of DSL Deployment and Multi-Pair Bonding

All xDSL ports can be configured independently for any DSL mode of operation. The length of the copper pair loop determines the DSL mode(s) that can be supported for each port. The loop length and gauge also determines the train rate and derived service rate that the Service Provider will be able to provide to the subscriber. The E7-2/E3-48C supports 3, 4, and 5 band VDSL2 capabilities using 8a, 8b, 8c, 8d, 12a, 12b, and 17a VDSL2 profiles. You can specify a profile setting or set for “Auto” which trains the line to the highest available performance. During the handshake process, the xDSL port and the CPE agree to determine the DSL mode automatically. The E7-2/E3-48C always tries to train the port at the DSL mode that provides the highest line performance possible, in order of priority as shown:

1. VDSL2 (PTM is a given, as it is the only transfer mode supported in this DSL mode)
2. ADSL2+ PTM mode
3. ADSL2+ ATM mode
4. ADSL2 PTM
5. ADSL2 ATM
6. ADSL; G.dmt, G.lite, ANSI T1.413. With traditional ADSL modes, ATM is the only transport supported.

Single pair DSL services are provisionable on all ports independently, and the line cards also support bonding of up to two pairs with any service type (that is ADSL/2/2+ and VDSL2 modes).


xDSL Bonding Groups

DSL bonding allows multiple loops to be aggregated into what appears to the applications as a single facility, roughly equal to the sum of the bandwidth of the individual lines. Calix E-series follows the bonding standards listed:

- ATM-based multi-pair bonding ITU-T G.998.1
- Ethernet-based multi-pair bonding (PTM) ITU G.998.2 G.bond

When a port in a bonding group fails or is administratively disabled, services continue on the ports that still function, within the bandwidth available to the remaining port(s). When the affected port recovers, the aggregate capacity of the bonding group resumes without operator intervention. Also, when both ports in an ATM-bonded interface fail or are administratively disabled, the interface will begin to pass traffic upon the first line recovering, depending on if the configured CPU allows it.

The E7-2 VDSL2 r2 cards running software version R2.5 and higher and the E3-48/C r2 running software version R2.6 and higher support DMT bonding of multiple pairs to Actelis ML-748-R or Positron AK624RU CPE devices.



Copper Bonded Uplinks

The E3-48 r2 Fiber + Copper models running software version R2.6 and higher support VDSL Transceiver Unit Remote (VTU-R) copper bonded uplinks. See Configuring Copper Bonded Uplinks.

Configuration guidelines

- Once the bonding group has been configured, the subscriber service provisioning is on the group interface.
- The Bonded DMT solution does not interfere with VDSL2 services that can be offered from the remaining available ports on the VDSL2 r2 line card.
- The bonding groups configured on VDSL2 r2 line cards should be configured in either 2, 4, or 8-pair bonding groups.
- The bonding groups should use contiguous ports and the starting ports should be configured as indicated below.

- 8-pair bonding groups can only be one of the following 6 groups:
 - Group 1: (v1, v2, v3, v4, v5, v6, v7, v8)
 - Group 2: (v9, v10, v11, v12, v13, v14, v15, v16)
 - Group 3: (v17, v18, v19, v20, v21, v22, v23, v24)
 - Group 4: (v25, v26, v27, v28, v29, v30, v31, v32)
 - Group 5: (v33, v34, v35, v36, v37, v38, v39, v40)
 - Group 6: (v41, v42, v43, v44, v45, v46, v47, v48)
- 4-pair bonding groups can only be one of the following 12 groups:
 - Group 1: (v1, v2, v3, v4)
 - Group 2: (v5, v6, v7, v8)
 - Group 3: (v9, v10, v11, v12)
 - ...
 - Group 10: (v41, v42, v43, v44)
 - Group 12: (v45, v46, v47, v48)
- 2-pair bonding groups:
 - For bonding groups with ports using PTM encapsulation, bonding of any xDSL port pairs is supported on the same card.
 - For bonding groups with ports using ATM encapsulation, bonding is restricted to pairs of odd/even adjacent ports. Select ports in a contiguous fashion, starting with an odd-numbered port "n," and its pair being the contiguous even-numbered port "n+1." For example, [1,2], [3,4], [5,6].
- The ports that are intended to be members of the same DMT bonding group must first be configured with matching Basic/Advance/PSD individual settings across all of the pairs.
 - Port profile settings must follow Best Practices indicated by the following:
 - *Calix xDSL Best Practices*
 - *Calix xDSL Profile Guidelines*
 - For bonded VDSL2 services, Calix recommends the following settings for the xDSL port:
 - Min Impulse Noise Protection (INP) = 2
 - Target SNR = 10 dB
 - Interleave Max Latency = 5 ms
 - All ports must have the same encapsulation (ATM-TC or PTM-TC).

- PTM bonding on VDSL2 can operate using any profile (8a/8b/8c/8d/12a/12b/17a), but 17a is limited to upstream tone 1216 (i.e. no limit for annex-B 998 bandplans).
- For VDSL bonding, use PTM Override = PTM (not Auto).
- For services to work correctly, only connect the CPE to cable pairs that are members of the Vectoring Group.
- Bonding group capacities:
 - Up to 24 2-port bonding groups per VDSL2 (r1, r2) card.
 - Up to 12 4-port bonding groups per VDSL2 r2 card.
 - Up to 6 8-port PTM bonding groups per E7-2 VDSL2 r2 cards.
 - ATM bonding is limited to 2-pair bonding. Greater than 2-pair DMT bonding requires PTM.
- It is important to match the transport mode desired with the modem's ability to support such a mode. For example, not all ADSL2+ capable modems support PTM mode.
- The length of the copper pair loop determines the DSL mode that can be supported for each port. The loop's length and gauge also determines the train rate, and derived service rate that the Service provider can provide to the subscriber. During the handshake process, the E-Series and the CPE agree to determine the DSL mode automatically, attempting to train the port at the DSL mode that provides the highest line performance possible.
- ATM Mode is supported for all ADSL-type modes. PTM Mode is supported only for ADSL2+, ADSL2, and VDSL2 modes.
- There is a 12-15% data transfer efficiency gain when choosing PTM as a service transport, due to avoiding the ATM cell tax (ATM overhead).
- Enhanced Impulse Noise Protection (INP) mode 'PhyR' is not supported on lines placed in a bonding group. If PhyR is enabled by the user, it will be automatically disabled by the hardware.
- "Bonding Scheme" is the terminology used in the ITU standards for the method used to bond. One can bond cells as per G.998.1 (scheme is atm), or one can bond packets as per G.998.2 (scheme is ptm).
- PTM Override set to Auto is not allowed as this results in an "Indeterminate Bonding Scheme."
- The Bonding Scheme is detected upon the provisioning of the ports, not by the negotiated encapsulation on the line.

The following table shows the settings that cause the bonding scheme to be indeterminate.

Service Type	PTM Override	Bonding Scheme	Result
VDSL2/ADSL2+ or VDSL2MM or Auto	ATM	Indeterminate	If the line trains in VDSL2 mode, it uses PTM-TC. If the line trains in a ADSL2/2+ mode, it uses ATM-TC. Therefore, the scheme is indeterminate at provisioning time.

Service Type	PTM Override	Bonding Scheme	Result
VDSL2/ADSL2+ or VDSL2MM or Auto	PTM	PTM	If the line trains in VDSL2 mode, it uses PTM-TC. If the line trains in ADSL2/2+ mode, it is forced to use PTM-TC.
VDSL2/ADSL2+ or VDSL2MM or Auto	Auto	Indeterminate	If the line trains in VDSL2 mode, it uses PTM-TC. If the line trains in an ADSL2/2+ mode, it could use ATM-TC or PTM-TC.
VDSL2	X	PTM	VDSL2 always uses PTM-TC.
Anything else	ATM	ATM	The line is provisioned as some kind of ADSL2/2+, and the line is forced to use ATM-TC.
Anything else	PTM	PTM	The line is provisioned as some kind of ADSL2/2+, and the line is forced to use PTM-TC.
Anything else	Auto	Indeterminate	The line is provisioned as some kind of ADSL2/2+, and the line could use ATM-TC or PTM-TC.

Any-to-Any Bonding Groups

Note: The bonding groups configured on VDSL2 r2 line cards should be configured in either 2, 4, or 8 contiguous-port bonding groups, and the starting ports should be configured as indicated in the Configuration Guidelines above.

Under certain circumstances you may encounter internal resource limitations, requiring you to intervene in any-to-any bonding group configurations. The message "Exceeds internal resources limit" indicates such a situation.

- When attempting to add ports to an any-to-any bonding group, you may encounter one of the following:
 - "No DSP resources available, please delete one of the existing >2-pair DMT bonding groups"**
The system has reached its limit of possible any-to-any bonding groups. You must delete one existing > 2 pair bonding group in order to accommodate a new > 2 pair bonding group.
 - "No DSP resources available, please re-configure Bonding Group 6 by using one of the following port ranges: v1~v8; v13~v20; v37~v44"**
This example message indicates you must re-configure the bonding group ports in a contiguous specified range.
 - "No DSP resources available, please delete group 7 and re-create (delete and re-configure) group 5 first before proceeding"**
This example message indicates you must delete and re-create the bonding group as instructed:
 - Delete the bonding group of which you are attempting to add a port.
 - Delete and recreate the existing bonding group in order to redistribute resources among DSPs.
 - Recreate the original bonding group with additional bonding ports.
- When attempting to remove a port from an any-to-any bonding group, you may encounter the following:

- **"failed: removing current DSL port from DSL bonded link will cause internal resource limits. First remove other ports"**

This message indicates that you must first remove other ports from the bonding group, and then remove the port that prompted the situation.

Parameters


You can provision the following parameters for a DSL bonding group:

Parameter	Description	Valid Options
Name*	Name of the xDSL bonding group.	text string
Admin State	Service state of the bonding group.	enabled (default), disabled
Subscriber ID	Subscriber ID information, such as phone number, or account number.	text string (blank) (default)
Description	Description of the DSL bonded interface, such as the subscriber address, name, or service.	text string (blank) (default)
IGMP Immediate Leave	Whether a multicast stream is dropped as soon as a Leave is received. The parameter is also set in the IGMP profile that is associated with a VLAN. Setting the parameter at the bonding group level overrides the setting in the IGMP profile.	enabled, disabled, use-vlan-setting
DS Min Rate	Downstream minimum rate in Kbps. Bonding group rates are controlled by the underlying port rates. The DS/US Min Rate parameters on groups are only used to generate low-rate alarms. If you configure these to some value between the single-port rate and bonded link rate, you will get a LOW RATE alarm on the group when one port goes down. These values are not related to physical train up rates or bandwidth profiles.	1-512000
US Min Rate	Upstream minimum rate in Kbps. Bonding group rates are controlled by the underlying port rates. The DS/US Min Rate parameters on groups are only used to generate low-rate alarms. If you configure these to some value between the single-port rate and bonded link rate, you will get a LOW RATE alarm on the group when one port goes down. These values are not related to physical train up rates or bandwidth profiles.	1-512000
Security Profile	Name of the Security Profile to use.	Any existing Security Profile.
DSCP-IP Precedence Map	Name of the DSCP or IP Precedence profile to use for mapping Layer 3 priority values to P-bits.	DSCPMap: access (default), IP PrecMap:access, any previously created profile
Force 802.1x	The 802.1x state.	none (default), authorized, unauthorized
VTU-Security	When the bonding group is connected to a copper bonded VTU-C Trunk interface, the VTU Security must be disabled (cleared checkbox) to disable some L2 security functionality as the subtended E3-48r2 units are trusted devices. When the bonding group is connected to a CPE (modem), the VTU Security must be enabled (selected) as the connected device is untrusted.	enabled (selected), disabled (cleared)
Member 1(-8)	The xDSL port to assign as a member of a bonding group. The xDSL ports that you want to add to a bonding group must first be configured with the valid combination of service type and PTM Override attributes and matching Basic/Advance/PSD individual settings.	any xDSL port

* Required

Bonded Pair Performance

Up to six 8-port VDSL2 bonded groups can be supported on a single 48 port VDSL2 r2 line card. The 48 subscriber ports of the VDSL2 r2 line cards can be configured in single, 2-pair bonded, or 8-pair VDSL2 bonded configurations or any combination in between as desired by the service provider.



The most important variable to guide deployments is the VDSL2 bonded throughput that can be provided to the remote site. The shorter the VDSL2 copper bonded loop the higher the trunk performance will be to the remote E3-48 r2 unit. The rate reach details for the E7 VDSL2 r2 line cards in DMT-8-pair bonded performance are provided below. The available throughput to the remote site will impact the oversubscription rate for the subscriber services provided from the remote site.

DMT 8-pair bonded performance results (Mbps)							
Loop Distance	0	1000	2000	3000	4000	5000	6000
DMT 8-pair bonded DS	1,120	1,008	649	430	304	221	142
DMT 8-pair bonded US	480	459	297	122	62	25	9

DMT bonding in Release 2.6 has been tested using the following CPE modem types:

- Actelis ML-748-R
- Positron AK624RU
- Actelis ML600

Copper Bonded Uplinks

The E3-48 r2 Fiber + Copper models have the capability of copper bonded VDSL2 trunk unlinks in addition to fiber uplinks, using the following port options:

- 8 VDSL2 pairs that are dedicated to the DMT bonded uplink capability
- 3 x 1GE / 2.5GE SFP and 2 x 1GE / 10GE SFP+ for fiber uplinks

Copper trunk connections to remote E3-48 r2 units offer the following:

- Re-purpose copper OSP for Trunk Applications
- Use DMT bonded ports for loops
- Allow migration from copper-fed to fiber-fed devices without a unit swap

DMT Bonded Uplink Trunks

The E3-48 r2 Combo and Overlay products support multiple pair bonding (DMT bonding) in both VDSL Transceiver Unit Central Office (VTU-C) and Remote Terminal (VTU-R) modes.

The "host" side (VTU-C) can be either an E7 line card or another E3-48 unit. There is no restriction of using both VTU-R and VTU-C functionality in the same E3-48 unit. The primary E3-48 can subtend a second E3-48 unit, via a GE point-to-point or a bonded interface.

Central Office (VTU-C) Mode

When the E7-2 or E3-48 r2 unit functions as the Central Office based unit, it can do the following:

- Connect to subscriber devices such as remote multiple pair bonded modems (multiple pair bonded modems from Actellis and Positron)
- Utilize DSL signaling to remotely tend E3-48 r2 units that are connected over copper bonded trunks


Remote Terminal (VTU-R) Mode

When the E3-48 r2 unit functions as the Remote Terminal base unit, it is used for northbound trunk transport connections. The remote VTU-R units can connect to Calix E7-2 VDSL2 r2 line cards as well as E3-48 VDSL2 r2.

The VTU-R DMT bonded trunk has ports specifically reserved so that 2, 4, 6, or 8 ports can be configured as one bonded group (Odd numbered DMT bonded trunks (3, 5, 7 port) configurations are not supported). These DMT bonded trunk ports do not deduct from the 48 subscriber-facing ports.

DMT Bonded Trunk Examples

The following diagram shows three example use cases for copper-fed E3-48 r2 units.



Example 1

An E7-2 VDSL2 r2 line card is using a copper bonded trunk to receive traffic from a remote E3-48 r2 unit.

- The E7-2 VDSL2 r2 line card (VTU-C) supports 48 subscriber ports and uses up to 8 ports in a bonding group, supporting the copper trunk to the remote device.
- The E3-48 r2 Fiber+Copper trunk unit has an additional set of dedicated copper ports for the trunk connections. The E3-48 r2 also supports 48 subscriber ports.

Example 2

An E3-48 r2 unit is subtending another E3-48 r2 unit.

- The E3-48 r2 unit on the 10GE ring (center) is using a 10GE fiber uplink and is subtending the remote E3-48 r2 unit (right).
- Similar to Example 1, the E3-48 r2 unit (VTU-C) supports 48 subscriber ports and uses up to 8 ports in a bonding group, supporting the copper trunk to the remote device.
- The subtended copper-fed E3-48 r2 unit (right) supports 48 subscriber ports.

Example 3

A cascaded deployment where a fiber-fed E3-48 (VTU-C) is subtending a copper-fed E3-48 r2 unit which in turn is subtending another copper-fed unit.

- The E3-48 r2 unit (middle) supports both VTU-C and VTU-R functionality.
- As shown in each example, the VTU-C functionality consumes some of the subscriber facing ports, while the VTU-R functionality uses the additional dedicated copper ports for the trunk.
- The VTU-R ports may not be re-configured for VTU-C functionality to support additional subtended devices.
- The services deployed over cascaded copper trunks are ideal for voice and data services.

Copper Bonded Uplinks and Remote Powering

The E3-48/C r2 introduces support for copper bonded uplinks with remote powering. The E3-48 r1 units currently support line powering, however the E3-48 r2 units expand on the capability by combining it with the copper bonded uplink.

- For information on providing remote power over the same pairs as DMT bonding, see *Calix Application Note: Designing an E3-48/E3-48C r2 Line-Powered Network*. Also see the video "Wiring an E3-48 r2 for Line Power + Copper Uplink" (<https://www.youtube.com/watch?v=S8Z1RfioNvE>).
- For instructions on enabling the line power pairs on the E3-48 r2, see Enabling Remote Powered Bonded Line Pairs.

Configuration Guidelines

SFP Constraints with E3-48 r2 Port Combinations	
Port Combinations	Support Description
• 2-10GE SFP+ • VTU-R Copper Uplink	Enabling 2 SFP+ ports and copper uplink in parallel is not supported.
• 2-10GE SFP+ • 2-1GE SFP Fiber	Requires the use of Direct Attach (DA) cables on the SFP ports (1Gbps).

- The SFP and SFP+ ports support the use of the copper direct attach (DA) cables available from Calix. Direct Attach cables are used to interconnect multiple E3-48 r2 units for node locations greater than 48 subscribers.
- IPTV services over copper bonded trunk uplinks (VTU-R) are not supported in this release (R2.6). It will be supported with some restrictions in a future release.
- A maximum of 6 service VLANs are supported over copper bonded links.
- DMT bonded links utilize DSL signaling to communicate to the remote device. Practically, the maximum distance for a DMT bonded trunk is the maximum distance of reADSL or 22,000ft. Realistically the ADSL2+ spectrum in concert with repeaters will extend the distance of the DMT bonded trunk. Without repeaters, ADSL2+ will be highly effective at loops greater than 8000 and less than 16000 ft. Below 8000 feet, the DMT bonded trunk can take advantage of VDSL2 signaling to increase the trunk capacity.

- VTU-C copper multi-pair bonding operates in VDSL2 and ADSL2+ modes. VTU-R bonded trunk uplink operates in VDSL2 mode only.
- When the bonding group is connected to a copper bonded VTU-C Trunk interface, the VTU Security must be disabled (cleared checkbox) to disable some L2 security functionality as the subtended E3-48r2 units are trusted devices.
- The maximum loop distance supported is 18kft using standard ADSL2 performance capabilities.
- The aggregate performance of the trunk link will be roughly a simple multiple of the ADSL2+ performance.
- A primary E3-48 can subtend a second E3-48 unit, via a GE point-to-point or a bonded interface.
- The E3-48 r2 supports a maximum of 24 VTU-C bonded connections (e.g. a maximum of 24 2-port bonded groups, 12 4-port bonded groups, 8 6-port bonded groups, 6 8-port bonded groups, or a mix of bonded group configurations).
- E3-48 r2 VTU-R line power is sourced from the remote E7-2 with E7-2 being the VTU-C.
- Remote power over the same pairs as DMT bonding is supported. See *Calix Application Note: Designing an E3-48/E3-48C r2 Line-Powered Network*.
- See the *Calix Customer Advisory Bulletin - E7-2, E3-48, E3-48C, E5-48/48C VDSL2 Modem Interoperability* for a list of the interop tested CPEs.

Parameters

You can provision the following parameters for a VTU-R copper bonded uplink:

Parameter	Description	Valid Options
Name*	Name of the VTU-R bonded trunk uplink.	text string
Admin State	Service state of the bonded trunk.	enabled (default), disabled
Description	Description of the VTU-R bonded trunk uplink, such as the subscriber address, name, or service.	text string (blank) (default)
DS Min Rate	Downstream minimum rate in Kbps. Bonded trunk uplink rates are controlled by the underlying port rates. The DS/US Min Rate parameters on bonded uplinks are only used to generate low-rate alarms. If you configure these to some value between the single-port rate and bonded link rate, you will get a LOW RATE alarm on the uplink when one port goes down. These values are not related to physical train up rates or bandwidth profiles.	1-512000
US Min Rate	Upstream minimum rate in Kbps. Bonded trunk uplink rates are controlled by the underlying port rates. The DS/US Min Rate parameters on bonded uplinks are only used to generate low-rate alarms. If you configure these to some value between the single-port rate and bonded link rate, you will get a LOW RATE alarm on the group when one port goes down. These values are not related to physical train up rates or bandwidth profiles.	1-512000

* Required

Before you start

Before you start configuring the copper bonded uplinks, you must have the connections in place for the ports of the VTU-C xDSL bonded group to the E3-48 r2 dedicated VTU-R Bond Trunk ports, including the remote power over the same pairs as DMT bonding. See the following references:

- *Calix Application Note: Designing an E3-48/E3-48C r2 Line-Powered Network*
- Enabling Remote Powered Bonded Line Pairs

Procedure Overview

The following sequence shows an overview of the copper bonded uplink configuration:


1. On the E7-2 or E3-48 r2 that will act as the VTU-C, do the following:
 - a. Create a VTU-C bonded group with a maximum of 8 ports and the VTU Security disabled (cleared checkbox).
 - b. Add services to the bonded group interface on the VTU-C.
2. On the E3-48 r2 that will act as the VTU-R, do the following:
 - a. Configure the VTU-R ports that will participate in the bonded uplink. The number of ports must match between the VTU-C bonded group and the VTU-R bonded uplink. That is, they must both be either 2, 4, 6, or 8 ports.
 - b. Add the service VLAN(s) that are configured on the VTU-C bonded group.

Remote Powered Bonded Line Pairs

The E3-48 r2 Fiber + Copper models introduce support for copper bonded uplinks with remote powering.

For information on providing remote power over the same pairs as DMT bonding, see *Calix Application Note: Designing an E3-48/E3-48C r2 Line-Powered Network*.

The remotely deployed E3-48 r2 unit accepts the Remote Powered + Copper bonded trunk pairs to power the unit and to provide the copper uplink trunk over 8 VDSL2 pairs that are dedicated to the uplink capability. For each 8-DMT pair group, 3 power pairs are combined with 6 of the DMT bonded pairs. The DMT signal/power combiner puts power from one power pair onto two DMT bonded pairs. The remaining two pairs of the 8-DMT pair group are not powered and only carry DMT signaling.



The powered pairs are broken into three channels (1-3) with two pairs in each channel.

- Channel 1 has Pairs 1-2
- Channel 2 has Pairs 3-4
- Channel 3 has Pairs 5-6


The E3-48C r2 (combo) units require a minimum of 3 channels (at least 6 powered copper pairs) for all applications. The E3-48 r2 (overlay) units can operate with 2 channels, where the third channel is provided for redundancy purposes, if one channel fails.

E-Series Release 2.6 introduces a power save feature for remote powered E3-48C (R1) and E3-48 (R2) Combo and Overlay units so when the power source changes from AC to battery backup during an AC power failure, the system switches-off xDSL services on selected nodes, leaving the remaining battery power available for higher priority voice connections to emergency services (for example, 911) or premium xDSL services (for example, schools, government, military installations, hospitals). See *Calix E-Series (E7 OS) Maintenance and Troubleshooting Guide*.

There are two supported applications as described in the figure below.


Example 1 shows remote powered pairs coming from the Central Office and Bonded Copper Uplinks coming from an E3-48 r2. The E3-48 r2 unit supports a multiple-pair bonded connections. The E3-48 r2 unit with VTU-C functionality (center) requires an external combiner / joiner that combines the VDSL2 bonded signal pairs with the copper remote powered pairs from the central office. The remotely deployed E3-48 r2 unit (right) accepts the Remote Powered + Copper bonded trunk pairs to power the unit and to provide the copper uplink trunk. The E3-48 r2 Fiber + Copper Uplink models have 8 VDSL2 pairs that are dedicated to the uplink capability. These pairs have been combined with the 6 powered pairs. The powered pairs are broken into three channels (1, 2, 3) with two pairs in each channel. Thus Channel 1 has Pairs 1-2; Channel 2 has Pairs 3-4; and Channel 3 has Pairs 5-6. The E3-48 r2 units require a minimum of 2 channels. The third channel is provided for redundancy purposes if one channel fails based on the distance and power requirements. Depending on the loop length and application, 3 channels may be required.

Example 2, in the diagram shows the remote powered pairs from the cabinet. This FTTN remote location (center) will require the -48vDC to +/- 190vDC upstream remote power converter, the Remote Power/Copper bonded uplink trunk combiner / joiner, and the battery back-up system to power both the FTTN location and the subtended device. The combiner / joiner is provided in rack mount and sealed form factors. The rack mount solution is designed to support up to 3 E3-48 r2 units with 8 pairs of VDSL2 bonding and 6 pairs of remote power (3 channels, 2 pairs per channel). The rack mount solution comes with the cabling and should be ordered with the ODC-100 power hub solution for cabinet installations. This ODC-100 power hub solution may also be ordered separately for cabinet retrofit installations. The sealed combiner / joiner solution is designed to support one E3-48 r2 unit supporting up to 8 pairs of VDSL2 bonding and 6 pairs of remote power (3 channels, 2 pairs per channel). The sealed version is installed in a NEMA-3 unit for ease of deployment in the outside plant and contains IDC based protection modules.




CAT3 Power Shutdown for Remote Powering

E-Series Release 2.6 introduces CAT3 Power Shutdown for remote powering. This feature extends the battery life of the power hub cabinet by disabling DSL on selected DSL ports. The extended battery life may be used to support POTS or premium DSL services (e.g. schools, government, military installations, hospitals, etc). This feature is supported on the E3-48C (R1), and E3-48 (R2) Combo and Overlay units. During an AC power outage at the power hub, the remote powered units are powered by the battery back up system of the power hub. The E7 at the power hub site sends a message to the remotely powered E3-48 units to power down DSL services on the affected lines. Upon AC power restoration, the E7 sends a message to the remotely powered E3-48 units to reestablish DSL services on affected lines.



Multiple VC Support for ADSL2+ Fallback

The E7-2/E3-48C products support up to 6 services per port that may be configured as Voice, Data, Video, TR-69 Management, Transparent LAN Services, or other VLAN assignments. Traditional POTS voice services in the case of the E7-2 VDSL2-48C Combo line card, E5-48C, and E3-48C products are handled utilizing the POTS DSP-to-VoIP processing hardware of the solution and are not segmented as a separate service. Digital Voice services that are created by an IAD or SIP VoIP device at the subscriber termination are routed through the E7-2 VDSL2-48D Data-Only Line Card as a data service. Support for multiple services is used when migrating from an ATM based DSLAM to the E-Series solutions. This feature allows service providers to migrate from the older DSL platforms to the E-Series solutions while continuing to support multiple services across the deployed ATM based CPE. In this context, services are assigned to a specific VLAN ID and may be applied to each port as shown in the diagram below. The E-Series support one VC assignment per service.



For installations using multiple PVC assignments, individual PVCs may be mapped to specific VLANs that are applied to the port. For example, installations that have multiple PVC assignments (for example, 0/33, 8/33, 0/35) may be mapped to individual data services (for example, "Data 1", "Data 2", "Data 3") that would then be applied to the port. TR-69 management in the example above is shown as a separate service but may also be integrated into a data service stream. VLAN ID assignments ("VLAN 3", "VLAN 4", etc.) are shown in quotations marks and may be assigned to any supported VLAN ID number.

Integrated Voice Services on the VDSL2 Units

The E7-2 VDSL2-48C line card and E3-48C/E3-48/E5-48/E5-48C ESAN support switched voice services from the combo POTS / DSL ports. An integrated Voice Frequency (VF) splitter terminates the analog POTS signal, and hands it to a voice Digital Signal Processor (DSP) for native VoIP conversion. The DSP performs the call handling and IP packetization required for VoIP implementation. The integration of VoIP in the device uses Session Initiation Protocol (SIP) as specified in RFC-3261. Supplemental protocol enhancements support H.248 voice services and TDM voice services (VoIP packet streams with legacy Class 5 switches in conjunction with C7 and B6 Gateways).

In addition to the DSLAM function, Quality of Service (QoS) for voice traffic is provided throughout the Ethernet transport network and uplink interfaces. VoIP traffic is tagged with an 802.1Q VLAN tag and prioritized using 802.1p priority bits.

VDSL2 based voice services are delivered in conjunction with either or both of the following:

- Calix C7 TDM switch interfaces (GR-303, TR-08 Mode I and Mode II) using a Voice over IP Resource card (VIPR) or Ethernet Gateway (EGW) card, as the VOIP to TDM gateway. The *Calix C7 VoIP Services Guide* provides a complete service description and provisioning instructions.
- Third party softswitches from MetaSwitch and Genband, for SIP signaling native VoIP services or H.248 VoIP services. Calix conducts softswitch interoperability testing through its engineering organization and System Quality Assurance (SQA) process. Calix continuously updates interoperability tests as new switches and protocols are added to the Calix VoIP portfolio.

The VDSL2 SIP solution supports the standard voice service offerings of a Class V switch, so the VoIP implementation is transparent to the subscriber. These features include Caller-ID, Call Waiting, Call Transfer, Three-Way Calling, Distinctive Ring, Ring Splash, and E911 support. The suite of supported CLASS features depends on the type of SIP server (softswitch or C7 TDM gateway) because feature support varies by vendor. Consult the above referenced documents for a complete set of ONT voice service attributes.

TDM interfaces to a Class5 switch can be GR-303 or TR-08 (Modes 1 and 2) in combination with the C7's TDM Gateway function.

G.INP Support

The E-Series supports ITU-T G.INP, which is an industry standard for DSL line physical layer retransmission to improve impulse noise protection. G.INP provides physical layer retransmission essential for the reliable delivery of high quality IPTV services in a deployment environment adversely affected by impulse noise. Both VDSL2 and ADSL lines are able to take advantage of the G.INP when supported by a G.INP capable CPE. This feature may be configured on lines where there is a G.INP CPE. There are no restrictions in supporting a mix of ITU-T G.INP modems and non-standard physical retransmission standard modems.

High-Speed Internet Access

The E-Series and 3rd party modems provide data interconnection between Internet service providers and subscribers. The method of transport over the copper loop is Packet (Ethernet) or Asynchronous Transfer Modes (ATM). PTM can be supported by VDSL2, ADSL2+ and ADSL2 modes, while ATM can only be supported in ADSL2+ / ADSL2 / ADSL / G.dmt modes. Each mode must also be supported by the specific 3rd party modem used. Service profiles provide tiered service offerings. Different models of 3rd party modems provide different types and numbers of Fast Ethernet (100 Mbps) and/or Gigabit Ethernet (1000 Mbps) ports, USB ports, WIFI interfaces, etc.


Annex B Co-Existence

E7 VDSL2 line cards support Annex B creating signal interoperability between ISDN and DSL services by applying a DSL signaling profile with a PSD limit mask to avoid ISDN frequencies, creating interoperability in the binder for DSL and ISDN subscribers.

Annex B is predominantly used in European markets where ISDN services have been deployed for subscriber internet applications.

- ISDN networks support voice and data services over a symmetric 128kbps circuit.
- DSL services supporting voice, data, and video applications support considerably higher bandwidth but use the same lower signaling spectrum.

In the case where the E7-2 VDSL2 Overlay or Data-Only line cards are deployed, an external splitter is required to route the ISDN traffic back to the ISDN host equipment. This deployment model is identical to providing data and video applications where there is an existing Class 5 switch providing POTS services. In the ISDN case, the ISDN switch sends the traffic toward the subscriber, the external splitter joins the ISDN and DSL frequencies and provides both in the binder toward the respective subscribers.



DSL Port Noise Protection

You can specify the following xDSL port parameters for noise suppression directly on a port or in a DSL Port Template that can be applied to multiple ports.

Seamless Rate Adaptation

Seamless rate adaption (SRA) targets optimal data rate performance without operator intervention, providing a safeguard from changes in the cross-talk levels in the binder caused by a variety of sources. These sources include POTS and other service signals, AM radio, disturbers, temperature changes, and water in the binder group. SRA is used to maintain good SNR (Signal to Noise Ratio) levels by automatically tweaking the line rate to avoid service interruption or high bit error rates.

SRA is part of the ITU-T G.992.2 and is supported on the BRCM 65300 chipset used in the E7-2/E3-48C. SRA can be configured in the Embedded Web Interface (EWI), under “Basic Settings > xDSL Profile Setup > xDSL Profile.”

PhyR and Impulse Noise Protection

Physical layer-retransmission or PhyR™ is a Broadcom technology used to set the best possible impulse noise protection (INP) value on the subscriber line in order to achieve low bit error rates. Other existing algorithms used to achieve improvements in INP are Reed-Solomon, trellis, interleaving schemes, etc. PhyR is a data retransmission scheme implemented above the DSL physical layer via Broadcom VDSL2 / ADSL2+ CO and CPE chipset firmware.

There is no status showing a broadband loop in PhyR mode, but if it's active the INP value will be considerably larger (typically 10x) than the interleave-based INP.

VDSL2 G.INP Support

Release 2.2 introduces support for ITU-T G.INP which is an industry standard for DSL line physical layer retransmission to improve impulse noise protection. G.INP provides physical layer re-transmission essential for the reliable delivery of high quality IPTV services in a deployment environment adversely affected by impulse noise. Both VDSL2 and ADSL lines are able to take advantage of the G.INP when supported by a G.INP capable CPE. This feature may be configured on lines where there is a G.INP CPE. There are no restrictions in supporting a mix of G.INP modems and non-standard physical retransmission standard modems.

Unit-Level Vectoring (E3-48C, E5-48C, E5-48)

Release 2.3 introduces unit-level vectoring support for the E3-48C, E5-48C, E5-48 products (not E7-2). Release 2.4 includes E7-2 in supporting unit-level vectoring capabilities that eliminate cross talk between VDSL2 lines, maintaining the available bandwidth. Vectoring also ensures a uniform level of performance from pair to pair in a vectored binder group. The E-Series vectoring complies with the ITU G993.5 standard. Vectoring-capable CPE must be deployed in conjunction with the E-Series units to achieve cross talk elimination using the vectoring function. In addition, the vectored VDSL2 lines cannot share binder groups with non-vectored VDSL2 lines or else the benefits of cross talk noise recovery are lost.

In high cross talk environments, a severe negative impact can occur on the achievable bandwidth of VDSL2. In a binder group full of VDSL2 circuits, as much as 50% of the achievable bandwidth is lost due to cross talk. Compounding the issue is the fact that the impact of cross-talk is not linear – it is not predictable.

This is due to the pseudo-random nature of VDSL2 distribution in a binder group. The following example illustrates this point.

- Consider a binder group with just two active VDSL2 pairs. At first glance it seems reasonable to assume that since only two pairs are active, the level of cross talk should be low, and thus minimal degradation would be observed. But even with just two pairs, the impact can be significant depending on the location of the two pairs relative to each other.
- If the two pairs are spaced far apart, the impact of cross talk would indeed be minimal. But if the two pairs are adjacent to each other, then severe crosstalk would occur, and significant bandwidth degradation would be observed.

The fact that cross-talk, and corresponding bandwidth degradation is unpredictable makes it a vexing customer service problem for which the only solution is to move the customer from pair to pair until by chance a sufficient spatial separation is achieved for cross talk minimization. Obviously this is a labor intensive solution with no certainty of success.

System level vectoring up to 384 lines on VDSL2r2

In release 2.6 the E7-2 VDSL2 r2 line cards add support for system level vectoring for subscriber densities up to 384 ports. Vectoring is a noise cancelation capability that restores the available bandwidth that is typically lost to cross talk in VDSL2 binder groups. Vectoring typically can improve performance out to distances of ~3,000 feet (915 meters) depending on a number of variables in the outside copper plant. There are two types of vectoring which are defined as unit level and system level vectoring. Unit level vectoring eliminates the cross talk within the ports associated with a specific unit or line card. Unit level vectoring has been supported on E-Series products beginning in 2013 and is currently supported on the E3-48, E3-48C, E5-48, E5-48C, and all E7-2 VDSL2 line cards. System level vectoring eliminates the cross talk across multiple units or line cards. System level vectoring typically requires a dedicated external vectoring processor that manages the DSL signaling profiles across an entire system. The E7 VDSL2 r2 line cards support system level vectoring in this traditional dedicated processor model and also introduces the unique capability of supporting 96-port system level vectoring without requiring a dedicated external vector control processor. This enables a more operational and capital efficient deployment option for smaller density locations up to 96-ports. For higher density locations a vectoring control processor can be added to support 192 ports (4 VDSL2 r2 line cards) or up to 384 vectored ports (up to 8 VDSL2 r2 line cards).

When Vectoring is Recommended

Vectoring is recommended for environments with high levels of cross talk. Conversely, vectoring is of little to no benefit in environments with low levels of cross talk. Just as important, vectoring is of little benefit if the target speeds are low – less than 50Mbps. And because vectoring is a double-ended technology, an installed base of non-vectoring capable CPE means vectoring cannot be deployed unless the CPE is replaced.

For example, a CSP planning to deploy FTTN and new CPE to offer 100Mbps Internet for \$40/month over 3,000 foot loops with bonding is a prime candidate for vectoring. At \$40 per month take rate will be high, and with bonding there will be a high density of VDSL2 in the binder groups. Indeed, without vectoring the CSP will be unable to meet his target speed of 100Mbps on loops between 2,000 and 3,000 ft.

On the other hand, a CSP that has been incrementally adding VDSL2 into his existing access network for the last few years to sell incrementally enhanced services (up to 50Mbps) will gain little benefit from vectoring due to long loop lengths, low VDSL2 take rates, non-vectoring capable CPE, and a varied mix of ADSL2+ and VDSL2.

The table below indicates when vectoring is appropriate:

Access Variable	Vectoring Benefit
High VDSL2 density (high cross talk)	High
Low VDSL2 Density (low cross talk)	Low
50Mbps or less	Low
50-75Mbps	Moderate
Up to 100Mbps	High

Access Variable	Vectoring Benefit
Long loops (3,000+ ft, 900+meters)	Low
Short loops (<3000 ft, <900 meters)	High
Mixed ADSL2+ and VDSL2	Varies according to loop length and service mix
VDSL2 bonding	High
Legacy VDSL2 CPE (non-vectorized)	Low

Impact of Non-Vectorized CPE

Because vectoring requires synchronization of the vectored lines, the DSLAM and connected CPE must both support vectoring in order for cross talk to be eliminated. However, in many existing environments, existing CPE may not support vectoring. While replacing the legacy CPE with vectoring CPE is the preferred solution, additional less invasive options exist for implementing vectoring in the presence of non-vectoring capable CPE.

Depending on the services being delivered, CSPs may have the option of force-training non-vectoring VDSL2 CPE in ADSL2+ mode. Because ADSL2+ operates a lower frequencies than VDSL2, ADSL2+ is a minor disturber relative to non-vectorized VDSL2. This approach is suitable for services of 25Mbps or less (the upper limit for ADSL2+).

Vectoring is disabled by default in the E-Series products. If vectoring is required, it can be activated via any of the user interface options.

Planning Considerations for Enabling Vectoring

When planning a unit level vectoring implementation a number of considerations and requirements are critical. The first concept which is different from traditional copper DSL deployments is that the binder group (the copper cable from the DSLAM to the subscribers) needs to be considered as shared spectrum much like a fiber in a GPON deployment. Transmission strength and frequencies on both ends of the line, the DSLAM and the modem, are managed to reduce the overall noise floor. For optimal performance, Calix recommends that all modems in a vector group be from the same modem manufacturer, use the same model type, and run the same firmware for a successful vectoring implementation. This ensures that all members of the vectoring group behave in the same manner when the vectoring group is trained. Vectoring technology is still in the early stages of commercial deployments and vectoring group train up time, modem behaviors, multiple modem interoperability within a binder group, etc., are expected to improve over time.

Train-Up Times

Vectoring group train up time varies from 1.5 – 4 minutes depending on the noise conditions in the binder and the number of modems in the vectoring group. When vectoring is enabled, the E-Series control the training of each of the modems that join the vectoring group. If a modem misses the initial vectoring group train cycle, the modem would then join the vectoring group in a subsequent cycle. Missing a vectoring group training cycle will add an average of 74-100 seconds to the train-up time for the modem. Ensuring that all modems are the same make, model, and firmware reduces the likelihood that a modem will miss the vectoring group train cycle.

Impact to Adjacent Lines

The initialization of the vectoring group is a cross talk rich process that may cause a retrain of a non-vector friendly modem. This is possible during the initial deployment of the vectoring group or after a system reboot. While this is possible, it is seen as a limited occurrence due to the distance variation between the modem types (VDSL2 vectoring groups are deployed less than 3000 feet while existing ADSL2+ modems are deployed on average greater than 5000 feet) and the frequency spectrum differences (VDSL2 uses 8MHz, 12MHz, and 17MHz while ADSL2+ uses 2.2MHz). Calix internal testing using outside plant loop distances of 1000, 2000, and 3000 feet for VDSL2 vectored modems and 6000 feet for ADSL2+ modems confirmed the low likelihood of a non-vector modem friendly retrain. However, a non-vector friendly VDSL2 modem deployed at the same loop distance as the vectoring group modems may experience a retrain as the vectoring group is initialized. In this case, Calix recommends the use of vector-friendly VDSL2 modems.

CPE Management and Considerations

As vectoring improvements are released it is also expected that existing VDSL2 modems that are non-vector friendly today may have vector-friendly firmware released in the future. While the modem may not have vectoring capability, vectoring-friendly may alleviate the amount cross talk generated in the binder as the modem is in service and or may reduce the likelihood of a retrain during the vectoring group initialization. With these concepts in mind Calix highly recommends TR-69 management to manage modem firmware as future modem firmware releases are expected to improve vectoring performance and interoperability. TR-69 would provide system wide modem firmware management which is a key success criterion for vectoring deployments.

Vectoring offers significant capability to increase bandwidth as well as to provide standardized service offerings. For vectoring deployments Calix recommends a minimum of 4 modems in a vectoring group before the vectoring benefits can be recognized. Fewer than 4 modems in the vectoring group will not yield the significant gains expected by enabling vectoring. In the ideal deployment, all vectored modems in the binder have the same make, model, and firmware and are deployed with roughly the same loop distance. This example represents the maximum gain possible from vectoring and Calix internal laboratory testing has shown up to 60% performance improvements for bandwidth. In mixed deployments supporting both VDSL2 vectored modems and non-vectored modems the presence of non-vector friendly VDSL2 modems represents uncontrolled disturbers in the binder group. The presence of disturbers in the binder reduces the overall vectoring performance improvement gain. The impact of the disturbers will vary based on the loop distance of both the disturbers and vectoring group members. With each of these considerations in mind vectoring offers a powerful new capability to provide new services over copper.

VDSL2 Vectoring

E-Series (E7 OS) units support vectoring on VDSL2 ports, eliminating crosstalk between VDSL2 lines thereby recovering bandwidth that would otherwise be “lost” due to crosstalk. Vectoring also ensures a uniform level of performance from pair to pair in a vectored binder group. The E-Series vectoring complies with the ITU G993.5 standard.

The scope of Recommendation ITU-T G.993.5 (Dynamic Spectrum Management Level 3) is specifically limited to the self-FEXT (far-end crosstalk) cancellation in the downstream and upstream directions. FEXT generated by a group of near-end transceivers and interfering with the far-end transceivers of that same group is cancelled. This cancellation takes place between VDSL2 transceivers, not necessarily of the same profile.

Note: To achieve crosstalk elimination with vectoring, vectoring capable CPE must be deployed in conjunction with the E-Series units. In addition, E-Series vectored lines must be in the same binder group. They cannot share binder groups with non-vectored VDSL2 lines or the benefits of cross talk noise recovery are lost. It is recommended that all configured modems are of the same make, model, and software or firmware. Using non-vectoring CPEs in a Vectoring Group can potentially be service affecting.

The following vectoring modes are supported for E7 release R2.6 and higher:

Vectoring Mode	#Ports	Shelves and Cards Supported
Board-Level (BLV)	Up to 48	- E7-2 MC Controller / VDSL2 r2 - E7-2 MC Expansion / VDSL2 r1 or VDSL2 r2 - E7-2 stand-alone / VDSL2 r1 or VDSL2 r2 - E3-48C, E5-48, E5-48C, E3-48 nodes
Dual-Level (DLV)	Up to 96	- E7-2 MC Controller / VDSL2 r2 - E7-2 MC Expansion / VDSL2 r2 - E7-2 stand-alone / VDSL2 r2
System-Level (SLV)	Up to 192	- E7-2 MC Controller / VDSL2 r2 - E7-2 MC Expansion / VDSL2 r2 - E7-2 stand-alone / VDSL2 r2 - E3-48 r2, E3-48C r2 nodes Required VDSL2 Vector Control Processor: - E7-2 stand-alone / E7 VCP-192 - E3-48 r2, E3-48C r2 nodes / E3 VCP-192
	Up to 384	- E7-2 MC Controller / VDSL2 r2 - E7-2 MC Expansion / VDSL2 r2 - E7-2 stand-alone / VDSL2 r2 Required VDSL2 Vector Control Processor: - E7-2 stand-alone / E7 VCP-384

See the *Calix Customer Advisory Bulletin - E7-2, E3-48, E3-48C, E5-48/48C VDSL2 Modem Interoperability* for a list of the interop tested vectoring CPEs.

High-density deployment models

High density VDSL2 nodes of 192 or 384 ports are typically supported using an E7-2 modular chassis which manages the node as a single IP address. The VCP units are subtended GE connected devices with a separate IP address. The management connection of the VCP enables software updates of the VCP units. VCP units are supported in any E7 topology as they are subtended as a management IP address associated with the node. Only one VCP per node is recommended. Multiple Vector Control Processors in a single installation are not recommended as the vector optimization of one vector processing group is not coordinated with the other vector processing group which re-introduces cross talk. The E7-2 Vector Control Processor should be connected to a GE port of the line card in the MCC shelf for management connections.

Bonded VDSL2 lines in a vectoring group

E-Series VDSL2 products support bonded VDSL2 lines in a vectoring group as they interoperate with all chipsets for VDSL2 bonded vectoring capable CPE (Broadcom, Ikanos, G.993.5 compliant Lantiq). Bonded vectored performance restores the bandwidth impacted by cross talk in a binder group and extends the reach of 100Mbps bonded VDSL2 modems. Bonded vectored ports can deliver 100Mbps downstream performance on service loops up to 4,000'. For additional information on Bonded Vectoring CPE interoperability please see the VDSL2 CPE Interoperability Customer Advisory Bulletin.

Configuration guidelines

- Vectoring can only be applied to xDSL ports that have a service type of VDSL2 or VDSL2 variant.
- You can add individual ports or a range of ports to a vectoring group.
- Calix DSL Best Practices recommend that non-vectored VDSL2 modems be configured in ADSL2+ mode to avoid active disturbers in the vectoring group.

Board-Level Vectoring (BLV):

- The E-Series supports board Board-Level Vectoring (BLV), where all ports on a VDSL2 card or VDSL2 unit work together as a 48-port vector group. Using vectoring on some lines but not all causes un-vectored lines to interfere with vectored lines. It is recommended to keep all ports groomed in the same binder.
- E7-2 VDSL2 r1 cards in the MCE shelves only support BLV vectoring.
- For BLV, there is only one vectoring group per VDSL2 card or VDSL2 unit.

Dual Line-card Vectoring (DLV):

- The E7-2 supports Dual Line-card Vectoring (DLV) that eliminates crosstalk across two VDSL2 r2 line cards, where all the ports on two cards within the same shelf work together as a 96-port vector group.

- The DLV feature is only supported on E7-2 VDSL2-48D r2 and VDSL2-48C r2 cards in the same shelf (standalone or modular chassis), without requiring a dedicated external vector control processor.
- For DLV, there must be two VDSL2 r2 cards in the same shelf, where:
 - One card must be configured as DLV-MASTER
 - One card must be configured as DLV-SLAVE.
- There is only one shared vectoring group located on the DLV-MASTER card.
- A modular chassis system supports up to 2 DLV groups, along with multiple BLV groups.

System-Level Vectoring (SLV):


- The E7-Series supports System-Level Vectoring (SLV) that eliminates crosstalk across multiple VDSL2 r2 line cards or units, where all the xDSL ports on all cards work together as a vector group.
- The SLV feature is only supported with a dedicated external Vector Control Processor (VCP) and interconnecting interface Vector Control cable.
 - A VCP-192 is required for vector groups with up to 192 ports across 4 VDSL2 r2 cards or units.
 - A VCP-384 is required for vector groups with up to 384 ports across 8 VDSL2 r2 cards.
- For the connections on each E7 Vectoring Control Y-cable:
 - The VDSL2 r2 overlay cards configured for SLV are located across multiple chassis (modular chassis or standalone shelves).
 - The VDSL2 r2 single-slot cards configured for SLV must be located in the same shelf.
- Once a VDSL2 r2 card or unit is set to SLV mode, the action of either connecting or disconnecting a VCC cable to the card or unit may impact all of the xDSL ports on that card or unit.
- For consistency in alarm reporting, the following E7 VCP (SLV Host)/ VDSL2 R2 (SLV Target) configuration with the Vector Control Y-cable is the Best Practice recommendation:


VCP (SLV Host)	E7-2 VDSL2 r2 (SLV Target) Single-slot Cards
VC 1-2	Shelf 1: • VCC Cable 'A' connects to Card 1 • VCC Cable 'B' connects to Card 2
VC 3-4	Shelf 2: • VCC Cable 'A' connects to Card 1 • VCC Cable 'B' connects to Card 2
(VCP-384 only) VC 5-6	Shelf 3: • VCC Cable 'A' connects to Card 1 • VCC Cable 'B' connects to Card 2

VCP (SLV Host)	E7-2 VDSL2 r2 (SLV Target) Single-slot Cards
(VCP-384 only) VC 7-8	Shelf 4: - VCC Cable 'A' connects to Card 1 - VCC Cable 'B' connects to Card 2


- Each E7 VCP unit has two SFP sockets at 1GE rate that are redundant ports (MGMT' GE) for out-of-band management.
- Each E3 VCP unit has two QSFP VCC ports that operate as redundant interfaces for in-band management.
- The 192.168.254.x/24 subnet is reserved for embedded channel communication for SLV and DLV applications. Therefore it cannot be used for management of the VCP unit.
- In-band management of the E7 VCP is not supported, however you can configure remote management using a tag-action that adds the system in-band management VLAN ID to untagged traffic at the MGMT' GE connecting GE port interface.
- E7 release 2.6.40 (with CMS R14.1) introduces the ability to manage an E7-VCP and an adjacent E7 Modular Chassis (MC) under a single management IP address, using different TCP/UDP ports. See *Calix E-Series (E7 OS) Turn-Up and Transport Guide*.
- All cards or units in an SLV configuration must be running the same software version.
- An E7 modular chassis system supports only one SLV group, along with multiple BLV groups.


E7 VCP-192 supported SLV configurations for E7 standalone and modular chassis are as follows:






E7 VCP-384 supported SLV configurations for E7 standalone and modular chassis are as follows:







Bonded VDSL2 Ports:

- Both single and bonded VDSL2 ports are supported in a vectoring group.
- For a vector group to have bonded VDSL2 port members, you must first enable vectoring on the ports, and then add the ports to the bonding group.

Vectoring Groups and State:

- A vectoring group exists by default per VDSL2 line card or VDSL2 unit.
- Individual lines can be added into the vector group.
- Vectoring on a xDSL port is disabled by default.
- The join-vectoring-group attribute is automatically changed to “disabled” when the port’s Vectoring Group is set to “none”. Another way to force a modem to train in and out of vectoring (new in R2.5) is to set the DSL port vectoring state to disabled, causing the DSL vectoring group to automatically change to “None”.
- The Upstream Vectoring and Downstream Vectoring parameters must be enabled for vectoring to be enabled on a port that has joined a vectoring group.
- A change in Vectoring Mode (DLV, BLV, SLV) results in a card or unit reset.

Note: When xDSL lines that are not part of a DSL vector group retrain, lines that are part of the DSL vector group may also retrain.

When is vectoring recommended?

Vectoring is recommended for environments with high levels of cross talk. Conversely, vectoring is of little to no benefit in environments with low levels of cross talk. Just as important, vectoring is of little benefit if the target speeds are low – less than 50Mbps. And because vectoring is a double-ended technology, an installed base of non-vectoring capable CPE means vectoring cannot be deployed unless the CPE is replaced.

Access Variable	Vectoring Benefit
High VDSL2 density (high cross talk)	High
Low VDSL2 density (low cross talk)	Low
50Mbps or less	Low
50-75Mbps	Moderate
Up to 100Mbps	High
Long loops (3,000+ ft, 900+meters)	Low
Short loops (<3000 ft, <900 meters)	High
Mixed ADSL2+ and VDSL2	Varies according to loop length and service mix
VDSL bonding	High
Legacy VDSL2 CPE (non-vectoring)	Low


Planning Considerations when Enabling Vectoring

When planning a VDSL2 vectoring implementation, a number of considerations and requirements are critical. The binder group (the copper cable from the DSLAM to the subscribers) needs to be considered as shared spectrum. Transmission strength and frequencies on both ends of the line, the DSLAM, and the modem are managed to reduce the overall noise floor. For optimal performance, Calix recommends that all modems in a vector group be from the same modem manufacturer, use the same model type, and run the same firmware for a successful vectoring implementation. This ensures that all members of the vectoring group behave in the same manner when the vectoring group is trained.

Please refer to the *Calix E7 Platform Release Product Planning Guide* for more information about planning considerations when enabling vectoring.

Deployment Models

High density VDSL2 nodes of 192 or 384 ports are typically supported using a modular chassis which manages the node as a single IP address. The VCP line cards are subtended GE connected devices with a separate IP address. The management connection of the VCP enables software updates of the VCP line cards. VCP line cards are supported in any E7 topology as they are subtended as a management IP address associated with the node. As described elsewhere in this document, only one VCP per node is recommended. Multiple Vector Control Processors in a single installation are not recommended as the vector optimization of one vector processing group is not coordinated with the other vector processing group which re-introduces cross talk. The E7-2 Vector Control Processor should be connected to a GE port of the line card in the MCC shelf for management connections.



See the *Calix E-Series (E7 OS) xDSL Applications Guide* for procedures on adding ports to a vectoring group and other related tasks.

System Level Vectoring

In release 2.6 the E7-2 VDSL2 r2 line cards add support for system level vectoring for subscriber densities up to 384 ports. Vectoring is a noise cancelation capability that restores the available bandwidth that is typically lost to cross talk in VDSL2 binder groups. Vectoring typically can improve performance out to distances of ~3,000 feet (915 meters) depending on a number of variables in the outside copper plant. There are two types of vectoring which are defined as unit level and system level vectoring. Unit level vectoring eliminates the cross talk within the ports associated with a specific unit or line card. Unit level vectoring has been supported on E-Series products beginning in 2013 and is currently supported on the E3-48, E3-48C, E5-48, E5-48C, and all E7-2 VDSL2 line cards. System level vectoring eliminates the cross talk across multiple units or line cards. System level vectoring typically requires a dedicated external vectoring processor that manages the DSL signaling profiles across an entire system. The E7 VDSL2 r2 line cards support system level vectoring in this traditional dedicated processor model and also introduces the unique capability of supporting 96-port system level vectoring without requiring a dedicated external vector control processor. This enables a more operational and capital efficient deployment option for smaller density locations up to 96-ports. For higher density locations a vectoring control processor can be added to support 192 ports (4 VDSL2 r2 line cards) or up to 384 vectored ports (up to 8 VDSL2 r2 line cards).

The E7-2 VDSL2 r2 line cards utilize the control communication channel that is provided in all E7-2 chassis to enable the 96 port system level vectoring capability. This dedicated vectoring communication channel is outside the data and management path so ring traffic and system management functions are not impacted by vectoring noise cancellation calculations. This 96-port system level vectoring capability is applicable to E7-2 VDSL2 r2 combo and data only line cards that are installed in the same chassis.


The 96-port system level vectoring capability is supported in stand alone and modular chassis configurations. Both Modular Chassis Controller and Modular Chassis Expansion configurations are supported.

Higher density system level vectoring capability is supported by adding a dedicated vectoring control processor. The E7-2 Vector Control Processor (VCP) line cards are equipped with a Vector Control Interface port that connects the E7-2 VDSL2 r2 line cards to the Vector Control Processor. Similar to the 96-port system level vectoring capability, the Vector Control Processor card is dedicated to handling the vectoring noise cancellation calculations and is outside the data and management path. The Vector Control Processors are equipped with a Vector Control Interface port that connects to the Vector Control Interface port on each of the E7-2 VDSL2-48 r2 line cards. The Vector Control Interface is a 10Gbps communication channel that passes the vectoring calculations to the vector control processor. As the Vector Control Processor line cards are outside the data path they can be deployed to support multiple E7-2 VDSL2 r2 line cards in Stand Alone or in Modular Chassis configurations. Typically, higher density E7-2 VDSL2 configurations will be in Modular Chassis configurations. The Vector Control Processor line cards are equipped with 2 GE interface ports for network management that are connected to the GE enabled ports of the E7-2 line cards or may be connected via the Ethernet management port of the E7-2 Chassis. The E7-2 VDSL2 r2 line cards offer complete scalability from a single line card Unit Level Vectoring implementation (48 Ports), to two line card System Level Vectoring (96 Ports), and multiple line card System Level Vectoring (Up to 384 Ports). Calix recommends that the FTTN sites have one vectoring implementation to simplify the installation and management. While it is technically feasible to have a mix of vectoring control engines (processors) at the same location (multiple ULV or SLV vector groups) each vectoring group would need to be groomed in the OSP cable to maximize the effectiveness of the vectoring implementation. Vectored VDSL2 signaling that is not managed by a common vector control processor is seen as uncontrolled noise within the binder group to the other vectoring control engine.

As service providers deploy vectoring across their networks the E7-2 offers an ideal architecture for vectoring deployments across their networks. Vectoring is supported on a single E7-2 VDSL2 r2 line card using unit level vectoring. As additional VDSL2 subscribers are turned up in that location, an additional E7-2 VDSL2 r2 line card can be added without the need for an additional vector control processor. The same VDSL2 r2 line cards can also support high density VDSL2 nodes in 192-port and 384-port configurations by adding a vector control processor. This flexible deployment architecture provides additional options and growth for network wide vectoring deployments. Please refer to the *Calix DSL Best Practices* document and Vectoring to Win white paper for additional information on vectoring deployments.

Interface Capacity

The Calix E7-2 VCP-192 and E7-2 VCP-384 line cards are equipped with two and four 20Gbps vector control interface ports that connect to the individual E7-2 VDSL2 r2 line cards. Each 20Gbps vector control interface port establishes a dedicated 10Gbps link to the supported E7-2 VDSL2 r2 line cards. The vector communication ports transfer vectoring signaling information from the supported VDSL2 r2 line cards to the Vector Control Processor. The Vector Control Processor calculates the VDSL2 power and signaling frequencies for all line cards associated with the vectoring group. The E7-2 Vector Control Processor manages the power and signaling frequency information and is outside of the data and content path. The E7-2 Vector Control Processor line cards are also equipped with two GE management ports and should be connected to the E7-2 line cards in the Modular Chassis Controller shelf that are co-located at the VDSL2 installation site. The management connections are used for software maintenance and alarm support on the E7-2 Vector Control Processor line cards.



Calix E7-2 VCP-192 with E7-2 VDSL2-48 r2 cards



Calix E7-2 VCP-384 with E7-2 VDSL2-48 r2 cards

Vectoring Fallback for an SLV Target

Calix DSL Best Practices recommend that non-vectorized VDSL2 modems be configured in ADSL2+ mode to avoid active disturbers in the vectoring group. The presense of these non-vectorized VDSL2 modems in the binder noise profile lowers the optimal bandwidth that can be gained by vectoring.

In R2.6.30 and earlier, when a Vectoring Control Cable (VCC) link failure is detected by a VDSL2 card configured as a System-Level Vectoring (SLV) target, all xDSL ports on the card fallback to non-vectoring mode after two minutes. The delay allows the system to disregard short interruptions in VCC connectivity.

In R2.6.40 the DSL Vector Fallback feature is introduced that applies a designated DSL template to members of an SLV vectoring group and retrains the xDSL ports to non-vectoring when a VCC failure is detected. You can also specify a delay time until the fallback behavior occurs.

The DSL template applied to the card during the fallback to non-vectoring mode should have the service type = adsl2+. This is because ports in fallback mode using VDSL2 service type can act as disturbers to adjacent binder groups using VDSL2 in vectoring mode. You can also adjust other line parameters in the applied template. However, the the Join Vectoring Group parameter must be disabled (unselected), and the PTM Override parameter is not used.

The DSL Vector Fallback feature applies to the following VDSL2 r2 cards and nodes:

- VDSL2-48 r2
- VDSL2-48C r2
- E3-48 r2
- E3-48C r2

Operation Events

When a VCC disconnect is detected on an SLV card, the “Fallback to Non-Vectoring” event is generated after the fallback delay. The xDSL ports start retraining in a non-vectorized mode using the parameters defined for the applied DSL template.

Note: If no DSL template is specified to apply to the card at a VCC link failure, all vectoring lines retrain in non-vectorized mode using the individual port settings.


When the VCC connection is restored, all xDSL ports immediately retrain to VDSL2 vectoring mode and the “Recovered from Fallback” event is generated.

Configuration Overview


- a.** Configure a VDSL2 SLV mode configuration.
- b.** Create a DSL template to apply at vectoring fallback.
- c.** Apply the DSL template to the VDSL2 card or node.
- d.** Set the Fallback delay. (Optional)

Vectoring Performance

VDSL2 Vectoring is most effective on short VDSL2 loops measuring up to 3,000 feet / 915 meters. Vectoring restores the performance that is lost due to cross talk. The performance curve below measures 384 VDSL2 vectored and non-vectored ports at the distances of 1,000, 2,000 and 3,000 feet / 305, 610, and 915 meters. The testing shows a 400 24AWG /0.5mm pair cable with 16 dead pairs. The performance curve is useful to estimate the potential vectoring gain at different loop lengths. The performance gains comparing the Calix test plant versus a service provider OSP cable will vary due to copper outside plant conditions including age, guage, splices, bridge taps, water in the outside plant, as well as other environmental considerations.



Vectoring performance is calculated by two critical variables measuring DS/US performance gains as well as the vector group join time. A typical VDSL2 modem trains in an average of 60 seconds after which it is joined to the vectoring group. The vectoring control processor synchronizes the VDSL2 signaling to optimize the vectored VDSL2 performance. During initialization all modems are not passing traffic until they are joined to the vectoring group. The vectoring join times are specified by TR-249 allowing the vector control processor to join all modems to the vector group. The VDSL2 vectored modem firmware is a critical element to optimize the vectoring performance. Calix has conducted testing of several different modems and observed join time differences across different modem manufacturers and modem firmware. The chart below summarizes the join times of several modem types. Several items can be identified in the chart below including the concept of the regular interval when vectored modems are joined to the vector group. The bulk of modems joined in roughly 250 seconds whereby the balance remained in a joining state. The 2nd group of modems were joined to the vector group roughly 200 seconds later. The amount of time required to join the vector group goes down as there are fewer modems that are joining the vector group. The 3 modems that did not join the vector group in the first 3 attempts may indicate an outside plant line condition or require a modem firmware update.



Third-Party Modem Interoperability

Calix continues to test VDSL2 and ADSL2+ modems against the different DSL Service Platforms supported. These modems must be commercially available and candidates for service deployment by Calix customers. Extensive rate / reach and data throughput performance testing is performed against all of modems. Qualification testing is conducted for physical layer, HIS / data as well as IPTV / Video service performance. Customer Advisory Bulletins (CAB) documents are available from Calix summarizing completed testing; new versions of the document will provide updates to tests that are in progress. Please refer to up to date CABs from Calix for the latest list of qualified and supported modems, which can be located on the Calix documentation portal.

Appendix A

Reference Information

This chapter provides reference information which may be useful for E7 engineering and planning.

For more information regarding Ethernet access network planning, see the document *Calix Engineering & Planning Guide: L2 Ethernet Access Networks*, available on the Calix Resource Center.

Topics Covered

This chapter covers the following topics:

- E7 System Support Capacities
- GPON Support Capacities
- xDSL Support Capacities
- Power Consumption and Heat Dissipation
- Protocols and Standards Supported
- ODC Cabinet Solutions

System Support Capacities

The Calix E-Series system support capacities follow.

Description	Capacity
General System	
User accounts, locally defined (not in the RADIUS)	100
Simultaneous Netconf sessions for web browser interface	15 The Netconf interface has a 30 minute timeout, which cannot be turned off. The web browser interface has an inactivity timer of approximately 30 minutes.
Simultaneous CLI sessions	5 There is an ENABLE/DISABLE flag in the CLI for session timeout, but the only ENABLE timeout duration supported is 30 minutes. There might temporarily be an extra Netconf (16) and an extra CLI session (6), for just the duration of the login process. If the extra session is accepted, then a previous, oldest session is dropped under the assumption that this session is most likely to be idle.
CLI Command Buffer	4000 characters for a single copy/paste operation
SNMP Trap destinations (defined by IP address)	5
RADIUS	4 authentication RADIUS servers and 4 accounting RADIUS servers are supported All of the authentication RADIUS servers are assumed to have the same authentication information, that is, server replication. The system communicates with the "best" server, and then only sends to the next server if it does not get a response. Where "best" is determined by the success rate of getting responses to recent requests.
MAC address Table	32,000 (E7 line cards share a common table.) 128,000 for E7-20 SCP2-10GE See GPON Capacities for MAC address limits on ONTs.
DHCP Leases	32,000 per system
Switching Capacity	Wire speed full duplex forwarding across all ports. Dedicated, non-blocking switch port to all VDSL2, GPON, GE, and 10GE interfaces.
Default "Native" VLAN for untagged traffic (GE and 10GE ports) on Trunk and Edge interfaces.	1 - This VLAN is utilized to switch all untagged traffic through the system. This VLAN cannot be deleted, but can be changed. Not supported on GE-24x GE, ONT Ethernet ports, or E7 Access interfaces.
Bandwidth Policing	1 Mbps up to line rate for GE and 10GE ports 64 kbps to 2 Mbps in 64kbps increments; 1 Mbps increments from 2 Mbps up to 1000 Mbps for GPON ONT Ethernet ports
Egress Priority Queues per 1GE or 10GE port (not ONT Ethernet ports)	8 per port (GE and 10GE) ports based on P-bit value with P-bit = 7 highest priority 4 per PON (GPON) P-bit values are mapped into four GPON CoS queues
Queue Scheduling Algorithm	Strict priority across 8 queues, with maximum and minimum guaranteed bandwidth per class. Tail drop is used when dropping packets from queue.

Description	Capacity
Broadcast Video Channel Support	- E7-20: 2000 (4000 for SCP2-10GE) - E7-2 xDSL: 2000 - E7-2 GPON: 4000 (for GE-24, GE-12, and 10GE-4 cards) By default 10 to 12 entries are reserved for control traffic on a per VLAN basis.
Ethernet port mirrors	1
VLANs	
Maximum number	4090 (VLANs 1002, 1003, 1004, and 1005 are reserved for system use but can be changed to another range, VLAN 1 is for untagged traffic)
Default Internal VLAN (GE and 10GE ports only)	1 - This VLAN is utilized to switch all untagged traffic through the system. This VLAN cannot be deleted, but can be changed. Not supported on ONT Ethernet ports.
Tag Actions per card (ONT service tag actions are not included in this limit)	768
Maximum number of service VLANs with DHCP Snoop enabled	per VDSL2 line card = 48 per system = 256
VLANs that may be IGMP enabled	32
VLANs that may have PPPoE profile assigned (E7-2 only)	24 Calix recommends less than 512 PPPoE sessions per line card.
VLANs that may have MAC-Forced-Forwarding and/or IP-Src-Verify enabled (per VDSL card only)	8
VLAN IP Host	2
10Gig SFP+/XFP modules and GPON Optical Interface Modules (OIMs)	Must be keyed
Ethernet services	
Services allowed on an xDSL port	6
Services allowed on an ONT Ethernet port	8 on an ONT Ethernet port 4 on an ONT RG interface 1 on an ONT FB interface
Services allowed on a PON interface	2048
Services on a VDSL card that may resolve to a VLAN marked as TLAN (this includes the same TLAN on multiple services)	24

Description	Capacity
Ethernet Frame Size	
MTU Maximum Transmission Unit size (bytes)	E7-20 SCP2-10GE (R3.0) supports the ability to set the MTU Maximum Transmission Unit size (bytes) on a GE and 10GE port interface to a maximum of 9390 bytes, not including Ethernet header and Two VLAN tags for Q-in-Q. The E7 supports the ability to set the MTU Maximum Transmission Unit size (bytes) on a GE and 10GE port interface to a maximum of 9600 bytes, not including Ethernet header and Two VLAN tags for Q-in-Q. GPON ONTs and xDSL ports have a fixed MTU value. • MTU = 9600 bytes (E7 Ethernet interfaces) • MTU = 9600 bytes (E7 Backplane links) • MTU = 2000 bytes (700GE and 760GX GPON ONTs) • MTU = 1600 bytes (700GX GPON ONTs) • MTU = 1500 bytes (GigaCenter ONTs) • MTU = 1500 bytes (E7 xDSL ports) The MTU is defined as the maximum size payload of the Ethernet frame, not the Ethernet frame size. In an IP network, this is the largest IP packet that can be transmitted on the Ethernet network without IP packet fragmentation. The Ethernet frame size varies depending on the number of VLAN tags applied to the payload, plus allowances for Preamble/Delimiter and interframe gap. • + 8 bytes (Preamble/Delimiter) • + 14 bytes (header not including VLAN tags) • + 4 bytes (inner VLAN tag) • + 4 bytes (outer VLAN tag) • + 4 bytes (trailer) • + 12 bytes (interframe gap) For example: • 2000 bytes = MTU • 2026 = Max Ethernet frame (with two VLAN tags without Preamble/Delimiter) • 2034 = Max Ethernet frame (with two VLAN tags including Preamble/Delimiter) • 2046 = Max Ethernet frame (with two VLANs, preamble/Delimiter, and Interframe Gap)
Maximum throughput: Protocol efficiency for Ethernet	Protocol efficiency = Payload size ÷ Frame size Maximum efficiency is achieved with the largest allowed payload size. For example: • 1500 bytes (Maximum payload size) • + 8 bytes (preamble) • + 14 bytes (header) • + 4 bytes (trailer) • + 12 bytes (interframe gap) • = 1538 • 1500 (payload size) ÷ 1538 (frame size) = 97.53%
Maximum throughput: Efficiency for optional 802.1Q tagged Ethernet packets, include 4 bytes in the frame size	1500 (payload size) ÷ 1542 (frame size) = 97.28%
Maximum throughput: Protocol overhead for Ethernet as a percentage	Protocol overhead = 1 — Protocol efficiency

Description	Capacity
Maximum throughput: IP payload throughput	Payload Throughput = Efficiency * Net bit rate Where the physical layer net bit rate (the wire bit rate) depends on the Ethernet physical layer standard, and may be 10 Mbit/s, 100 Mbit/s, 1 Gbit/s or 10 Gbit/s. Maximum throughput for 100BASE-TX Ethernet is consequently 97.53 Mbit/s without 802.1Q, and 97.28 Mbit/s with 802.1Q.
ERPS/G.8032v2 Ring	
ERPS/G.8032v2 rings per E7-2 system	6 of each
ERPS/G.8032v2 rings per E7-20 system	2 6 for a pair of SCP2-10GE
ERPS/G.8032v2 rings per VDSL card	2
ERPS/G.8032v2 domain per E7 interface	1
"units" per ERPS/G.8032v2 ring	32 —or— 16 (for E7-2 MC nodes with VDSL2-based MCCs) This number counts each E7-20 SCP card located in the ring as 1 unit, and each E7-2 card located in the ring as 1 unit, whether in a dual- or single-card E7 shelf. This number does not include devices or E7s subtended from the ring, including the MCE shelves of an E7-2 modular chassis (MC) node.
Interconnected rings for a "chain of ERPS rings"	3
Unit count per E7-2	1
Unit count per E7-20 SCP card	1 (up to 2 units per E7-20 shelf)
Unit count per E7-2 line card	1 (up to 2 units per E7-2 shelf)
Link Aggregation Groups	
Link Aggregation Groups per shelf using GE ports	16
Active ports per Link Aggregation Group using GE ports	8
Ports per Link Aggregation Group using GE ports	8
Link Aggregation Groups per shelf using 10GE ports	2
Ports per Link Aggregation Group using 10GE ports	4 8 for E7-20 SCP2-10GE
Traffic Rate Limiters	
Per E7 card	1500
Independent policy rules	1500
Per card with RSTP protection	1500
Profiles and Templates per E7	
Policy Maps	256
Policies per Policy Map	1500
Policy Map Match entries for a VDSL card	24
Policy Map Match entries for all other card types	1536
Class Maps	1500
Class Rules per Class Map	100
Subscriber bandwidth profiles (GPON and VDSL together)	300
Multicast profiles	32
MVR profiles	16 at the system level (1 video provider per MVR profile)

Description	Capacity
Multicast Whitelists	128 16 Multicast Whitelists to be associated with an individual subscriber 32 Ranges per Whitelist X 16 Whitelists per service = 512 Ranges ONTs do not support 512 Ranges per service: • T-Series SFU support 64 ranges • T-Series MDU support 128 ranges • 700 Series (G, GX, GE) support 128 ranges • 836GE RSGs support 128 ranges • GigaFamily supports 128 ranges • VDSL supports QQQ ranges
Match lists (GPON and VDSL together)	255 A match list can contain both "tagged" and "untagged" match rules, up to 12 tagged rules and up to 16 untagged rules for each Ethernet port. 32 OUIs are allowed in an untagged match list.
Service tag actions	256
DSL port templates	512
SIP Remote Configuration Profiles (GPON)	512
SIP Gateway Profiles (VDSL)	512
TDM Gateway profiles	32
Dial plans	20 Maximum Network Dial Plan Table size is 100 rows x 28 (2800 bytes). A dial plan rule cannot exceed 28 bytes (or characters), because a rule must fit in a single row. The required " " character at the end of each rule limits the rule to 27 characters. A rule is not allowed to overlap rows.
Ds1PWE3 profiles	8
DSCP maps	10
IP Precedence maps	10
H.248 Gateway profiles	32
MGCP Gateway profiles	32
ONT PWE3 profiles	4
ONT profiles	200
PPPoE profiles	50
VLAN IGMP profiles	20
VLAN DHCP Proxy Service profiles	2
Security profiles	16 Security profiles used by VDSL interfaces (including bonded-Links) must have DHCP lease limit = 10 or less

xDSL Support Capacities

This topic covers the following subjects:

- E7-2 VDSL2 Deployment Capacities
- General xDSL Support Capacities

E7-2 VDSL2 Deployment Capacities

When planning the number of VDSL2 subscribers to support on a given node and ring, you must consider the entire E7 network deployment (the sum of all E7 nodes and subscribers, regardless of topology and access type). For guidelines regarding the recommended deployment capacities of an access network, see the document *Calix Ethernet Access Networks Engineering & Planning Guide*.

In general, if not in conflict with other network planning considerations, you may use the deployment capacities shown in the table below. Doing so will help ensure that E7-2 VDSL2 deployments are scalable, supportable, and perform as intended. Supported deployment configurations include single E7-2 shelves or E7-2 modular chassis nodes with GPON/AE-based MCCs, VDSL2-based MCCs, stand-alone or within ERPS rings.

Note that for E7-2 modular chassis with VDSL2r1-based MCCs (with VDSL2r1 cards) in an ERPS ring, the maximum VDSL2 combo or data-only subscribers/node and nodes/ring are lower than in other configurations. These capacity differences are shown in **bold** and apply to all nodes in a ring when at least one VDSL2r1-based MCC is present in the ring. These limitations don't apply if all nodes in the ring use either GPON/AE-based or VDSL2r2-based MCCs (with VDSL2r2 cards).

Supported E7-2 VDSL2 Deployments

	ERPS Ring, Single Shelves	Stand-Alone, GPON/AE MCC	Stand-Alone, VDSL2 MCC	(Note 1) ERPS Ring, GPON/AE MCC	(Note 1) ERPS Ring, VDSL2r2 MCC	ERPS Ring, VDSL2r1 MCC
Max Shelves/ Node	1	10	10	10	10	combo/data: 5 overlay: 10
Max Nodes/ Ring	16	N/A	N/A	16	16	8
Max Subs/ Node	combo/data: 96 overlay: 48	combo/data: 864 overlay: 432	combo/data: 960 overlay: 480	combo/data: (Note 2) overlay: 432	combo/data: (Note 2) overlay: 480	combo/data: 480 overlay: 480

Note 1) For modular chassis nodes with GPON/AE-based or VDSL2r2-based MCCs in an ERPS ring, this table assumes that the entire E7 network deployment consists of one ERPS ring, 16 nodes, < 8000 subscribers, and an average of 2.5 MACs per subscriber. This limits the total number of MAC addresses in the network to < 20,000 (the recommended best practice).

Note 2) Supports the maximum configuration of 864 (GPON/AE-based MCC) or 960 (VDSL2r2-based MCC) VDSL2 subscribers per node. However, in line with the example/assumption in note 1, Calix recommends that 480 VDSL2 subscribers per node be used for planning purposes on a 16-node ring.

DMT Trunk Speeds

Based on 24awg wire performance testing, the following results are determined.

DMT 8-pair bonded performance results (Mbps)							
Loop Distance	0	1000	2000	3000	4000	5000	6000
DMT 8-pair bonded DS	1,120	1,008	649	430	304	221	142
DMT 8-pair bonded US	480	459	297	122	62	25	9

General xDSL Support Capacities

Description	Capacity
MTU Packet Size (DSL services)	1536 Bytes total frame size
VDSL2 Bandwidth	Up to 140 Mbps downstream, 60 Mbps upstream (Profile 17a). Also applicable to downstream VDSL2 bonded capacity.
ADSL2+ Bandwidth	Up to 24 Mbps downstream, 3 Mbps upstream
VDSL2 Profiles	8a-d, 12a-b, 17a
ADSL Modes	ADSL2+, ADSL2, RE-ADSL, ADSL (G.dmt, G.lite, ANSI T1.413)
DSL Transport Modes	PTM and ATM (PTM only in VDSL2, ADSL2+ and ADSL2 service modes)
VC handling in ADSL Fallback	Single VC or Multi VC Single VC combines all services onto a single VC in between port and CPE. Multi VC supports up to 6 VCs between port and CPE with 1 to 1 mapping of VC per service (i.e. IPTV, Data, TR-69 mgmt., etc.)
Bonding Group Capacity	Up to 24 2-port bonding groups per card Up to 6 8-pair bonding groups per VDSL2 r2 Adjacent ports recommended for DSL Bonding. Bonding of ADSL2+ Fallback in ATM mode, requires adjacent or contiguous port bonding; [Odd(n), Even(n+1)]: [1,2], [3,4], [5,6], etc.

Description	Capacity
Adjacent xDSL ports bonded train-up rates	For E7-2 VDSL2 r1 cards With 12a, 17a profiles: • 120 Mbps Downstream • 60 Mbps Upstream With 8a profile: • 120 Mbps Downstream • 37 Mbps Upstream Note: Although bonded ports can achieve downstream train-up rates of up to 60 Mbps, bonded together, the group has a train-up rate of 100 Mbps. For E7-2 VDSL2 r2 cards With 17a profile and 8-pair bonding CPEs: • 1120 Mbps Downstream • 480 Mbps Upstream Note: The maximum DS rate configurable in the bandwidth profile can be up to 1200 Mbps, however, currently there are no CPEs available that can achieve this throughput.
Dual-Level Vectoring (DLV)	96 VDSL2 r2 ports in a vectoring group--two cards in the same shelf: • E7-2 stand-alone shelves • E7-2 Modular Chassis Controller (MCC) shelves • E7-2 Modular Chassis Expansion (MCE) shelves
Unit-Level Vectoring (ULV) or Board-Level Vectoring (BLV)	48 xDSL ports in a vectoring group: • E7-2 MC Controller / VDSL2 r2 • E7-2 MC Expansion / VDSL2 r1 or VDSL2 r2 • E7-2 stand-alone / VDSL2 r1 or VDSL2 r2 • E3-48C, E5-48, E5-48C, E3-48 nodes
System-Level Vectoring (SLV)	192 VDSL2 r2 ports in a vectoring group: • E7-2 MC Controller / VDSL2 r2 • E7-2 MC Expansion / VDSL2 r2 • E7-2 stand-alone / VDSL2 r2 • E3-48 r2, E3-48C r2 nodes Required VDSL2 Vector Control Processor: • E7-2 stand-alone / E7 VCP-192 • E3-48 r2, E3-48C r2 nodes / E3 VCP-192 384 VDSL2 r2 ports in a vectoring group: • E7-2 MC Controller / VDSL2 r2 • E7-2 MC Expansion / VDSL2 r2 • E7-2 stand-alone / VDSL2 r2 Required VDSL2 Vector Control Processor: • E7-2 stand-alone / E7 VCP-384
Bandwidth Shaping (DSL ports)	64 kbps increments

Description	Capacity
Number of Egress Priority Queues (DSL)	4 per port (P-bit values are mapped into four CoS queues)
Queue Scheduling Algorithm	Strict priority with maximum and minimum guaranteed bandwidth per class. Tail drop is used when dropping packets from queue.
MAC addresses per DSL port	128 maximum MAC addresses per VDSL2-48/-48C port
Service VLANs with MAC-Forced-Forwarding and/or IP-Source-Verify enabled	8 maximum per VDSL2 line card
Service VLANs with DHCP Snoop enabled	48 maximum number per VDSL2 line card
DSL Port Templates	512 maximum

VDSL2 Unit Support Capacities

The Calix E-Series VDSL2 unit system support capacities follow.

Description	Capacity
General System	
User accounts, locally defined (not in the RADIUS)	100
Simultaneous Netconf sessions for web browser interface	15 The Netconf interface has a 30 minute timeout, which cannot be turned off. The web browser interface has an inactivity timer of approximately 30 minutes.
Simultaneous CLI sessions	5 There is an ENABLE/DISABLE flag in the CLI for session timeout, but the only ENABLE timeout duration supported is 30 minutes. There might temporarily be an extra Netconf (16) and an extra CLI session (6), for just the duration of the login process. If the extra session is accepted, then a previous, oldest session is dropped under the assumption that this session is most likely to be idle.
SNMP Trap destinations (defined by IP address)	5
RADIUS	4 authentication RADIUS servers and 4 accounting RADIUS servers are supported All of the authentication RADIUS servers are assumed to have the same authentication information, that is, server replication. The system communicates with the "best" server, and then only sends to the next server if it does not get a response. Where "best" is determined by the success rate of getting responses to recent requests.
MAC address Table	32,000
Switching Capacity	Wire speed full duplex forwarding across all ports. Dedicated, non-blocking switch port to all GE and 10GE interfaces.
Default "Native" VLAN for untagged traffic (GE and 10GE ports) on Trunk and Edge interfaces.	1 - This VLAN is utilized to switch all untagged traffic through the system. This VLAN cannot be deleted, but can be changed. Not supported on GE-24x GE, Ethernet ports, or E7 Access interfaces.
Bandwidth Policing (rate limiting)	• 1 Mbps up to line rate for GE and 10GE ports
Egress Priority Queues per 1GE or 10GE port	• 8 per port (GE and 10GE) ports based on P-bit value with P-bit = 7 highest priority
Queue Scheduling Algorithm	Strict priority across 8 queues, with maximum and minimum guaranteed bandwidth per class. Tail drop is used when dropping packets from queue.
Broadcast Video channels using IGMP snooping	800
Ethernet port mirrors	1
VLANs	
Maximum number	4090 (VLANs 1002, 1003, 1004, and 1005 are reserved for system use but can be changed to another range, 1 is untagged)
Default Internal VLAN (GE and 10GE ports only)	1 - This VLAN is utilized to switch all untagged traffic through the system. This VLAN cannot be deleted, but can be changed.
Tag Actions per unit	768
VLANs that may be set to DHCP Snoop or Proxy	48
VLANs that may have PPPoE profile assigned	24
VLANs that may have MAC-Forced-Forwarding and/or IP-Src-Verify enabled	8
Ethernet services	
Services allowed on an xDSL port	6

Proprietary Information: Not for use or disclosure except by written agreement with Calix.

Description	Capacity
Services on a VDSL card that may resolve to a VLAN marked as TLAN (this includes the same TLAN on multiple services)	24
Ethernet Frame Size	
MTU Maximum Transmission Unit size (bytes)	The E7 supports the ability to set the MTU Maximum Transmission Unit size (bytes) on a GE and 10GE port interface to a maximum of 9000 bytes. • MTU = 9000 bytes (E7 Ethernet interfaces) • MTU = 1500 bytes (E7 xDSL ports) The MTU is defined as the maximum size payload of the Ethernet frame, not the Ethernet frame size. In an IP network, this is the largest IP packet that can be transmitted on the Ethernet network without IP packet fragmentation. The Ethernet frame size varies depending on the number of VLAN tags applied to the payload, plus allowances for Preamble/Delimiter and interframe gap. • + 8 bytes (Preamble/Delimiter) • + 14 bytes (header not including VLAN tags) • + 4 bytes (inner VLAN tag) • + 4 bytes (outer VLAN tag) • + 4 bytes (trailer) • + 12 bytes (interframe gap) For example: • 2000 bytes = MTU • 2026 = Max Ethernet frame (with two VLAN tags without Preamble/Delimiter) • 2034 = Max Ethernet frame (with two VLAN tags including Preamble/Delimiter) • 2046 = Max Ethernet frame (with two VLANs, preamble/Delimiter, and Interframe Gap)
Maximum throughput: Protocol efficiency for Ethernet	Protocol efficiency = Payload size ÷ Frame size Maximum efficiency is achieved with the largest allowed payload size. For example: • 1500 bytes (Maximum payload size) • + 8 bytes (preamble) • + 14 bytes (header) • + 4 bytes (trailer) • + 12 bytes (interframe gap) • = 1538 • 1500 (payload size) ÷ 1538 (frame size) = 97.53%
Maximum throughput: Efficiency for optional 802.1Q tagged Ethernet packets, include 4 bytes in the frame size	1500 (payload size) ÷ 1542 (frame size) = 97.28%
Maximum throughput: Protocol overhead for Ethernet as a percentage	Protocol overhead = 1 — Protocol efficiency
Maximum throughput: IP payload throughput	Payload Throughput = Efficiency * Net bit rate Where the physical layer net bit rate (the wire bit rate) depends on the Ethernet physical layer standard, and may be 10 Mbit/s, 100 Mbit/s, 1 Gbit/s or 10 Gbit/s. Maximum throughput for 100BASE-TX Ethernet is consequently 97.53 Mbit/s without 802.1Q, and 97.28 Mbit/s with 802.1Q.
ERPs Ring	
ERPS rings per E3-48C system	6

Description	Capacity
ERPS rings per E3-48C system	2
ERPS rings per VDSL card	2
ERPS domain per E3-48C interface	1
“units” per ERPS ring	32
Interconnected rings for a “chain of ERPS rings”	3
Unit count per E3-48C	1
Link Aggregation Groups	
Link Aggregation Groups per shelf using GE ports	6
Ports per Link Aggregation Group using GE ports	8
Link Aggregation Groups per shelf using 10GE ports	2
Ports per Link Aggregation Group using 10GE ports	4
Traffic Rate Limiters	
Per E3-48C	1500
Independent policy rules	1500
Per card with RSTP protection	1500
Profiles and Templates per E3-48/C r2	
Policy Maps	256
Policies per Policy Map	1500
Policy Map Match entries for a VDSL card	24
Policy Map Match entries for all other card types	1536
Class Maps	1500
Class Rules per Class Map	100
Subscriber bandwidth profiles	300
Multicast profiles	32
MVR profiles	5
Match lists	255
Service tag actions	256
DSL port templates	264
SIP Gateway Profiles (VDSL)	512
TDM Gateway profiles	32
Dial plans	20
Ds1PWE3 profiles	8
DSCP maps	10
IP Precedence maps	10
H.248 Gateway profiles	20
MGCP Gateway profiles	20
PPPoE profiles	50
VLAN IGMP profiles	20
Security profiles	16
	Security profiles used by VDSL interfaces (including bonded-Links) must have DHCP lease limit = 10 or less

Power Consumption and Heat Dissipation

The power consumption and heat dissipation of an E7-2 system are determined by the card lineup and the location of the E7-2 system. In a typical central office (CO) environment, the E7-2 fans run at low speed and draw less power than in a remote terminal (RT) location where temperatures may peak at 65° C.

Integrated power management on Calix E7 line cards:

- Redundant (A and B) DC battery feeds
- Independently monitored battery voltage
- Input Range: –42.5 VDC to –72 VDC
- Fuse: 7.5 Amps (A and B) – assumes standard rerating of fuses to no more than 80%

System Component	Power Draw (Watts)	Heat Dissipation (Watts)
10GE-4	55	55
GPON-8	75	75
GPON-8x	70	70
GPON-4 r2	75	75
GPON-4	75	75
GPON-4x	50	50
GE-12	50	50
GE-24	50	50
VDSL2-48C (Combo)	92	79
VDSL2-48 (Overlay)	72	69
VDSL2-48D (Data-Only)	72	69
VDSL2-48C r2 (Combo)	92	79
VDSL2-48 r2 (Overlay)	72	69
VDSL2-48D r2 (Data-Only)	72	69
VDSL2-192 VCP	72	72
VDSL2-384 VCP	119	119
E5-48 (Overlay)	72	69

System Component	Power Draw (Watts)	Heat Dissipation (Watts)
E5-48C (Combo)	92	79
FTA2 (CO nominal)	8.5	8.5
SCP-10GE	67	67
SCP2-10GE	65	64
E7-20 FTA (CO nominal)	30	8
E7-20 FTA (CO max)	200	55
E7-20 FTA2 (CO nominal)	30	8
E7-20 FTA2 (CO max)	350	83
E5 FTA (CO nominal)	8.5	8.5
E5 FTA (RT max)	48	48
FTA2 (RT max)	48	48
FTA (CO nominal)	22	22
FTA (RT max)	65	65
E3-8G	90	90

Powering Solution

The E3-48 r2 and E3-48C r2 units support international AC power input: 110-240VAC at 50/60Hz, -48vDC, or +/- 190vDC for Remote Powering. The table below shows the consumed power using different powering solutions and the different E3-48 r2 models. As a general rule, POTS/VDSL2 combo products require more power than Overlay (VDSL2 Splitter) enabled models and units with copper bonded uplinks draw more power than fiber only models. Nominal wattage calculations are based on a12 active POTS calls (9CCS) usage assumption with all DSL ports engaged.

	Power Input		
	-48vDC	AC	Remote Power
E3-48C r2 Combo with copper bonded uplinks	114	130	131 (164 at 36CCS)
E3-48C r2 Combo with fiber uplinks	108	124	124 (158 at 36CCS)
E3-48 r2 Overlay with copper bonded uplinks	81	93	94
E3-48 r2 Overlay with fiber uplinks	76	87	88 (158 at 36CCS)

Proprietary Information: Not for use or disclosure except by written agreement with Calix.

Protocols and Standards Supported


Standard	Title
TR-059	DSL Evolution - Architecture Requirements for the Support of QOS-Enabled IP Services
TR-101	Migration to Ethernet-Based DSL Aggregation
IEEE 802.1ac	Carrier Class Wi-Fi on 5 GHz band
IEEE 802.1ad	Bridged VLANs - 802.1Q Amendment 4, Provider Bridges (Q-in-Q or Stacked VLANs)
IEEE 802.1ag	Connectivity Fault Management
IEEE 802.1D	MAC Bridges
IEEE 802.1p	Priority
IEEE 802.1Q	Virtual Bridged Local Area Networks (VLAN)
IEEE 802.1w	Rapid Spanning Tree Protocol (RSTP)
IEEE 802.6	LAN/MAN Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications, Ethernet
IEEE 802.6ad	Link Aggregation
IEEE 802.6ae	10GBASE-xR/W, 10Gb/s Ethernet over fiber
IEEE 802.6ah	Ethernet in the First Mile
IEEE 1588 v2	Precision Clock Synchronization Protocol
IETF RFC 2132	DHCP Options and BOOTP Vendor Extensions
IETF RFC 2666	Definitions of Object Identifiers for Identifying Ethernet Chip Sets
IETF RFC 2827	Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing
IETF RFC 3046	DHCP Relay Agent Information Option
IETF RFC 3261	Session Initiation Protocol (SIP)
IETF RFC 3376	Internet Group Management Protocol, Versio 3 (IGMP v3)
IETF RFC 3393	IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)
IETF RFC 3619	Ethernet Automatic Protection Switching (EAPS) Version 1
IETF RFC 3637	Definition of Managed Objects for the Ethernet WAN Interface Sublayer
IETF RFC 3942	Reclassifying Dynamic Host Configuration Protocol version 4 (DHCPv4) Options
IETF RFC 3985	Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture
IETF RFC 4197	Requirements for Edge-to-Edge Emulation of Time Division Multiplexed (TDM) Circuits over Packet Switching Networks
IETF RFC 4243	Vendor-Specific Information Sub-option for the DHCP Relay Agent Option
IETF RFC 4251	Secure Shell (SSH) Protocol Architecture
IETF RFC 4541	IGMP Snooping
IETF RFC 4553	Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)
IETF RFC 4836	Definitions of Managed Objects for IEEE 802.3 Medium Attachment Units (MAUs)
ITU-T G.8021	Characteristics of Ethernet transport network equipment functional blocks
ITU-T G.8032v2	Ethernet Ring Protection Switching
ITU-T G.984	Gigabit-capable Passive Optical Networks (GPON)
ITU-T Y.1730	Requirements for OAM functions in Ethernet-based networks and Ethernet services
ITU-T Y.1731	OAM Functions and Mechanisms for Ethernet-based Networks
MEF 9	Abstract Test Suite for Ethernet Services at the UNI
MEF 14	Abstract Test Suite for Traffic Management Phase 1
IETF RFC 1901-1908	Community-based SNMPv2c
IETF RFC 3411-3418	User-based SNMPv3
IETF RFC 2578	Structure of Management Information Version 2 (SMIv2)

Standard	Title
IETF RFC 2863	Interfaces Group MIB
IETF RFC 5650	Definitions of Managed Objects for VDSL2
NFPA 72	National Fire Alarm and Signaling Code (only supported on Calix P-Series GE ONTs)

ODC Cabinet Solutions

The Calix ODC-100 can house multiple E-Series systems and accommodate deep fiber FTTP deployment, as shown in the following graphic. With the following features, GPON splitters and fiber termination can optionally be integrated to provide a complete E7 GPON or optical Ethernet cabinet configuration:

- Small cabinet for space restricted and lower subscriber locations
 - Up to 4 ft in height
 - Optional PON splitters
 - AC or remote powered options
 - Integrated optional EDFA
- Copper and/or fiber subscriber access
 - Supports E5-100's for ADSL2+/VDSL2 subscriber access
 - Supports E7 w/ GPON or active Ethernet and 10GigE transport
- Mounting options
 - Pad, pole, vault
 - Cross connect mountable



Larger Cabinet Solutions

The Calix ODC-1000, ODC-2000, and ODC-3000 cabinets are designed to serve larger subscriber densities and accommodate a wide variety of additional networking equipment in addition to the E7. The ODC-1000E and ODC-2000E configurations are designed for high density E7 configurations. Consult a Calix account representative for more information on cabinet configurations.



ODC-1000E and ODC-2000E Description

The ODC-1000E and ODC-2000E cabinets support high-density fiber and copper plant terminations, enabling deployment of high-capacity broadband applications from a single remote node location.


The Calix ODC-1000E cabinet is a small enclosure that can house multiple E-Series in a compact space, providing several mounting options. The ODC-1000E supports local or remote power configurations.

The Calix ODC-2000E cabinet is a mid-size enclosure that can house multiple E-Series, with sufficient space for additional equipment. The ODC-2000E is equipped with hinged front and rear racks that swing outward up to 90 degrees, providing easy rear access to all mounted equipment.

The ODC-1000E and ODC-2000E cabinet design supports incremental system expansion, lowering initial deployment costs while maintaining the capacity for future growth.

ODC-100 Cabinet Solutions

The Calix ODC-100 is an ultra-compact outdoor cabinet (ODC) enclosure, specifically designed for the Calix E-Series product family. The basic building blocks of the Calix ODC-100 are shown in Figure 15 (below). The ODC-100 is an enclosure that meets stringent GR-487 requirements, which helps ensure OSP longevity. The Calix ODC-100 is equipped with two doors – a plain door and a door with a standard 300-watt capacity heat exchanger – to ensure that the ODC-100 is ready for service deployments of up to 288 subscriber ports.



ODC-100 features and functionality:

- Dimensions: 37"H x 18"W x 24"D outside plant enclosure designed to meet GR-487 compliance. The optional battery base adds 11" to cabinet height (48"H) and a door mounted heat exchanger adds 4.75" to the cabinet width (22.75"W)
- Mounting options: Wall, pole, pad, or vault mount (pour in place template or pre-cast pad available)
- Powering options: Local AC (redundant 30A rectifier), or +/-190 VDC remote power (aka line or network power)
- Up to six field installable 48-pair 5-pin protector panels (CAT-5), each interchangeable with 48-port fiber distribution frames (SC connectors).
- Cabinet options: Battery base, batteries, battery warmer, generator connector, and integrated cross-connect
- Convenience outlet provided when the AC load center is included in the cabinet
- Integrated remote power hold-over capacitor

The Calix ODC-100 is designed to accommodate multiple E7-2 systems and is ideal for copper access deployment. All the necessary options can be integrated in this compact cabinet to provide a complete POTS and DSL cabinet configuration:

- Support for up to three E7-2 shelves equipped with any combination of line cards
- Integrated Copper or Fiber Backhaul
- AC Powered (shown below) or Line Powered
- Redundant 30A rectifiers
- 40Ah, 60Ah or 120Ah Battery Backup Capacity
- Integrated hold-over capacitor for remote power configurations
- Additional space for T1/E1, etc.
- Optional integrated 300 or 150 copper pair cross-connect
- Optional Generator Connector
- Pole, wall, vault or pad mount options
- Integrated FTTH fiber distribution options for GPON and AE



ODC-100 w/ 288 VDSL2 Combo Ports



ODC-100 w/ 192 VDSL2 Combo Ports and Copper Transport

Larger Cabinet Solutions

The Calix ODC-1000, ODC-2000, and ODC-3000 cabinets are designed to serve larger subscriber densities and accommodate a wide variety of additional networking equipment in addition to the E7. ODC-1000E and ODC-2000E configurations are designed for high-density E7 configurations.



Calix Outdoor Cabinets

Remote Terminal / Cabinet E7-2 DSL Deployments

The E7-2 is very well suited to meet the requirements of remote terminal deployments.

- Environmentally hardened, redundant architecture
- 1RU design is small yet can expand with an extra slot for very low first install cost
- Very high port concentration
 - 96 DSL / POTS Combo ports in 1RU
 - 48 DSL Overlay ports in 1RU
- Additional modular chassis are added with subscriber growth yielding a near linear cost curve
- Mix and match copper and fiber access / transport line cards in a common chassis meet needs of diverse remote access network
- Line cards are managed as a single chassis for operational efficiency
- Subscribers are aggregated and network resources efficiently shared across protected trunk facilities
- Resilient, hot-swappable fan tray requires less power

- Integrated aggregation and transport reduces the number of system that need to fit in the RT
- The E7 can be easily integrated into most existing outdoor cabinets to;
 - Upgrade the existing RT's capacity to VDSL2 ultra broadband copper service
 - Expand the existing RT's capacity in number of subscriber access ports with a compact 96 Combo or 48 Overlay DSL ports in spare 1RU
 - Extend cost effectively the reach of the RT with FTTx or mixed copper / FTTx service network.

ODC Cabinet Powering Options

The Calix E-series ODC cabinets offer the following power configurations:

Cabinet	Available Power Options
ODC-100	Local or Remote power
ODC-1000E	Local or Remote power
ODC-2000E	Local power only

Local (AC) power option

Locally powered systems are supplied by commercial AC power, and use an AC-to-DC rectifier to supply the equipment. You typically use a local power configuration when a power source or utility is readily available at the cabinet installation site.

Cabinets equipped for local power include an AC load center, rectifier shelf with integrated distribution, and a separate base compartment for housing batteries.

Remote (DC/line) power option

Remote powered systems are supplied by ± 190 VDC power, using a converter shelf inside the cabinet for ± 190 VDC to -48 VDC conversion.

Network remote power is a method for powering remote equipment over standard twisted-pair copper wires commonly used for telephony signaling, but using pairs that do not carry active telephony signals. Unlike local AC powering schemes that require an input source to be collocated near the equipment, remote power allows you to separate the power source from the equipment, enabling you to locate the equipment in places where commercial AC power is too difficult or too costly to reach. Further, remote power allows you to reuse existing copper plant infrastructure, bypass local utilities, and centralize battery maintenance.

Remote power solutions consist of counterpart upstream and downstream power shelves that convert between -48 VDC and ± 190 VDC power.

Cabinet Mounting Options

Mounting options for the Calix E-Series ODC cabinets include:

ODC-100

- Concrete pad mounting:
 - Pre-cast concrete pad
 - Site-cast concrete pad using Calix template
- Pole mounting using Calix pole-mount kit
- Adjunct enclosure mounting (for cross-connect cabinets and other large enclosures)
- Foundation vault mounting (third-party supplied)

ODC-1000E

- Concrete pad mounting:
 - Pre-cast concrete pad
 - Site-cast concrete pad using Calix template
- Pole mounting using Calix pole-mount kit
- Wall/H-frame mounting using Calix wall-mount kit

ODC-2000E

- Concrete pad mounting:
 - Pre-cast concrete pad
 - Site-cast concrete pad using Calix template
- Wall/H-frame mounting using Calix wall-mount kit

Cabinet Retrofit Kits

Calix has developed retrofit kits for installing E-Series units into legacy Calix cabinets and various third-party cabinets. These retrofit solutions have been thoroughly tested to characterize powering, heat dissipation / cooling and engineering requirements in an effort to minimize service provider installation and maintenance activities. Retrofit kits are available for both augment and 'clear and fill' installations.

Retrofit kits include the following features and options:

- Mounting solutions for the Calix E-Series equipment
- Copper line protection for E-Series service lines
- Cabling to connect to existing power supplies or options to upgrade the power system
- Heat exchanger/colling upgrade options, as required

For more details or ordering information, contact your Calix representative.