Calix

E7-2/E5-48/E3-48C R2.4 xDSL Applications Guide

December 2015

#220-00811, Rev 12

Contents

About This Guide

This Calix E7, E3-48, E5-48 and E5-48C Ethernet service access platforms leverage VDSL2 and ADSL2+ technology, along with an Ethernet switching fabric, to perform L2 aggregation / switching and layer-3-aware service delivery.

The Calix topics in this guide describe how to configure VDSL2-based services for the following platforms:

- E7-2 in stand-alone and modular chassis configurations
- E5-48, E5-48C and E3-48C in stand-alone configurations

Note: Calix E7-20 products do not support VDSL2-based cards or services.

Throughout this guide, E-Series is used to refer to the specific product of E7-2 that supports VDSL2 cards and the E3-48C, E5-48 and E5-48C products that supports the equivalent of the VDSL2 line cards.

For specific procedures on provisioning Ethernet elements, such as VLANs, Ethernet ports and interfaces on the E5-48, E5-48C and E3-48C, refer to the *Calix E3-48/E5-48/E5-48C User Guide*.

Intended Audience

This document is intended for use by network planning engineers, CO technicians, and craft and support personnel responsible for network equipment turn-up, service configuration, and maintenance. The procedures in this guide are of a technical nature and should only be performed by qualified personnel. Familiarity with standard telecom and datacom terminology and practices, as well as standards-based Ethernet technologies and conventions, is recommended.

Related Documentation

You can access Calix product documentation from the Calix Resource Center online at *www.calix.com* (*https://www.calix.com/portal/site/resourcecenter/*).

The Calix E7 documentation set includes:

Engineering and Planning

- *Calix Ethernet Access Networks Engineering & Planning Guide*

  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.

- *Calix E7 Engineering and Planning Guide*

  This document provides engineering and planning information for the Calix E7 Ethernet Service Access Platform (ESAP) and Calix E3-48C Ethernet Service Access Node (ESAN). It describes the features and capabilities of each system, and provides engineering guidelines to assist engineers and network planners effectively deploy the E7 and E3-48C.

Installation

- *Calix E7-2 Installation Guide*
- *Calix E7-20 Installation Guide*
- *Calix E3-48C Installation Guide*
- *Calix E5-48/E5-48C Installation Guide*

  These documents provide a general installation practice for the Calix Ethernet service access platform, including guidance for planning, power installation, cabling, and maintenance.

User Guide

- *Calix E7 User Guide*
- *Calix E3-48C, E5-48/E5-48C R2.3 User Guide*

  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

- *Calix E7/E5-48/E3-48C System Upgrade Guide*

  This document describes how to perform software/firmware upgrades as well as database backup and restore operations.

- *Calix E7 GPON ONT Upgrade Guide*

  For E7 GPON systems, this document describes how to perform upgrades for compatible P-Series, GigaCenter, and T-Series GPON ONTs.

Application Provisioning

- *Calix E7 GPON Applications Guide*
- *Calix E7-2/E5-48/E3-48C xDSL Applications Guide*
- *Calix E7 Active Ethernet Applications Guide*

  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

- *Calix E7 Maintenance and Troubleshooting Guide*

  This guide includes procedures for monitoring E7 network operation, general troubleshooting, and replacing or installing equipment.

Command Line Interface

- *Calix E3-48C, E5-48/E5-48C CLI Reference*

  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 E7 GPON*
- *Calix P-Series ONT Model/Feature Matrix*
- *Calix 800G GigaCenter Embedded Web Interface (EWI) User's Guide*
- *Calix T1 Pseudowire Applications Guide*
- *Calix C7 VoIP Services Guide*
- *Calix Application Note: Using the ONT VoIP Configuration File*
- *Completing Residential Gateway and SIP Configuration File Intake Forms*
- *Calix P-Series VoIP Configuration File - Template*
- *Calix 836GE RSG Wi-Fi Best Practices Guide*
- *Calix GPON RF Overlay Deployment Guidelines*
- *Calix Application Note: GPON Interface Adaptor*
- *Calix E7 Pluggable Transceiver Module Support*
- *CAB-12-023 - Pairing Bidirectional SFPs (to Support Single-Fiber Ethernet Links)*

Chapter 1

VDSL Applications Overview

This chapter describes VDSL2 applications and the Calix E-Series VDSL2 solution.

Topics Covered

This chapter covers the following topics:

- An overview of the E-Series xDSL support
- Description of modular chassis deployment with VDSL-based cards

xDSL Subscriber Access Support

The Calix E-Series Ethernet service platform leverages VDSL2 and ADSL2+ technology, along with an Ethernet switching fabric, to perform L2 aggregation / switching and layer-3-aware service delivery. Up to two dedicated 10 Gigabit and four 2.5 or 1 Gigabit Ethernet uplinks can be leveraged to provide high capacity service delivery on a per-chassis basis.

E7-2 VDSL2-48 Overlay Line Card

Interfaces:

- 48x DSL Overlay ports with corresponding 600- or 900-Ohm integrated splitters
- Four 1GE / 2.5GE SFP MSA sockets, and two 1GE / 10GE SFP+ MSA sockets

Features:

- The VDSL2-48 also supports pair bonding in both ADSL2+ and VDSL2 modes.
- One VDSL2-48 line card can be installed in an E7-2 chassis.
- Integrated splitters provide a means to extract analog voice form 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.
- Rev 13 and above of the VDSL2-48 card supports 600/900 Ω impedance, compatible with all E7 releases. There is no need to select between 600 and 900 Ω impedance and the splitter design accommodates either impedance without configuration.

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

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

E7-2 VDSL2-48 r2 Overlay Line Card

This card is backward compatible with the original E7-2 VDSL2-48 Overlay line card.

E7-2 VDSL2-48C Combo Line Card

Interfaces:

- 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. Pair bonding is supported in both ADSL2+ and VDSL2 modes.

E7-2 VDSL2-48C r2 Combo Line Card

This card is backward compatible with the original E7-2 VDSL2-48C Combo Line Card.

E7-2 VDSL2-48D Data-Only Line Card

Interfaces

- 48 ports VDSL2/ADSL2+ Fallback data-only
- 2 SFP sockets at 1GE/2.5GE rates
- 2 SFP+ sockets at 1GE/10GE rates

The Calix E7-2 VDSL2-48D card can be plugged into one or both of the two universal slots within a Calix E7-2 shelf to create a compact, high-density DSL node, with Ethernet aggregation and transport, ideal for copper based delivery of IP services across the access network. The line card supports pair bonding in both ADSL2+ and VDSL2 modes.

The Calix E7-2 VDSL2-48D Data-Only line card is ideally suited to meet new deployment requirements for the "All-Digital loop" as well as traditional Class 5 TDM Voice Switch consolidation applications. The All-Digital Loop provides high speed data services to all subscribers with the VoIP POTS functionality provided as an integrated component of the home residential gateway or as a separate IAD device attached to the modem. VoIP traffic is packetized in the home and then forwarded across the Calix network to the SIP enabled Softswitch. This application completely removes the POTS functionality from the service loop as the VoIP traffic is forwarded as a prioritized data packet. Using the E7-2 VDSL2-48D Data-Only line card Service Providers may continue to provide triple play services with the voice element being classified as "digital voice" as compared to "Life-line POTS".

In this application, the TDM Class 5 POTS and Splitters are external to the E7 solution. The E7-2 VDSL2-48D Data-Only solution does not include splitter functionality. POTS services in this application are provided by the Class 5 switch and the splitter function is external to the E7 solution. In this application, high speed digital traffic is routed across the E7 while POTS signaling is forwarded across the installed TDM network.

E7-2 VDSL2-48D r2 Data-Only Line Card

This card is backward compatible with the original E7-2 VDSL2-48D r2 Data-Only line card.

The Calix E7-2 VDSL2-48 card combines forty-eight VDSL2/ADSL2+ subscriber ports and corresponding integrated splitters, with four Gigabit Ethernet SFPs and two 10GE SFP+ ports, to provide high speed copper services with integrated Ethernet transport. Integrated splitters provide a means to extract analog voice on a 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, very high density DSL node, with Ethernet aggregation and transport, ideal for copper based delivery of IP services across the access network. The E7-2 VDSL2-48 supports a full set of subscriber services and network topology protocols.

Interface Capacity

The Calix E7-2 VDSL2-48C line card combines VDSL2 / ADSL2+ fallback subscriber ports and corresponding 900Ω integrated splitters, with GE SFPs and 10GE SFP+ ports, to provide high speed copper services with integrated Ethernet transport.

DSL Ports per card	Cards(s) per E7-2 Shelf	DSL Ports per E7-2 Shelf	Integrated VF Splitters	1GE/2.5GE SFP Ports	10GE/1GE SFP+ Ports
48	1	48	48	4	2

Integrated splitters provide a means to extract analog voice form 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. The E7-2 VDSL2-48 supports a full set of subscriber services and network topology protocols.

Broadband Overlay

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.

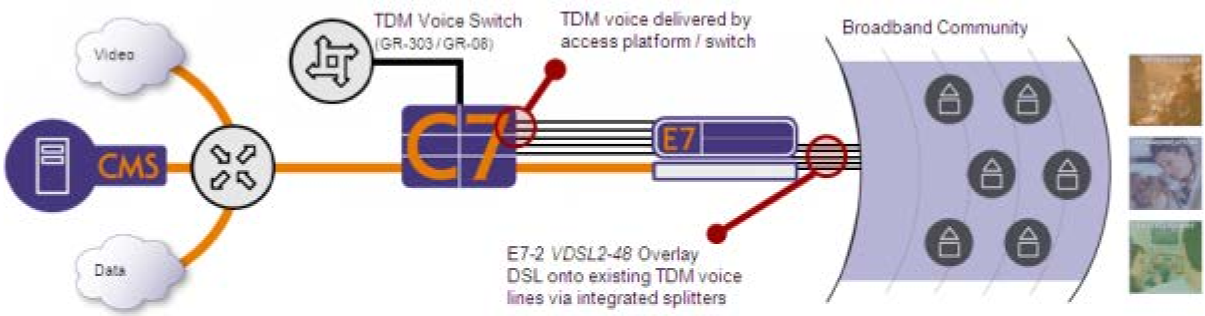

Overlay POTS Services with VDSL2-48

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 Ω + 2.16 µf

Loop current of up to 25mA (during normal operation), and 18mA (during battery backup)

E3-48C Node

- 48x Combo DSL and POTS ports, two 1GE / 2.5GE SFP MSA sockets, and two 1GE/10GE SFP+ MSA sockets
- The Calix E3-48C is a compact, hardened, high capacity sealed ESAN designed for outside plant applications.
- The E3-48C may be AC, DC, or Line powered.

E5-48 Overlay Node

- Fixed 1RU form factor
- 48 x VDSL2/ADSL2+ Fallback
- Integrated POTS splitter per port

- 4 x 1G ports for uplink/downlink
- Identical to E7-2 with a VDSL2-48 card, with the following exceptions:
  - Modular chassis is not supported
  - LAG supports only two ports in a group, either ports 1&2 or ports 3&4. LAG group may NOT include ports across the pair boundary. For example, a LAG with ports 1&3, or 2&3 are not supported.

E5-48C Combo Node

- Fixed 1RU form factor
- 48 x VDSL2/ADSL2+ Fallback ports
- 48 x POTS ports
- 4 x 1G ports for uplink/downlink
- Identical to E7-2 with a VDSL2-48C card, with the following exceptions:
  - Modular chassis is not supported
  - 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.

Multiple modes of DSL deployment

All xDSL ports on the VDSL2 line cards can be configured independently, for any of the xDSL modes of operation. The length and gauge of the copper pair loop determines the DSL mode(s) that can be supported for each port, and the train rate, and the derived service rate that the service provider is able to deliver to the subscriber. During the "handshake" process, the xDSL ports and the CPE agree to determine the DSL mode automatically. The E-Series trains the port at the DSL mode that provides the highest line performance possible, in the following order of priority:

1. VDSL2 PTM
2. ADSL2+ PTM
3. ADSL2+ ATM
4. ADSL2 PTM
5. ADSL2 ATM
6. ADSL ATM

xDSL can operate using either of two transmission convergence layers (PTM-TC, ATM-TC) and in one of several modes, listed above.

- Packet mode utilizes PTM encapsulation (Ethernet services) for xDSL port operation. All traffic on the port shares a single path and requires that all Basic Packet Functions (BPF) be performed in order to support varied services (VLAN) on the port. This is possible when using any of VDSL2, ADSL2+ or ADSL2 physical layers. PTM mode is not possible with ADSL1 physical layers (for example, G.dmt).
- ATM mode utilizes ATM encapsulation for xDSL port operation. In this case, the E-Series performs a SAR function to interwork packets to/from cells. ATM mode can be used with all xDSL physical layers, except VDSL2.
- Fallback mode is where the E-Series automatically determines that it cannot train up in VDSL2 mode, and then uses ADSL2+ instead. If the port can use PTM-TC, then it does use it. If the xTU-R does not support PTM-TC, the E-Series uses ATM-TC and performs the SAR function. Note that the E-Series currently supports up to six PVCs per xDSL port for ADSL operation.
- Legacy mode (this mode is only supported on the C7 and E5-series Calix products, and NOT supported on the E-Series) is distinctly different than the other modes in that ATM VC are connected directly to individual ATM VC on the DSL line, just as with a traditional ADSL card. As such, QoS is provided by the ATM VC, rather than on a per-priority, per VLAN port basis. Traditional CAC functions are performed to ensure that the aggregate capacity of the ATM VC feeding the port do not exceed the ports capacity. The exception being handling of multiple UBR VC on egress, wherein some rate shaping and priority queuing are required.

Bonding

The VDSL card supports bonding wherein a single “bonded” port can be created from multiple physical ports. Bonded ports are treated the same as individual ports. Therefore, all aspects of Basic Packet Functions (BPF) that are applicable to individual ports are also applicable to Bonded ports.

Vectoring

E-Series supports unit level vectoring on all VDSL2 equipment. Vectoring eliminates cross talk between VDSL2 lines and thus recovers 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.

Modular Chassis Deployment

The E7-2 MC platform operates as a ‘node’ such that it is managed as a complete entity; adding and deleting service, updating within a node concept, operationally managing troubles and diagnosis as a node. Configuration and installation of ‘stacked’ units to the node is automatically handled with ease of visual indicators to ensure cables are connected properly.

See the following for related information:

- "Turning Up a Modular Chassis System" in the *Calix E7 User Guide* for information on initially configuring several E7s into a modular chassis
- *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
- *Calix E7 Maintenance and Troubleshooting Guide* for instructions on how to add, replace, or delete components in a modular chassis system.

Considerations for VDSL2-based cards in a modular chassis

- The VDSL2-48C / VDSL2-48 line cards are supported components of the modular chassis architecture and platform deployment model.
- Modular chassis with VDSL2 cards support SFP ports at 2.5GE data rates for inter-chassis stacking ring connectivity.
- The VDSL2-48C card shares a common EXA Powered architecture and feature set common to all E7 cards and can be mixed interchangeably with other E7 cards to provide an application specific solution.
- Single-card Modular Chassis Controller (MCC) is supported when deploying a multi-shelf E7-2 with VDSL2-48 Overlay cards.

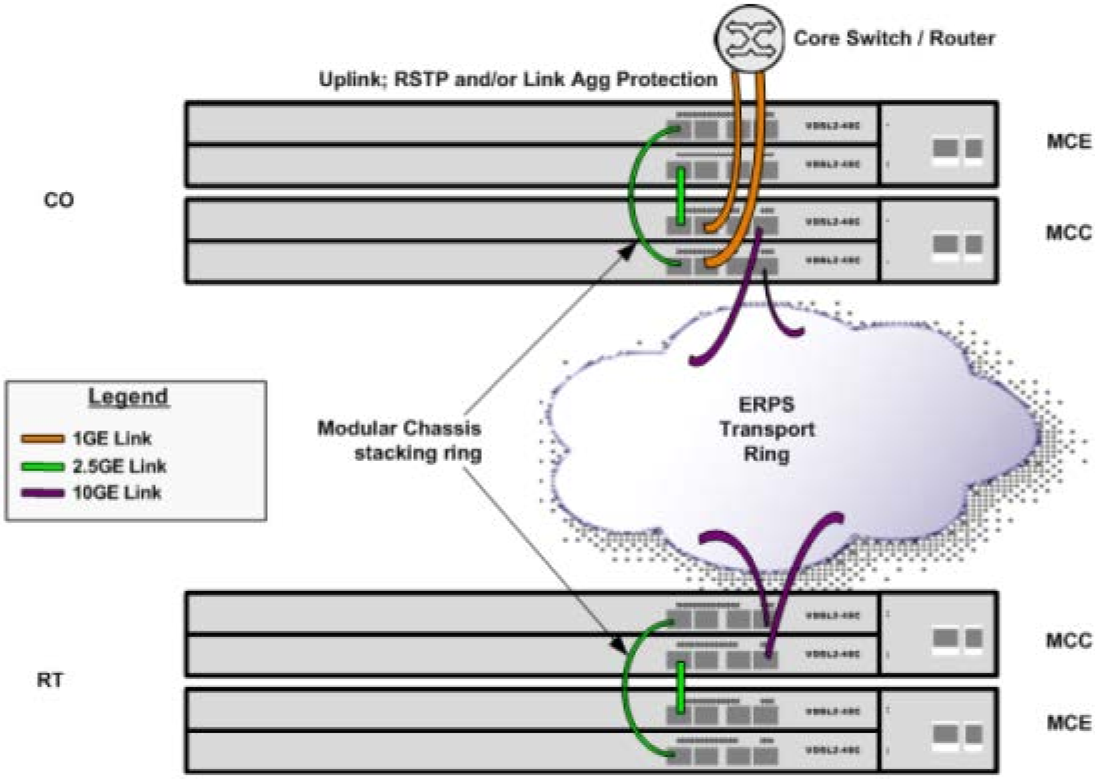

In the diagram above, VDSL2-48C line cards in the MCC and MCE shelves lose a forward facing SFP+ socket to the chassis' backplane. This is not the case with the VDSL2-48 Overlay line card, since it requires no backplane communication to a second card in the E7-2 same chassis. Both the RT and CO nodes are fully and only equipped with E7-2 VDSL2-48C line cards, while they could also use transport cards (for example 10GE-4) in the MCC to aggregate multiple ERPS rings or take advantage of longer range XFP optics. Both nodes support protected 10GE ERPS Transport and / or RSTP or Link Aggregation network protocols from the MCC shelf. See "Turning Up a Modular Chassis system" in the *Calix E7 User Guide* for configuration guidelines and detailed instructions for configuring a modular chassis system.

Chapter 2

Configuring VDSL2 Applications

This chapter describes how to setup VDSL2 applications, including configuring the E-Series system, creating profiles, and then adding subscriber services.

Note: For procedures on initial turnup and network configuration specific procedures, see the *Calix E7 User Guide* or the *Calix E3-48/E5-48/E5-48C User Guide*. See the *E7 Maintenance and Troubleshooting Guide* for a procedure on VDSL2 Single Ended Line Test (SELT) that is is a DSL pre-provisioning test tool used to assess the line capability prior to installing a modem at the customer premise.

Topics Covered

This chapter covers the following topics:

1. Configure network uplinks for VDSL2 services
2. Creating system profiles that support VDSL2 applications
3. Configure subscriber services

Step 1. Configuring the Network Uplink(s) for VDSL2 Services

This section describes how to configure the network uplinks for provisioned VDSL2 services.

Topics Covered

This section covers the following **topics in bold** that are part of the overall VDSL2 services configuration process:

1. **Configure network uplinks for VDSL2 services**
   - **Configuring Ethernet port interfaces**
   - **Configuring Ethernet ports**
   - **Creating service VLANs**
   - **Adding interfaces to VLAN memberships**
2. Creating system profiles that support VDSL2 applications
3. Configure subscriber services

Overview: Configuring the Network Uplink(s)

This chapter describes how to configure the network uplink(s) for E-Series VDSL2 services.

If the network contains a number of E-Series nodes, such as in a 10GE transport ring, the network uplink(s) may reside in a different shelf from the VDSL2 ports and may include multiple uplinks per service. Network uplink(s) are typically located in a shelf closest to the core or headend.

Note: For information on configuring system-level objects, such as NTP servers and SNMP traps, see the *Calix E7 User Guide* or the *Calix E3-48/E5-48/E5-48C User Guide*.

Configuration process

The network uplink configuration process follows:

1. Configure the 10GE or GE uplink port for service.
   - Set the Admin status for enable.
2. Configure the Ethernet interface on the uplink port.
   - Set the interface role to Trunk.
   - Enable RSTP for link protection as required.

3. Create the service VLAN(s).
   - Create one VLAN per subscriber for the 1:1 provisioning model.
   - Create one VLAN per service for the N:1 provisioning model.
   - For video service, enable IGMP Snooping.
4. Add the Ethernet uplink interface to the service VLAN memberships.

Configuring the Ethernet Uplink Port

The topic describes how to configure an E-Series Ethernet port for an uplink.

E-Series 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 *Configuring an Ethernet or LAG Interface* (on page 33).

The physical characteristics of the underlying ports include:

- Speed
- Duplex setting
- Interface type

ETH-Port names:

- g(port number) = Gigabit Ethernet Ports (GE)
- x(port number) = 10Gigabit Ethernet Ports (10GE)

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. (Before you can assign a port to a LAG interface, you must disable the port's default associated interface.)
- The interface provisioning (for example, VLAN membership) applies to the port when the interface is assigned to the port.
- 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.
- 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.

- The E7-2 supports 2.5 Gbps pluggable module interfaces in the SFP ports of the 10GE-4, GPON-4, VDSL2-48C, and VDSL2-48 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.

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

- The following rules apply for stand-alone E7 shelves and Modular Chassis Controller (MCC) shelves:
  - SFP and SFP+ Network ports; support protection protocols; ERPS, LAG, RSTP:
    - Transport ports; ERPS or RSTP protection
    - Trunk/Edge ports; RSTP and/or LAG protection, towards the edge switch/router or subtended device
  - SFP sockets:
    - 1GE or 2.5GE data rate modes
    - At 2.5GE data rate mode, supported port roles are Stacking/Edge/Trunk/Access
    - At 1G data rate mode; supported port roles are Edge/Trunk/Access
  - SFP+ sockets:
    - 1GE or 10GE data rate modes
    - At 10G data rate mode; supported port roles are Stacking/Edge/Trunk/Access
    - At 1G data rate mode; supported port roles are Edge/Trunk/Access
  - Support for subscriber DSL drops; supported port roles are Access
- The following rules apply for Modular Chassis Expansion (MCE) shelves:
  - SFP and SFP+ Access ports; no support for protection protocols
    - Aggregation; unprotected pt-to-pt links to subtended devices from SFP and SFP+ sockets

- SFP sockets
  - 1GE or 2.5GE data rate modes
  - At 2.5GE data rate mode; supported port roles are Stacking/Access (unprotected pt-to-pt)
  - At 1GE data rate mode; supported port roles are Access (unprotected pt-to-pt)
- SFP+ sockets
  - 1GE or 10GE data rate modes
  - At 10GE data rate mode; supported port roles are Stacking/Access (unprotected pt-to-pt)
  - At 1GE data rate mode; supported port roles are Access (unprotected pt-to-pt)

Before starting

Before starting the configuration process, check that the following conditions are met:

- The class map and class rules are configured.
- The policy map and policies are configured.
- The Ethernet or LAG interface is configured.
- The Ethernet port grade-of-service (GoS) profile is created.
- The class-of-service (CoS) profile is created.

Parameters

You can provision the following parameters for E-Series GE or 10GE Ethernet ports:

Parameter	Description	Valid Options
Admin State	Service state of port. While troubleshooting a port that has Admin State = "enabled-no-alarms," either use the CLI command "show alarm include suppressed," or from the web browser interface temporarily set the Admin State to "enabled," and then refresh the alarm panel manually or wait for default refresh rate to see the suppressed alarms.	enabled disabled ‡ enabled-no-alarms
Interface	Name of interface to select whether to leave the default connection to the port's logical (associated) Ethernet interface, or assign the port to a link aggregation group (LAG) interface. **Note:** Before you can assign a port to a LAG interface, you must disable the port's default associated interface. When a port is assigned to a LAG, the LAG interface provisioning (for example, VLAN membership) applies to the port.	EthIntf LagIntf
GOS Profile	Name of Grade-of-Service (GoS) profile to use that has been previously defined. The GoS profile specifies the operation thresholds for the Ethernet port.	Any established Ethernet GoS profile
COS Config	Name of class-of-service (CoS) profile to use that has been previously defined. The class-of-service profile specifies the queuing of packets.	Any established Ethernet port CoS profile

Parameter	Description	Valid Options
Broadcast Max Rate	Select whether to disable this parameter (off), or enter a value to specify the maximum rate for broadcast traffic (packets/second). This is an ingress rate limiter.	off, 0-16383999 Use "k" and "m" to multiply the rate.
Unknown Mcast Max Rate	Select whether to disable this parameter (off), or enter a value to limit the rate for unknown multicast traffic (packets/seconds). Use "k" and "m" to multiply rate.	off ‡ Enter Value = 0–16,383,999
DLF Max Rate	Select whether to disable this parameter (off), or enter a value for the maximum ingress 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.	off ‡ Enter Value = 0–16,383,999
LACP Priority	Priority value to use for determining which port to activate in a LAG. The lower value takes priority. For example, in a cross-card protection LAG, set the ports on the active card to a common LACP priority value, and then set the LAG ports on the standby card to another common priority value. The priority value on the active card ports must be lower than the value set for the ports on the standby card.	0-65535 32768 ‡
LACP Timeout	The length of timeout for LACP. **Note:** Avoid the use of LACP long timeouts, unless there is a very specific need.	short ‡ long
Duplex	Duplex mode for the port. • **Half-duplex** uses Carrier Sense Multiple Access (CSMA) to detect collisions and recover from them. • **Full-duplex** transmits and receives at the same time. **Note:** Use this setting for ERPS ring ports. • **Auto**: If the link is auto-negotiated, the duplex attribute is negotiated with the link partner. If the link speed is forced to a set value, full duplex is the default.	half, full, auto‡ (auto-negotiate duplex value with the link partner)
Flow control	Applies back pressure to a transmitter that is outrunning the receiver's capacity to process incoming data. • **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. **Note:** Use this setting for ERPS ring ports. • **Auto**: 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.	rx-tx, rx-pause, tx-pause, auto‡, none

Parameter	Description	Valid Options
LLDP Mode	Link Layer Discovery Protocol (LLDP) mode for the port. LLDP defines a set of information to be transmitted and received periodically on an Ethernet interface to and from connected devices. This information can be leveraged by the management interfaces to build a "network" topology view and identify all connected access nodes. **Note:** LLDP is enabled by default for all GE Access interfaces configured as Edge or Trunk.	disabled, tx-only ‡
Ethernet Speed (Mb/s)	Data rate of port (bits/s). Auto setting: • If the link supports auto-negotiation, the link partners auto-negotiate the speed while advertising the duplex and flow control parameters specified. • If the link does NOT support auto-negotiation, the setting is for the fastest rate that the module can support. Module-rate is for SFP+ ports, which supports both 10GE and 1GE modules. 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. Fixed speed setting forces the speed to the value specified. (**Note:** See the configuration guidelines above for the forced speed setting supported for various ports.)	auto‡, module-rate (native speed of pluggable module), 10mbps, 100mbps, 1gbps, 2.5gbps, 10gbps

‡ Default

To configure an E-Series GE port for service

1. On the Navigation Tree, double-click an E-Series line card, and then click a **GE** or **10GE** port.
   - Alternatively, you can access an E-Series Ethernet port using the following methods:
   - Click the triangle-arrow to the left of the service card in the Navigation tree, and then click the specific Ethernet port from the tree.
   - On the Navigation Tree, click **E-Series**, and then click **System** > **Ethernet Ports** and click a listed Ethernet port, or click in the row between the columns to edit the row. You can select multiple ports to edit using the **Control+click** and **Shift+click** keys. Click **Apply** when the parameter settings are complete.
2. Reference the table above to configure the parameters.
3. From the menu, click **Apply**.

For CLI:

```
set eth-port <port ID> [speed|duplex|flow-ctrl|interface|eth-gos|cos-queue-
cfg|bcast-max-rate|unk-mcast-max-rate|dlf-max-rate|lacp-priority|lacp-
timeout|admin-state]
```

Bulk Modifying Ethernet Ports

Modifying Ethernet ports in selected groups results in rapid configuration changes across a node.

To create a range of Ethernet ports

1. On the Navigation Tree, select the node, and then click **Ethernet Ports**.
2. In the work area, select the Ethernet ports to modify using **Shift+click** or **Ctrl+click**.

   To select a row from the table, click on the portion of the row in between the columns that has no text, as indicated by the red rectangles below.

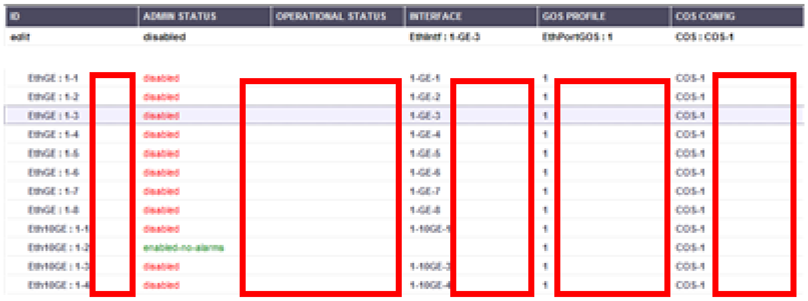

ID	ADMIN STATUS	OPERATIONAL STATUS	INTERFACE	GOS PROFILE	COS CONFIG
edit	disabled		EthIntf : 1-GE-3	EthPortGOS : 1	COS : COS-1
EthGE : 1-1	disabled		1-GE-1	1	COS-1
EthGE : 1-2	disabled		1-GE-2	1	COS-1
EthGE : 1-3	disabled		1-GE-3	1	COS-1
EthGE : 1-4	disabled		1-GE-4	1	COS-1
EthGE : 1-5	disabled		1-GE-5	1	COS-1
EthGE : 1-6	disabled		1-GE-6	1	COS-1
EthGE : 1-7	disabled		1-GE-7	1	COS-1
EthGE : 1-8	disabled		1-GE-8	1	COS-1
Eth10GE : 1-1	disabled		1-10GE-1	1	COS-1
Eth10GE : 1-2	enabled-no-alarms			1	COS-1
Eth10GE : 1-3	disabled		1-10GE-3	1	COS-1
Eth10GE : 1-4	disabled		1-10GE-4	1	COS-1

3. In the edit row at the top of the work area, select the parameter to modify, and select the new value.
4. Click **Enter** to identify each row with a modified value as having a pending change, indicated by an orange arrow.
5. In the Toolbar, click **Apply** to commit the changes.

For CLI:

```
set eth-port <port ID> [speed|duplex|flow-ctrl|interface|eth-gos|cos-queue-
cfg|bcast-max-rate|unk-mcast-max-rate|dlf-max-rate|lacp-priority|lacp-
timeout|admin-state]
```

Configuring an Ethernet Interface

This topic describes how to configure an Ethernet port's associated interface. E-Series Ethernet interfaces are logical objects that represent the service-related attributes of an Ethernet port.

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 *Configuring a Link Aggregation Interface* (on page 33) for information.

Interface names

Eth interfaces (Non-LAG related interfaces) share the same name as the E-Series Ethernet Ports (card 1/Eth port g1, for example 1/g1)

- g(port number) = GE(port number)
- x(port number) = 10GE(port number)

Interface roles

Each GE and 10GE interface in the system 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. 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. 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 or E-Series	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.

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 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.

Note: Node protection can be configured by using two E7-2 nodes on an ERPS ring. See Configuring RSTP Settings for the RSTP parameters to apply to the nodes.

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

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 LAG interface can be created, deleted, and modified.
- A port is always a member of exactly one interface, even when it is being used in a standalone manner.

- 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.
- 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.

Before starting

Before starting the configuration process, check that the following conditions are met:

- The policy map that you want to associate to the E-Series Ethernet port interface is already created.

Parameters

You can provision the following parameters for Ethernet interfaces:

Parameter	Description	Valid Options
Name*	Name of LAG interface or associated Ethernet port (card/Ethernet port). (**Note:** This is a case-sensitive string.)	String 31 characters
Role*	Role of E-Series Ethernet port interface. • **trunk** and **edge** are supported on E7-2 standalone, E7 modular chassis controller, and E7-20 SCP cards. • **access** is supported on E7-2 standalone, E7 modular chassis controller, E7 modular chassis expansion cards, and E7-20 linecards. See the rules and restrictions for each role in the above paragraphs.	trunk edge access E7-2, E7 Modular Chassis Controller, and E7-20 SCP cards: • 10G = trunk ‡ • 1G = edge ‡ E7-20 and E7 Modular Chassis Expansion: • 10G = access ‡ • 1G = access ‡ Default (LAG) = NA
Admin State	Service state of E-Series port interface. Select whether the interface is in service. **Note:** Before you can assign a port to a LAG interface, you must disable the port's default associated interface. When a port is assigned to a LAG, the LAG interface provisioning (for example, VLAN membership) applies to the port.	enabled ‡ disabled
Description	Descriptive name for the interface.	String 31 characters

Parameter	Description	Valid Options
RSTP	Whether the interface is running rapid spanning tree protocol (RSTP). 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. • RSTP (enabled) is only supported for trunk and edge interfaces. • RSTP (tunneled) is supported for all interface roles and should only be used when configuring TLAN service. • RSTP and BPDU Guard cannot be enabled on an interface, simultaneously. • For Modular Chassis systems, RSTP is only supported on MCC shelf interfaces. • For E7-20 systems, RSTP is only supported on SCP interfaces. • For cross-card LAG interfaces, RSTP is not supported and must be disabled in order to configure a cross-card LAG.	enabled (‡ E7-2) disabled (‡ E7-20) tunneled
STP Priority	Spanning tree protocol (STP) priority of this port interface. **Note:** Leave the default value unless you are certain of a modified value to assign as an STP priority for the interface.	0, 16, 32, 48, 64, 80, 96, 112, 128‡, 144, 160, 176, 192, 208, 224, 240
STP Path Cost	Spanning tree protocol (STP) path cost is the cost of transmitting a frame on to a network through that port. It is assigned according to the speed of the bridge. The slower the media, the higher the cost. **Note:** Leave the default value unless you are certain of a modified value to assign as the cost of service.	1-200000000 4 ‡
Policy Map	Name of policy map used on ingress traffic packets. The associated map describes what actions to take when the ingress traffic packets match the listed criteria.	name of map
Subscriber ID	Identification information of subscriber, such as phone number, or account number.	String 0-63 characters
MTU (Bytes)	Maximum Transmission Unit size (bytes). E7 supports the ability to set the MTU size on an interface to a maximum or 9000 bytes, not including the Ethernet header, two VLAN tags for Q-in-Q, and the frame check sequence (32-bit CRC). • MTU = 9600 bytes for E7 Ethernet interfaces • MTU = 2000 bytes for 700GE and 760GX GPON ONTs • MTU = 1600 bytes for 700GX GPON ONTs 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)	1500-9600 2000 ‡
Ether Type	The Ethertype indicates the protocol being transported in the Ethernet frame. • 0x8100 - IEEE 802.1Q-tagged • 0x88a8 - IEEE 802.3ad provider bridging • 0x9100 - Q-in-Q (double tagged) The recommended value of 0x8100 should be used for all interfaces (Ethernet and LAG). **Note:** The VLAN tagged frames are identified as having a tag by utilizing the Ethertype field.	0x8100 ‡ 0x88a8 0x9100

Parameter	Description	Valid Options
Native VLAN	Native VLAN to use for untagged user traffic on this interface. VLANs can be specified by name or by numeric VLAN ID. Supported for trunk and edge interfaces, only. To forward untagged traffic on E-Series Ethernet ports with an access interface, an add-tag action must be applied to untagged frames, assigning the traffic to a designated VLAN.	numeric value (range 1-4093) 1 ‡
Split Horizon Forwarding	Enable or disable split-horizon forwarding on this Edge interface of a standalone E7-2 or modular chassis controller shelf. (The Split Horizon Forwarding is not supported on the E7-20 system.) The default is “Enabled” or "Y" and should only be disabled when setting up TLAN service between multiple edge or access ports on the same E7. When Split-horizon forwarding is enabled on an edge port, traffic from that port will only route to trunk links. The Split-horizon forwarding flag has no effect on trunk links. 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. By default, E7 edge ports have the split horizon feature enabled which isolates port traffic from other edge ports within the same E7 line card.	disable enable (‡ E7-2)
BPDU MAC Mode	MAC for rapid spanning tree protocol (RSTP) bridge protocol data units (BPDUs). • **1d** results in the E7 transmitting BPDUs with a DA of 01:80:C2:00:00:00. Use this selection when the E7 is connected to an 802.1d compliant switch with redundant link. • **1ad** results in the E7 transmitting BPDUs with a DA of 01:80:C2:00:00:08. Use this selection when the E7 is connected to an 802.1ad compliant switch with redundant link.	1d ‡ 1ad
LACP Tunnel	LACP protocol packets are forwarded when set to enabled (Y) on an Ethernet interface and should only be used when configuring TLAN service.	disable ‡ enable
RSTP Auto Edge Link	Enables or disables automatic RSTP detection and protocol negotiation for other connected bridges on E-Series Ethernet ports. Each RSTP edge link transitions immediately to the RSTP forwarding port state, since there is no possibility of it participating in a loop. If another connected bridge is detected on a port with RSTP Edge Link enabled, the port immediately transitions to point-to-point and negotiates the RSTP topology.	disable ‡ enable
Trusted	Whether the interface is a trusted source of DHCP option 82/LDRA data. • Ethernet interfaces used for LAG or ERPS links can only be set to Trusted = Y. • Access interfaces can only be set to Trusted = N.	selected (yes) ‡ unselected (no)
BPDU Guard	Enables or disables BPDU guard mode. When enabled, this prevents topology loops by preventing the participation of an interface in a spanning tree. When the interface receives a BPDU (STP, RSTP, MSTP), it is put into a disabled state where an operator must manually disable, and then re-enable the interface to put the interface back into service. BPDU Guard and RSTP cannot be enabled on an interface, simultaneously.	disable (default for E7-2 standalone, MCC shelf, and E7-20 SCP with RSTP enabled on GE/10GE ports) enable (default for all Access interfaces, where RSTP is not supported)
IGMP Immediate Leave	Enables or disables IGMP immediate leave. When enabled, checks are omitted that would see if there are other hosts interested in the multicast group.	use-vlan-setting ‡ enabled disabled

Parameter	Description	Valid Options
Interface Quality Audit Mode	Mode to periodically check 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 one of the following modes: • **no-audit** - disables the Interface Quality Audit (IQA) mode • **alarm-only** - generates an alarm, but, does not take any action on the interface • **disable-interface** – Disable the interface when the threshold is exceeded • **protocol-action** – For ERPS and LAG, only disable the interface if there is an alternate path that is up and available. For non-ERPS and non-LAG interfaces, this is interpreted as "alarm only."	no-audit, alarm-only ‡, protocol-action, disable-interface
Polling Interval	Number of seconds between interface quality audits that compare errored frames to total received frames.	1-60 1 ‡
Error Threshold	Number of errored frames per million total frames to consider a specific interval as failed.	1-100000 1000 ‡
Polling Window	Number of interface quality audit intervals to consider for failure determination.	10-60 60 ‡
Errored Interval Count	Number of failed audit quality intervals within the polling window that will indicate an interface failure for IQA to take an alarm or OOS action.	1-60 10 ‡
Interval Min Frames	Minimum number of frames that must be received per interval for a specific interval to be considered valid.	1-2147483647 100 ‡

* Required field
‡ Default

To configure an Ethernet port associated interface for service

1. On the Navigation Tree, click a **GE** or **10GE** port where you want to configure an associated interface.
2. In the Work Area, click **Associated Interface > Provisioning**.
3. Reference the table above to configure the parameters.
4. Click **Apply**.

For CLI:

```
set interface <interface name> [eth-svc|role|description|subscriber-
id|native-vlan|mtu|rstp-active|rstp-prio|rstp-path-cost|rstp-bpdu-mac|rstp-
edge|bpdu-guard|immediate-leave|ingress-policy-map|split-horizon-fwd|lag-
mode|lacp-role|lacp-hash-method|lacp-min-ports|lacp-max-ports|lacp-system-
priority|lag-cross-card|lag-cross-card-revert|trusted|ethertype|admin-state]
```

Configuring a Link Aggregation Interface

This topic describes how to configure an Ethernet interface and a Link Aggregation Group (LAG) interface. E7 Ethernet interfaces are logical objects that represent the service-related attributes of an Ethernet port.

E7 Ethernet ports and the associated Ethernet interfaces always exist and can only be modified. LAG interfaces and their association with E7 Ethernet ports can be created, deleted, and modified.

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 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. 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 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 or E-Series	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 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 same 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

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 a 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 a 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 E7 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.

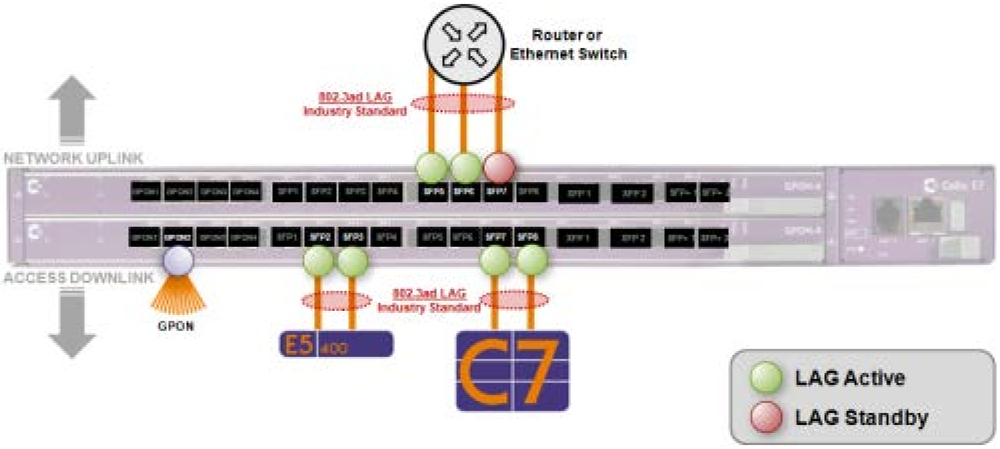

Example Cross-Card LAGs

The 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

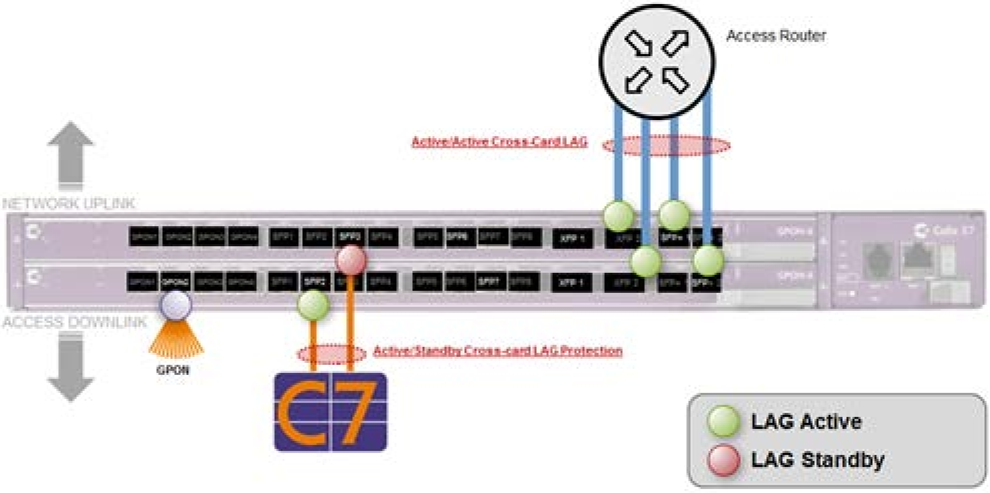

LAG per card and RSTP

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

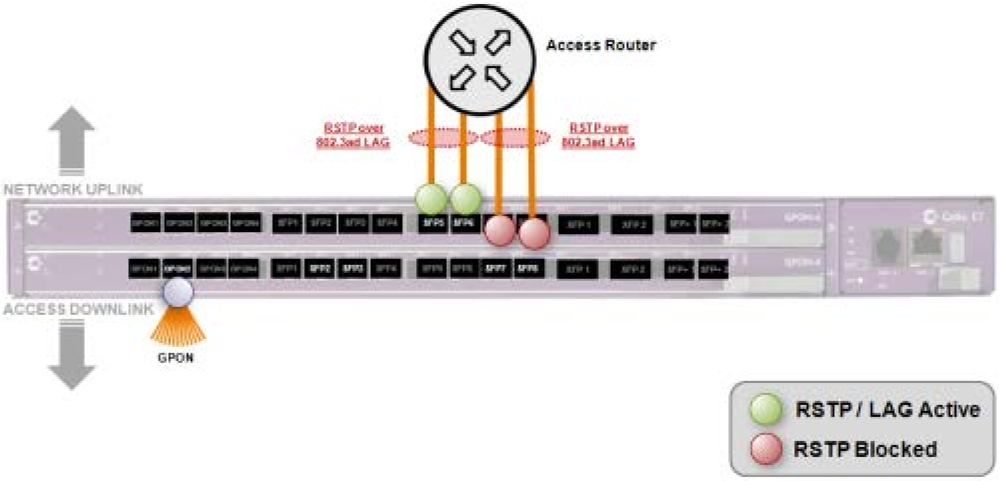

Facility and equipment protection using RSTP

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 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 on the E7, and the LAG end connected to a switch must also have RSTP disabled. If RSTP is enabled on the connected system, then that system will take 20-30 seconds to pass traffic again while waiting for RSTP to timeout.

Note: Node protection can be configured by using two E7-2 nodes on an ERPS ring. See Configuring RSTP Settings for the RSTP parameters to apply to the nodes.

To view the E7 Ethernet port interfaces that are actively participating in RSTP, click **E7** on the Navigation Tree, and then click the **RSTP** > **Interfaces** tabs.

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 LAG interface can be created, deleted, and modified.
- A port is always a member of exactly one interface, even when it is being used in a standalone manner.
- 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.
- 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.
- **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 the LAG 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* (on page 21) 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.)

Before starting

Before starting the configuration process, check that the following conditions are met:

- The policy map that you want to associate to the E7 Ethernet port interface is already created.

Parameters

You can provision the following parameters for LAG interfaces:

Parameter	Description	Valid Options
Name*	Name of LAG interface or associated Ethernet port (card/Ethernet port). (**Note:** This is a case-sensitive string.)	String 31 characters
Role*	Role of E7 Ethernet port interface. • **trunk** and **edge** are supported on E7-2 standalone, E7 modular chassis controller, and E7-20 SCP cards. • **access** is supported on E7-2 standalone, E7 modular chassis controller, E7 modular chassis expansion cards, and E7-20 linecards. See the rules and restrictions for each role in the above paragraphs.	trunk edge access E7-2, E7 Modular Chassis Controller, and E7-20 SCP cards: • 10G = trunk ‡ • 1G = edge ‡ E7-20 and E7 Modular Chassis Expansion: • 10G = access ‡ • 1G = access ‡ Default (LAG) = NA
Admin State	Service state of E7 port interface. Select whether the interface is in service. **Note:** Before you can assign a port to a LAG interface, you must disable the port's default associated interface. When a port is assigned to a LAG, the LAG interface provisioning (for example, VLAN membership) applies to the port.	enabled ‡ disabled
Description	Descriptive name for the interface.	String 31 characters

Parameter	Description	Valid Options
RSTP	Whether the interface is running rapid spanning tree protocol (RSTP). The E7 supports port-level RSTP. Therefore, ensure the far-end device is configured similarly (port-level RSTP) and not using VLAN-level RSTP. • RSTP (enabled) is only supported for trunk and edge interfaces. • RSTP (tunneled) is supported for all interface roles and should only be used when configuring TLAN service. • RSTP and BPDU Guard cannot be enabled on an interface, simultaneously. • For Modular Chassis systems, RSTP is only supported on MCC shelf interfaces. • For E7-20 systems, RSTP is only supported on SCP interfaces. • For cross-card LAG interfaces, RSTP is not supported and must be disabled in order to configure a cross-card LAG.	enabled (‡ E7-2) disabled (‡ E7-20) tunneled
STP Priority	Spanning tree protocol (STP) priority of this port interface. **Note:** Leave the default value unless you are certain of a modified value to assign as an STP priority for the interface.	0, 16, 32, 48, 64, 80, 96, 112, 128‡, 144, 160, 176, 192, 208, 224, 240
STP Path Cost	Spanning tree protocol (STP) path cost is the cost of transmitting a frame on to a network through that port. It is assigned according to the speed of the bridge. The slower the media, the higher the cost. **Note:** Leave the default value unless you are certain of a modified value to assign as the cost of service.	1-200000000 4 ‡
Policy Map	Name of policy map used on ingress traffic packets. The associated map describes what actions to take when the ingress traffic packets match the listed criteria.	name of map
Subscriber ID	Identification information of subscriber, such as phone number, or account number.	String 0-63 characters
MTU (Bytes)	Maximum Transmission Unit size (bytes). E7 supports the ability to set the MTU size on an interface to a maximum or 9000 bytes, not including the Ethernet header, two VLAN tags for Q-in-Q, and the frame check sequence (32-bit CRC). • MTU = 9600 bytes for E7 Ethernet interfaces • MTU = 2000 bytes for 700GE and 760GX GPON ONTs • MTU = 1600 bytes for 700GX GPON ONTs 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)	1500-9600 2000 ‡
Ether Type	The Ethertype indicates the protocol being transported in the Ethernet frame. • 0x8100 - IEEE 802.1Q-tagged • 0x88a8 - IEEE 802.3ad provider bridging • 0x9100 - Q-in-Q (double tagged) The recommended value of 0x8100 should be used for all interfaces (Ethernet and LAG). **Note:** The VLAN tagged frames are identified as having a tag by utilizing the Ethertype field.	0x8100 ‡ 0x88a8 0x9100

Parameter	Description	Valid Options
Native VLAN	Native VLAN to use for untagged user traffic on this interface. VLANs can be specified by name or by numeric VLAN ID. Supported for trunk and edge interfaces, only. 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.	numeric value (range 1-4093) 1 ‡
Split Horizon Forwarding	Enable or disable split-horizon forwarding on this Edge interface of a standalone E7-2 or modular chassis controller shelf. (The Split Horizon Forwarding is not supported on the E7-20 system.) The default is “Enabled” or "Y" and should only be disabled when setting up TLAN service between multiple edge or access ports on the same E7. When Split-horizon forwarding is enabled on an edge port, traffic from that port will only route to trunk links. The Split-horizon forwarding flag has no effect on trunk links. 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. By default, E7 edge ports have the split horizon feature enabled which isolates port traffic from other edge ports within the same E7 line card.	disable enable (‡ E7-2)
BPDU MAC Mode	MAC for rapid spanning tree protocol (RSTP) bridge protocol data units (BPDUs). • **1d** results in the E7 transmitting BPDUs with a DA of 01:80:C2:00:00:00. Use this selection when the E7 is connected to an 802.1d compliant switch with redundant link. • **1ad** results in the E7 transmitting BPDUs with a DA of 01:80:C2:00:00:08. Use this selection when the E7 is connected to an 802.1ad compliant switch with redundant link.	1d ‡ 1ad
LACP Tunnel	LACP protocol packets are forwarded when set to enabled (Y) on an Ethernet interface and should only be used when configuring TLAN service.	disable ‡ enable
RSTP Auto Edge Link	Enables or disables automatic RSTP detection and protocol negotiation for other connected bridges on E7 Ethernet ports. Each RSTP edge link transitions immediately to the RSTP forwarding port state, since there is no possibility of it participating in a loop. If another connected bridge is detected on a port with RSTP Edge Link enabled, the port immediately transitions to point-to-point and negotiates the RSTP topology.	disable ‡ enable
Trusted	Whether the interface is a trusted source of DHCP option 82/LDRA data. • Ethernet interfaces used for LAG or ERPS links can only be set to Trusted = Y. • Access interfaces can only be set to Trusted = N.	selected (yes) ‡ unselected (no)
BPDU Guard	Enables or disables BPDU guard mode. When enabled, this prevents topology loops by preventing the participation of an interface in a spanning tree. When the interface receives a BPDU (STP, RSTP, MSTP), it is put into a disabled state where an operator must manually disable, and then re-enable the interface to put the interface back into service. BPDU Guard and RSTP cannot be enabled on an interface, simultaneously.	disable (default for E7-2 standalone, MCC shelf, and E7-20 SCP with RSTP enabled on GE/10GE ports) enable (default for all Access interfaces, where RSTP is not supported)
IGMP Immediate Leave	Enables or disables IGMP immediate leave. When enabled, checks are omitted that would see if there are other hosts interested in the multicast group.	use-vlan-setting ‡ enabled disabled

Parameter	Description	Valid Options
Interface Quality Audit Mode	Mode to periodically check 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 one of the following modes: • **no-audit** - disables the Interface Quality Audit (IQA) mode • **alarm-only** - generates an alarm, but, does not take any action on the interface • **disable-interface** – Disable the interface when the threshold is exceeded • **protocol-action** – For ERPS and LAG, only disable the interface if there is an alternate path that is up and available. For non-ERPS and non-LAG interfaces, this is interpreted as "alarm only."	no-audit, alarm-only ‡, protocol-action, disable-interface
Polling Interval	Number of seconds between interface quality audits that compare errored frames to total received frames.	1-60 1 ‡
Error Threshold	Number of errored frames per million total frames to consider a specific interval as failed.	1-100000 1000 ‡
Polling Window	Number of interface quality audit intervals to consider for failure determination.	10-60 60 ‡
Errored Interval Count	Number of failed audit quality intervals within the polling window that will indicate an interface failure for IQA to take an alarm or OOS action.	1-60 10 ‡
Interval Min Frames	Minimum number of frames that must be received per interval for a specific interval to be considered valid.	1-2147483647 100 ‡
LACP Cross Card	(LAG only) Whether LAG cross-card protection is enabled, allowing ports on two cards to be configured into a LAG: • **active-standby** - one card's ports are active and the other card's ports are standby. • **active-active** - both cards' ports are active. **Note:** This only applies to cross-card protection LAGs.	disabled ‡ active-standby active-active
LAG Cross Card Revertive	(LAG only) Whether to have the LAG interface revert to the ports on the active card after a failure is found and fixed. **Note:** This only applies to active-standby cross-card protection LAGs. The active-active LAG does not support the revertive mode.	N (unselected) ‡ Y (selected)
LACP Hash	(LAG only) Individual traffic flows will only use a single link in the Link Aggregation Group (LAG). The link used for each packet is based on a hash algorithm. • **src-dest-mac** will likely give the best results for typical configurations that have many downstream client MACs, and few upstream server/router MACs. • **src-mac** or **dest-mac** may give positive results for a condition of badly uneven hashing with the typical configuration of many downstream MACs and few upstream MACs. Try src-mac hashing on the upstream LAG link, or try dest-mac hashing on a downstream LAG link.	src-mac dest-mac src-dest-mac ‡
LACP Min Ports	(LAG only) Minimum number of ports required for LAG activation. When the number of active ports falls below this value, the group is taken out of service. • When two LAGs are used with RSTP node protection, the system switches LAG operation to the other LAG. • For active-standby cross-card LAGs, the system switches the LAG operation to the ports on the standby card.	1-8 1 ‡

Parameter	Description	Valid Options
LACP Max Ports	(LAG only) Maximum number of active ports participating in the LAG. • For single-card LAGs, the ports added to the LAG that exceed this value are designated as standby ports and come online when an active LAG port fails. • For active/standby cross-card LAGs, this value indicates the number of ports on each card, assuming that there are an equal number of active and standby ports in the LAG. You cannot add more than this number of ports from any one card. The limit is 4. • For Active/Active LAGs, this value indicates the total number of ports in the LAG where all available ports must be active, so you cannot add more than the Max ports. The limit is 8.	1-8 8 ‡
LAG Mode	(LAG only) Mode for LAG interface. • **lacp-enable** - The LACP protocol is used to control the LAG ports. It is required for an Active/Standby LAG, or for any LAG that might have standby ports. • **manual** - The LACP protocol is not used. Therefore, if a port is added to the LAG, it is automatically made an active link when the port is enabled. This is also known as a static LAG.	lacp-enable ‡ manual
LACP System Priority	(LAG only) 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). • When the LACP system-priority is changed, a 2-second downtime will occur where no traffic passes through the LAG. • When provisioning active/standby LAG, the port priority should **NOT** be provisioned to 0. • For both ends of a cross-card LAG, the LAG ports on each card must all have the same LACP Priority value for the Ethernet port parameter, and the priority values must be different between the two cards. For active-standby cross-card LAGs, priority 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. • The port priorities on each side of the LAG should be set to the same values.	0-65535 32768 ‡
LACP Role	(LAG only) Role for this end of the LAG. • Active control mode actively initiates the LACP negotiations on a link. • Passive mode does not initiate LACP negotiations, but will respond. **Note:** Avoid the use of LACP passive role, 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.	active ‡ passive

* Required field
‡ Default

To create or configure a LAG interface for service

1. On the Navigation Tree, click **Interfaces**.
2. Either create or modify a LAG interface:
   - To create a LAG interface, click **Create**.
   - To modify an existing LAG interface, in the table of Ethernet port interfaces, double-click the row that shows the interface you want to configure.
3. Reference the table above to configure the parameters.

4. Click **Create** or **Apply**.

Note: Before you can assign a port to a LAG interface, you must disable the port's default associated interface. 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* (on page 21) for details on how to add ports to the LAG interface.

For CLI:

- `create interface <interface name> (name is case sensitive) role <trunk|edge>`
- `disable interface <interface name> (name is case sensitive)`
- `delete interface <interface name> (name is case sensitive)`
- `set interface <interface name> (name is case sensitive)`

Example:

```
create interface “MyLAG” role trunk description “LAG to 3750”
```

Creating the Service VLAN(s)

This topic shows you how to create a virtual LAN (VLAN) or transparent LAN (TLAN) to segregate traffic for different services or subscribers on the same network and separate device management control messages from traffic. See *Provisioning VLAN Ranges* (on page 62) if you want to create a set of VLANs with similar settings.

Native VLAN and untagged traffic

By default, every E-Series 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 E-Series 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 E-Series management traffic. Do not put management traffic on the Native VLAN (default is VLAN 1).

Note: For E-Series Ethernet services, the E-Series only passes VLAN tags provisioned on an Edge or Access link interface.

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.

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.

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.

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.

Calix VDSL2 service model

When a service is provisioned at an xDSL 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): 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.

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 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 Traffic flow in VDSL2 Cards

Unlike the Ethernet ports, xDSL Ethernet ports on the E-Serieshave no interface association. Consequently, 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 port. The VLAN must already be created on the E-Series for an xDSL port to pass traffic carried on a VLAN.

- To view the VLANs associated with specific port, click the port of interest on the Navigation Tree, and then click the **Associated Interface** > **VLANs** tabs.
- 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 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 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
  - 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 SNOOPERS	224.0.0.106

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.

- **IP Source Verification** ensures that only data from IP addresses learned by DHCP snooping or static provisioning are allowed to ingress 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.

Configuration guidelines

- 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 E-Series requires these VLANs always be identified. The default values for these VLANs are 1002, 1003, 1004, and 1005. If required, these four VLAN ID values may be reconfigured to be some other set of four consecutive VLANs, using the basic system settings.

- For an E-Series 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 E-Series Ethernet port interface:
  - VLAN membership
  - Tag-action (edge and access interfaces only)
  - Native VLAN (edge and trunk interfaces only)
- For an xDSL port to pass traffic carried on a VLAN, a service tag action must be created specifying the VLAN.
- For xDSL VLANs that need to carry IPv6 traffic, enable the VLAN TLAN parameter for transparent passthrough of IPv6 traffic.
- 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 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).
- 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 E-Series drops all DHCP server communication originating from subscriber Ethernet interfaces. Calix recommends DHCP snooping be enabled for all residential subscriber services.
- In the VDSL2 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.
- When DHCP snooping is enabled, DHCP requests for VLANs do not appear in an E-Series 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.
- Simultaneous operation of DHCP Snooping and PPPoE are not supported. When a PPPoE profile is selected, the DHCP features are disabled. See *Configuring PPPoAPPPoE Operation for a Data Service* (on page 184) for more information.
- For E-Series VDSL2, 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 following capacities apply to Static IP Addresses/Subnets within the E-Series 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 E-Series adding any subscriber MAC addresses to the E-Series 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.
- 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 E-Series VDSL2, MAC Forced Forwarding (MAC FF) can be used with the following IP hosts:
  - 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)
- MAC Forced Forwarding supports a single source MAC address with multiple source IP addresses, when the MAC address is not specified.
- IP Source Verify is not supported for a source MAC address with multiple IP addresses.
- For E-Series VDSL2, MAC FF can be used with static IP subnets.

- 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).
- The default behavior for data services is to filter all multicast traffic upstream from an 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.
- 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.
- For the double-tagged VLAN, the same inner tag cannot be used with different outer tags. Any given (inner) tag received upstream on the VDSL2 card can only be told to push a single (outer) tag. For example, VLAN 100 received on the VDSL2 card can push tag 200; 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 would require VLAN 300 to push an outer tag.
    - It cannot not be used as the outer tag in a double-tagged VLAN because that would mean it would get popped in the downstream direction.
- Multicast addresses on an E7 VDSL2 card must be unique across all IGMP-enabled VLANs on the VDSL2 card, including MVR VLANs in up to 4 MVR profiles. For example, multicast address w.x.y.z cannot exist in both VLAN-A and VLAN-B on the same VDSL2 card, where IGMP snooping/proxy is enabled on VLAN-A and VLAN-B.
- An MVR VLAN cannot also be used as an Ethernet services VLAN.
    - DSL only supports a single MVR VLAN if the IGMP mode is set to snoop-suppress and will support multiple MVR VLANs if the IGMP mode is set to proxy.

Parameters

You can provision the following parameters for VLANs:

Parameter	Description	Valid Options
ID*	VLAN ID	2-4093 (Except for 1002-1005 which are reserved for E-Series operation.)
Name	Name of VLAN	String of 31 characters

Parameter	Description	Valid Options
IGMP Mode	IGMP mode for the VLAN. **snoop-suppress** (This parameter does not apply to the E7-20, as it only supports proxy for IGMP): 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. 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 E-Series 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._ **Note: Calix recommends IGMP proxy for the multicast VLAN in all nodes in an ERPS ring.** **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. Note: Recommended for VLANs not carrying multicast traffic.	snoop-suppress proxy flood ‡
IGMP Profile	Name of the IGMP profile to associate with the VLAN. This profile only applies if the IGMP mode is set to proxy.	system-default ‡ any existing profile
DHCP Snoop	Whether to enable DHCP snooping to track all DHCP activity on that VLAN and create/update a table of DHCP leases granted. Calix recommends DHCP snooping be enabled for all residential subscriber services. **Note:** • When DHCP Snoop is enabled, MAC Learning is disabled by default. • Maximum number of service VLANs with DHCP Snoop enabled, per VDSL2 line card = 48.	cleared (disabled) ‡ selected (enabled)

Parameter	Description	Valid Options
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. 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.	cleared (disabled) ‡ 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) ‡ selected (enabled)
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) ‡
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) ‡ 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 limit 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) ‡ 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. 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.	cleared (disabled) ‡ selected (enabled)

Parameter	Description	Valid Options
PPPoE Profile	Assign a previously-created PPPoE profile. When a PPPoE profile is selected, the DHCP features are disabled. See *Creating a PPPoE Profile* (on page 146) for instructions. 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). **Note:** Simultaneous operation of DHCP Snooping and PPPoE are not supported.	any available PPPoE profile

*Required field
‡ Default

To create a VLAN

1. On the Navigation Tree, click **VLANs**.
2. In the Work Area, click **Provisioning** > **Create** to open the Create VLAN dialog box.
3. In the ID box, enter a VLAN ID for the VLAN you are creating.
4. In the Name box, enter a name that is descriptive of the VLAN use. For example, VoIP or Management.
5. In the IGMP Mode list, select the IGMP mode. Calix recommends that VLANs used for video services be set to suppress-snoop or proxy (enable snooping).

   Note: DSL only supports a single MVR VLAN if the IGMP mode is set to snoop-suppress and will support multiple MVR VLANs if the IGMP mode is set to proxy.

   Note: The suppress-snoop parameter only applies to standalone <e-type> and Modular Chassis systems, as the E7-20 only supports proxy.

6. (Only applies if the IGMP mode is set to proxy.) In the IGMP Profile list, select whether to accept the system default profile or another existing IGMP profile. Also see *Creating an IGMP Profile* (on page 119).
7. In the DHCP Snoop list, select whether to enable DHCP snooping. Calix recommends DHCP snooping be enabled for all residential subscriber services.
8. In the MAC Forced-Fwd list, select whether to enable MAC Forced Forwarding. Calix recommends MAC Forced Forwarding be enabled for all residential and business internet access services.
9. In the IP Src Verify list, select whether to enable the IP source verification (station validation).
10. In the MAC Learning list, select whether to enable the MAC address learning for the VLAN.

    Note: This setting only applies to standalone E7-2 systems, as the E7-20 and modular chassis VLANs only support MAC Learning enabled.

11. In the AE Discover Event list, select whether the E7 sends an event to CMS whenever a new AE ONT is discovered.
12. In the TLAN list, select whether to enable transparent LAN service for an ONT Ethernet port or an xDSL port.

    Note: 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.

13. In the PON Hairpin list, select whether to enable PON hairpinning.

    Note: Hairpin of services in the PON requires additional resource allocation in the GPON subsystem and should not be enabled unless hairpin service is required.

14. In the PPPoE Profile list, select the name of the profile to use.
15. Click **Create**.

For CLI:

- `create vlan <vlan ID> [name|mac-learning|mac-forced-forwarding|ip-source-verify|pon-hairpin|pon-tlan|igmp-mode|igmp-profile|dhcp-snooping|ae-ont-discovery|pppoe-profile]`
- `delete vlan <vlan ID>`
- `show vlan [vlan id|detail|igmp-counters|members|onts]`
- `show vlan <vlan id> [detail|igmp-counters|mac|mcast|members|onts]`

Configuring DHCP Relay Option 82 and LDRA

This topic shows you how to configure the global DHCPv4 L2 Relay Agent (Option 82) and Lightweight DHCPv6 Relay Agent (LDRA) features. Option 82 applies to all VLANs with DHCP Snooping enabled, and LDRA applies to VLAN per Service data models with DHCP Snooping enabled. These features are enabled simultaneously 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.

The E-Series 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. Both IPv4 and IPv6 are supported for data services simultaneously on VLAN per Service models in all Layer 2 topologies when the following conditions:

- The "DHCP Snoop" is enabled on the S-VLAN
- “Option 82/LDRA” is enabled on the E7 or E3/E5 unit or shelf
- The Ethernet uplink interface is set for "Trusted=Y"

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, and are not user-defined.

When DHCP Option 82/LDRA is enabled, Relay Agent identification information is inserted into DHCP messages captured at the "Untrusted" interfaces and sent to the DHCPv6 server by applying a system-defined access-identifier profile.

- The "eth-system-default" profile is applied to Ethernet and xDSL ports.
- The "gpon-system-default" profile is applied to E7 GPON ONT Ethernet ports.

Ethernet interfaces are identified as either "Trusted" or "Untrusted."

- 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.

When DHCP snooping is enabled on a VLAN, the E-Series tracks all DHCP activity on that VLAN within the system 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: VLAN ID, IP address, MAC address, 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.

Upstream (client to server) DHCP packets captured on Untrusted interfaces, will have the following Option 82 information inserted:

Note: LDRA sub-options are derived from the Option 82 Circuit-ID and Remote-ID sub-options using the same format.

- 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)

- E7 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>:<cetag>[-<tag-Id>]
        - Calix-format-2: <system-ID>:<shelf>/<slot>/<port>/<OntID>/<Ontport>/
    - Remote-ID options:
        - 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').

Downstream (server to client) DHCP packets will be captured and examined on Trusted interfaces.

- If a session match is found for an interface, the Option 82/LDRA string is removed if Option 82/LDRA is enabled globally, 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.

Configuration guidelines

The following additional constraints only affect VLANs with DHCP Snooping that are provisioned on Ethernet and xDSL port interfaces and where the global DHCP relay Option 82/LDRA is enabled:

- DHCP snooping and Option 82/LDRA are not supported on the Management VLAN.
- Ethernet interfaces used for LAG or ERPS links cannot be set to “Untrusted.”
- If a line card reboot or reset occurs, the DHCP lease database for all ports (xDSL, GE/10GE, GPON) persists.

The following guidelines apply to LDRA only:

- LDRA cannot be enabled independently from Option 82.
- LDRA is automatically enabled on any VLAN with DHCP Snoop 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.
- E7-2 and E3-48C 10GE and E5-48 and E5-48C GE 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.

To configure DHCP Option 82/LDRA insertion

1. On the Navigation Tree, select **E-Series**.
2. In the work area, click **DHCP > Provisioning** to open the DHCP Configuration form.
3. In the Option 82/LDRA Enabled checkbox, select the checkbox to enable this feature.
4. In the Option 82 Policy list, select whether to drop or overwrite packets with Option 82 on ingress packets.
5. In the toolbar, click **Apply**.
6. To specify the Remote-ID or Circuit-ID attributes on the global option 82 profile for E-Series networks, use the procedure shown below.

For CLI:

```
set dhcp-cfg option-82 [enabled|disabled]
set dhcp-cfg option-82-policy [drop|overwrite]
```

To configure the global access-identifier profiles

1. On the Navigation Tree, select **E-Series**.
2. In the work area, click **Profiles > Access Identifier** to view the table of default access-identifier profiles.
3. Double-click the name of the profile that you want to configure:
   - **eth-system-default** is used for xDSL and GE ports.
   - **gpon-system-default** is used for GPON ONT ports.
4. In the Access Identifier Profile form, select the parameter from the attribute list:
   - Circuit ID list parameters:
     - **calix-format**
     - **tr101-format**
     - **calix-format-2**

- Remote ID list parameters:
  - **ONT FSAN serial number** 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 16 characters of the Subscriber ID text field are inserted.
  - **MAC Address** (for DOCSIS provisioning) of the port 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.
  - **none**

5. In the toolbar, click **Apply**.

For CLI:

```
set access-identifier-profile <eth-system-default|gpon-system-default>
remote-id [subscriber-id|fsan-serial-number|mac-addr|none]

set access-identifier-profile <eth-system-default|gpon-system-default>
circuit-id [calix-format|tr101-format|calix-format-2]
```

Provisioning VLAN Ranges

Provisioning VLANs in ranges allows you to efficiently create, update, modify, and delete VLANs. This feature also allows you to quickly create VLAN memberships to multiple Ethernet port interfaces.

To create a range of VLANs

1. On the Navigation Tree, select and expand the **E-Series** node, and then click **VLANS**.
2. In the Toolbar, click the Range Operations icon ▦▾, and then select **Range Create** to open the Create VLAN Range dialog box.
3. In the From and To boxes, enter the values of the first VLAN ID and last VLAN ID in the range.
4. Select the Stop On Failure box to stop the range creation operation if a failure occurs. For example, if VLANs within the specified range already exist.
5. Enter the parameters to apply to the VLANs in the range, and then click **Create**. The Task Progress dialog appears to show the status of each operation.
6. Click **ok**.

For CLI:

```
create vlan <vlan ID>-<vlan ID> [name|mac-learning|igmp-mode|igmp-
profile|dhcp-snooping|ae-ont-discovery]
```

To update a range of VLANs

1. On the Navigation Tree, select and expand the **E-Series** node, and then click **VLANS**.
2. In the Toolbar, click the Range Operations icon, and then select **Range Update** to open the Update VLAN Range dialog box.
3. In the From and To boxes, enter the values of the first VLAN ID and last VLAN ID in the range.
4. Enter only the parameter values that you want to update for the VLANs in the specified range, and then click **Update**. The Task Progress dialog appears to show the status of each operation.
5. Click **ok**.

   - Alternatively, to select an E-Series VLAN to edit, click in the VLAN table row between the columns where there is no text. You can select multiple VLANs to edit using the **Control+click** and **Shift+click** keys. Make any modifications using the top edit row. Click **Apply** from the toolbar when the modifications are complete.

For CLI:

```
set vlan <vlan ID>-<vlan ID> [name|mac-learning|igmp-mode|igmp-profile|dhcp-
snooping|ae-ont-discovery]
```

To delete a range of VLANs

1. On the Navigation Tree, select and expand the **E-Series** node, and then click **VLANS**.
2. In the Toolbar, click the Range Operations icon, and then select **Range Delete** to open the Delete VLAN Range dialog box.
3. In the From and To boxes, enter the values of the first VLAN ID and last VLAN ID in the range to delete.
4. Select the Stop On Failure box to stop the range deletion operation if a failure occurs. For example, if VLANs within the specified range do not exist.
5. Select the Force Deletion box to ensure VLANs and associated VLAN memberships are deleted.
6. Click **Delete**.

Note: A VLAN can not be deleted if it is referenced by a tag action, a service tag action, an ONT Ethernet Service, or an ONT IP Host.

- Alternatively, to select an E-Series VLAN to delete, click in the VLAN table row between the columns where there is no text. You can select multiple VLANs to delete using the **Control+click** and **Shift+click** keys. Click **Delete** from toolbar to delete the selected VLANs.

For CLI:

```
delete vlan <vlan ID>-<vlan ID>
```

To add interfaces or an ERPS domain to a range of VLAN memberships

1. On the Navigation Tree, select and expand the **E-Series** node, and then click **VLANS**.
2. In the Toolbar, click the Range Operations icon ▦, and then select **Range Action** > **Add VLAN Members** to open the Add VLAN Members dialog box.
3. In the From and To boxes, enter the values of the first VLAN ID and last VLAN ID in the range.
4. In the Add Interfaces/ERPS column, select the interfaces to add to the specified range of VLAN memberships.

   Note: Use **Shift+click** or **Ctrl+click** to select multiple interfaces from the column.

5. Click the direction button > to add the interfaces to the range of VLAN memberships, and then click **ok**.
6. The Task Progress dialog appears to show the status of each operation. Click **ok**.

For CLI:

```
add interface <interface name> to-vlan <vlan ID>-<vlan ID>
```

Related topic

- Creating VLANs

Adding the Uplink Interface(s) to VLAN Memberships

This topic shows you how associate E-Series Ethernet port interfaces, LAG interfaces, multicast router interfaces, and ERPS domains to a VLAN, thereby creating a VLAN membership. A VLAN membership enables traffic on a specific VLAN to be forwarded to or accepted on an E-Series interface or ERPS domain. For any VLAN-tagged traffic to flow through an E-Series 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.

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).
- An interface role cannot be modified when it is a member of a VLAN.
- ERPS domains can only be associated with VLANs through membership.
- 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 membership:

For incoming E-Series 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 membership 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 add it to a VLAN membership.

To make an interface a VLAN member

1. On the Navigation Tree, click **VLANs**.
2. Click the **Provisioning** tab.
3. In the table of existing VLANs, double-click the row showing the VLAN which you want to add members.
4. Click **Action** and then select the action to perform:
   - **Add/Remove VLAN Members** enables traffic on a specific VLAN to be forwarded to or accepted on an E-Series interface or ERPS domain.
   - **Add/Remove VLAN Router Interfaces** sets 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.
5. In the dialog box, do the following:
   a. In the Available Interfaces/ERPS scrolling list, click the interface or the ERPS domain to add to the VLAN. You can select more than one item, using the Ctrl+click or Shift+click key combinations.

   b. Click the **>** button to add the selections to the Current Members or Current Interfaces box.

Note: See the Configuration guidelines above.

Note: If the E7 system is set to Modular Chassis mode, the Ethernet port interfaces are indicated with a shelf/card/port location. For example, EthIntf:1-2-G1.

6. Click **Ok**.

For CLI:

```
add interface <interface-id> to-vlan <vlan-id>
add erps-domain <domain-id> to-vlan <vlan-id>
add static-mcast-src interface <interface-id> to-vlan <vlan-id>
```

Step 2. Creating VDSL2-Related Profiles

This section describes how to create various profiles that are necessary for provisioned VDSL2 services.

Topics Covered

This section covers the following **topics in bold** that are part of the overall VDSL2 services configuration process:

1. Configure network uplinks for VDSL2 services
2. **Creating system profiles that support VDSL2 applications**
   - **Creating Quality-of-Service (QoS) for VDSL2 traffic management**
   - **Creating profiles for data and video services**
   - **Creating profiles for voice services**
3. Configure subscriber services

Overview of VDSL2-Related Profiles

This section describes how to create system profiles to use for VDSL2 applications. Creating profiles allow you define common provisioning attributes that can be reused many times and applied to multiple service ports.

Note: For information about system profiles unrelated to VDSL2 services, refer to the *Calix E7 User Guide* or the *Calix E3-48/E5-48/E5-48C User Guide.*

Profiles referenced in service provisioning

Provisioning services at the VDSL2 card port allows you to set up information common to multiple subscribers, maintained via profiles. The service definitions at the subscriber port level reference existing profiles and sometimes provide subscriber-specific information to complement the profiles.

Note: When a profile is changed, all objects referring to that object are affected. For example, if the rate in a bandwidth profile is changed, all provisioned services referencing the profile are changed to the modified bandwidth rate. Profile objects may not be deleted while they are being referenced.

Global profiles in CMS

Global profiles automate synchronizing profiles across multiple E-Series nodes. They support cross-network capabilities such as bulk provisioning. You create a profile once within CMS and apply it across all targeted E-Series nodes to ensure consistency across large deployments.

When you create a global profile, the profile is automatically downloaded to the networks in the CMS management domain that enable global profile updates.

1. On the Navigation Tree, click **CMS**.
2. In the Work Area, click **Profile > E5-48/E3-48C/E7/ONT** (or **E3-48C/E5-48/E7**), and then select the profile to create.

Service tag actions:

Each service provisioned at the subscriber port includes a reference to a match list and a tag action that specifies the classifying and marking of packets from the subscriber port into the service VLAN.

A service carried on an N:1 VLAN applies to multiple subscriber ports, so a single match list and tag action can be used to describe the service.

A service carried on a 1:1 VLAN is the same for each subscriber except the customer tag is unique per subscriber. In this case, it would be ideal to 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 in the tag action that indicates the value of the output tag is subscriber specific. The customer-specific tag is contained in the service definition.

Bandwidth profiles:

Specifies the upstream and downstream rate limits for an individual service member.

A multicast profile references two previously defined profiles:

- **Multicast Address Map**

  Identifies the optional global allowable multicast IP ranges.
- **Multicast VLAN Registration (MVR) Profile**

  Identifies the optional MVR address ranges associated with specified multicast VLANs.

Security profiles:

Specifies security attributes of the xDSL port: DHCP lease limit, upstream broadcast/multicast limit, L2CP filter.

Static IP host addresses and subnets:

Configures a static Ethernet service IP address or subnet, to associate with an xDSL port service.

Multicast video profiles:

A video service is identified by the presence of a service definition, referencing a multicast profile. The multicast profile defines the following attributes of a video service.

- Maximum number of simultaneous streams on subscriber port.

Multiple video profiles will typically differ in the number of streams allowed. A small number of video profiles will be created and applied to many subscriber ports. For example, bronze, gold, and platinum video services.

IGMP profiles:

Defines a profile for VLAN association that sets configuration attributes of the Internet Group Management Protocol (IGMP) snoop used to establish membership in a multicast video services group.

SIP Gateway Voice profiles:

Specifies the SIP service configuration information that previously resided in a SIP configuration file in a primary and a secondary location.

TDM Voice Gateway profiles:

Specifies the port numbers for the UDP control plane traffic, and RTP voice traffic, as well as the IP address of the C7 viper card.

H.248 Gateway profiles and H.248 Gateway:

Specifies the H.248 gateway properties for the VDSL2 H.248 gateway services.

IP host profiles:

The IP host profiles define attributes for static IP hosts, and are used for SIP and TDM gateway services. It defines how the service obtains an IP address for communication.

Grade of Service (GOS) profiles:

Specifies reporting thresholds for certain monitored attributes of an xDSL port. For example, any time a particular count exceeds a specified threshold within a certain period (either 15 minutes or one day), a threshold-crossing alert is generated. For information on creating the profile, see the *Calix E7 Maintenance and Troubleshooting Guide* or the *E3-48/E5-48/E5-48C User Guide.*

Creating the Quality of Service for VDSL2 Traffic Management

The E-Series supports the ability to classify traffic based on the P-bit, DSCP, IP Precedence, or VLAN values of the incoming traffic. Additionally, the VDSL2 card can classify traffic based on the customer equipment MAC OUI value. The E-Series then marks the traffic with a new priority value.

Using the newly marked P-bit values, the egress queues on the port then deliver the traffic to the network using the strict priority queuing scheme (with minimum bandwidth guarantees if desired). Strict priority queuing means that all traffic (in the queues) with the highest priority is delivered to the network first, then lower priority queue traffic is delivered.

- The E-Series xDSL Access ports use match lists and service tag actions to classify subscriber traffic.
- The E-Series Trunk and Edge ports use class maps and policy maps to classify incoming network traffic.

This section describes how to create the following traffic-control objects for VDSL2 traffic management:

- **Match list and rules:** Defines a set of criteria (matching rules) for classifying traffic on an xDSL port. The match list defines how the VDSL2 card classifies subscriber traffic into a service VLAN. The match list is applied to a service tag action that is then associated with an xDSL port service.
- **Service tag actions:** Marks the traffic with the appropriate VLAN and priority value required by the service.
  - See *Creating an Ethernet Bandwidth Profile* (on page 105) for instructions on creating another traffic-control object for services provisioned on xDSL ports.
- **Class map and rules:** Defines the matching criteria for traffic on an Ethernet port that are specified by an associated classification map which lists rules to identify packets, such as VoIP traffic.
- **Policy map and policies:** Contains lists of QoS-related actions to perform on packets at the Ethernet port that match certain criteria.

Creating 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 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.

See "Traffic Management" and "Data Path" in the *Calix E7 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

Parameter	Description	Valid Options
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
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

To create a class of service

1. Access the profile page:
   - From CMS:
     - On the Navigation Tree, click **CMS**.
     - In the Work Area, click **Profile > E3-48C/E5-48/E7 > CoS Profile**.
   - Locally on E-Series:
     - On the Navigation Tree, click the unit.
     - In the Work Area, click **Profiles** > **CoS** > **Ethernet** > **Create**.
2. Reference the table above to configure the parameters.
3. Click **Create**.
4. Associate the class of service profile to an Ethernet port so traffic is sent at a particular rate by controlling the queueing of packets.
5. See *Configuring an Ethernet Port* (on page 21).

For CLI:

```
create cos-queue-cfg <queue name> [queue-*-rate|queue-*-min-bw|port-
rate|port-burst-size]
```

Creating VLAN Tag Actions on an Ethernet Interface

This topic describes how to create VLAN tag actions on an E-Series edge or access link interface. A tag-action allows you to add or change, or add and change VLAN tags on packets that are received 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.

See the *Calix E-Series Engineering and Planning Guide* for an example of an E-Series 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.

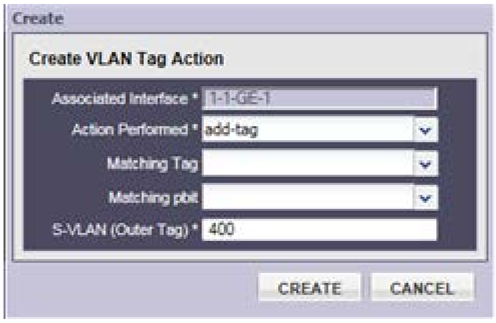

Create

Create VLAN Tag Action

Associated Interface * 1-1-GE-1
Action Performed * add-tag
Matching Tag
Matching pbit
S-VLAN (Outer Tag) * 400

Create an action to add a second tag (VLAN 400) to all VLAN 401 traffic on the associated interface.

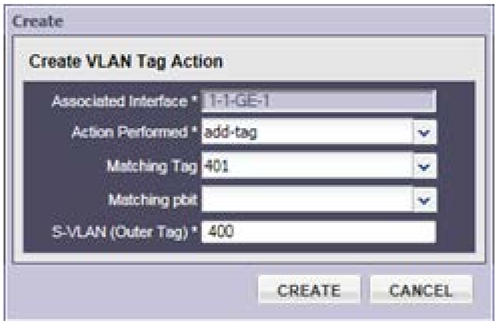

Create

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 an action to add a tag (VLAN 400) to all traffic based on the P-bit for the associated interface.

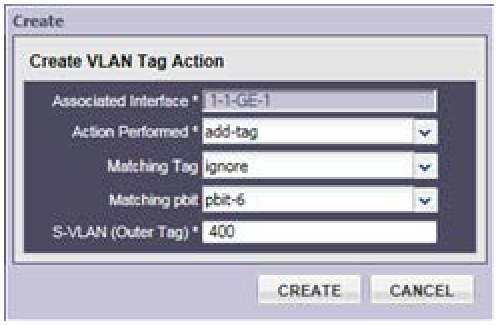

Create

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 an action to add a tag (VLAN 400) to all traffic based on the P-bit and outer tag on the associated interface.

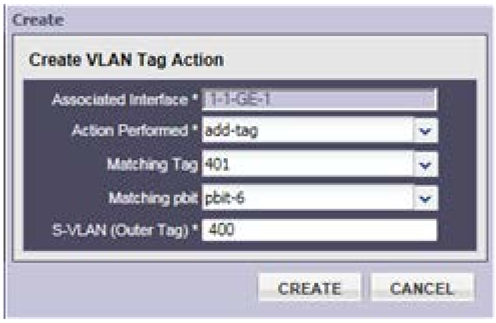

Create

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

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 E-Series 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 E-Series operation.)	2-4093
C-VLAN (Inner Tag)	New VLAN ID. (VLAN IDs 1002-1005 are reserved for E-Series 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

To add VLAN tag actions to an E-Series Ethernet interface

1. On the Navigation Tree, click **Interfaces**.
2. Double-click the specific interface where you want to add a tag action.

   Note: The tag actions only apply to interfaces assigned as edge or access link.

3. Click the **Tag Actions** tab.
4. From the menu, click **Create**.
5. Reference the table above to configure the parameters.
6. Click **Create**.

For CLI:

- To add a tag to all packets on an interface:

```
create tag-action add-tag <vlan ID> interface <interface name>
```

- To add a tag only to the packets that match specific criteria on an interface, append one of the following to the command above:

```
match-pbit <pbit value>
match-tag <vlan ID>
match-tag <vlan ID> match-pbit <pbit value>
```

Note: Using match-pbit criterion without using match-tag criterion will match only 802.1P packets (reserved VLAN 0).

- To add an outer tag and an inner tag to untagged traffic on an interface:

```
create tag-action add-2-tags outer <vlan id> inner <vlan id> interface
<interface name> match-untagged
```

- To change the outermost tag on packets that match a specific VLAN ID.

  ```
  create tag-action change-tag <vlan ID> interface <interface name> match-tag
  <match vlan ID>
  ```

- To change the tag and add an outer tag on packets that match either a specific VLAN ID or priority-tagged (P-bit) frames:

  ```
  create tag-action add-and-change outer <vlan id> inner <vlan id> interface
  <interface name> match-tag <vlan id>

  create tag-action add-and-change outer <vlan id> inner <vlan id> interface
  <interface name> match-prio-tag p-bit any
  ```

Creating Match Lists and Rules

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.

This topic describes how to create a match list and then add an ordered collection of matching rules to associate with a service tag action. The match list defines how the VDSL2 card classifies subscriber traffic to determine the service in which it belongs. 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 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
  - The Ethernet type to distinguish video and High-Speed Internet (HSI) untagged traffic
  - All untagged traffic by using the system default match list "all-untagged".
  - PVC VPI and VCI port numbers.
    - The E7-2 VDSL2 line cards, E5-48C, E5-48, and E3-48C products support up to six Services per DSL port that may be configured to support Voice, Data, Video, TR-69 Management, Transparent LAN Services, or other VLAN assignments. Support for multiple services is used when migrating from an ATM based DSLAM to the E-Series solutions. If an E-series DSL product is placed in an existing network that uses multiple PVC assignments, individual PVCs may be mapped to specific VLANs that are applied to the port. For example, PVC assignments (e.g. 0/33, 8/33, 0/35) may be mapped to individual data services (“Data 1”, “Data 2”, “Data 3”) that are applied to the port. The E7-2, E5-48C, E5-48, and E3-48C support one VC assignment per service.
    - TR-69 management channel (if present) is handled in-band over data VC.

    - The DSL port MUST be set to PTM-Override=ATM in order for Multi-VC support to work
    - The System will not allow a tag action that specifies a PVC match list to be added if PTM-Override=Auto
    - If PTM-Override=Auto, only the provisioned Fallback PVC will work.
- **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

Configuration guidelines

- If a residential gateway (RG) at the subscriber's premises classifies traffic into VLANs, then for each service type, create a match list with a tagged rule that matches the service VLAN.
- In E7 R2.1 and earlier, downstream priority-tagged traffic that does not meet any match criteria will be passed through unmetered. In E7 R2.2 and above, this is also true for VDSL2 and GPON (if the PON CoS setting allows it).

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 = ignore
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 Default = p-bit-none
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:x) • none (no MAC mask) • oui (mask OUI fields (FF:FF:FF:FF:FF:FF)	six hexadecimal digits in the range 0-FF, optionally separated by colons, none ‡, oui

Parameter	Description	Valid Options
Ethertype	Ethernet type to match (untagged match rules only). • any = default • pppoe = 0x8864 (for HSI) • arp = 0x0806 (for video) • ipv4 = 0x0800 (for video) • ipv6 = 0x86DD (for video)	any ‡, pppoe, arp, ipv4, ipv6
VPI	Specify the VPI for the PVC (untagged match rules only) used by the subscriber's modem (ADSL modems only).	0-255 0 ‡
VCI	Specify the VCI for the PVC (untagged match rules only) used by the subscriber's modem (ADSL modems only).	0-255 0 ‡

*Required field

To create a service match list

1. Access the profile page:
   - From CMS:
     - On the Navigation Tree, click **CMS**.
     - In the Work Area, click **Profile > E5-48/E3-48C/E7/ONT > Service** > **Tagging** > **Match Lists**.
   - Locally on the E-Series:
     - On the Navigation Tree, click the E-Series unit.
     - In the Work Area, click **Profiles** > **Service** > **Tagging** > **Match Lists** > **Profiles**.
2. In the menu, click **Create**.
3. In the Name box of the Create Match List dialog box, enter the name of the service match list that you are creating.
4. Click **Create**.

For CLI:

- `create svc-match-list <list name>`
- `delete svc-match-list <list name>`
- `show svc-match-list [list name]`

To add a rule to a service match list

1. If you have not already done so, create a service match list. See "To create an Service match list," above.
2. Access the profile page:

- From CMS:
  - On the Navigation Tree, click **CMS**.
  - In the Work Area, click **Profile** > **E7/E5-48/E3-48C > Service** > **Tagging** > **Match Lists**.
- Locally on E7:
  - On the Navigation Tree, click **E7**.
  - In the Work Area, click **Profiles** > **Service** > **Tagging** > **Match Lists** > **Profiles**.

The table of the previously created match lists appears.

3. Double-click the row that shows the match list in which you want to add a rule.
4. In the menu, click **Create** and select the type of rule:
   - **Tagged Match Rule**
   - **Untagged Match Rule**
5. If you chose to create a tagged match rule, do the following:

   a. In the Outer Tag list, select one of the following:
      - **Enter value** and then enter the VLAN ID of the outer tag that can be specified by name or by numeric VLAN ID, which ranges from 1-4095.
      - **untagged** indicates that only untagged traffic should be matched.
      - **ignore** indicates that the VLAN ID should not be examined.

   b. In the Outer P-bit list, select one of the following:
      - pbit value (range 0-7) that specifies the VLAN priority value to match.
      - pbit-none indicates the P-bit value is not considered. If this parameter is not specified, "pbit-none" is the default behavior.

   c. Click **Create**.
6. If you chose to create an untagged match rule, do the following:

   a. In the Source MAC list, do one of the following:
      - Enter the source MAC address (six hexadecimal digits in the range 0-FF, optionally separated by colons).
      - Leave the "ignore" to leave blank.

   b. In the Source MAC list, do one of the following:
      - Select Enter Source MAC Mask, and then enter the source MAC mask (six hexadecimal digits in the range 0-FF, optionally separated by colons).
      - Select oui to mask the OUI fields.
      - Select none to ignore.

c. In the Ethertype list, select the type of Ethernet frames that you want to match.

d. In the VPI and VCI boxes, enter values for the Virtual Path Identifier (VPI) and Virtual Circuit Identifier (VCI). E-series DSL nodes offer support for up to six services per port in ADSL2+ fallback applications.

e. Click **Create**.

7. Associate the service match list with a service tag action to use on E-Series ingress traffic packets.

For CLI:

- `add tagged-rule to-svc-match-list`
- `add untagged-rule to-svc-match-list`
- `remove tagged-rule <r-index> from-svc-match-list`
- `remove untagged-rule <r-index> from-svc-match-list`

Creating Service Tag Actions

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 E-Series 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. • NOT 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* (on page 96).

  - Only one service per port can use DSCP-Bit mapping function.
  - Works with Add-Tag and Add-2-Tags 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* (on page 98) or create an *IP Precedence to P-Bit Mask* (on page 100), if you do not want to use the system-default.

  b. Assign the map to the *xDSL port associated interface* (on page 182), 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 E-Series 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 table shows 802.1p to DiffServ DSCP 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 E-Series 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 E-Series 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 E-Series 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 E-Series 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 E-Series 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 E-Series 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 E-Series.
- 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 E-Series.
- 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.

- 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 (ony 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
  - 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.

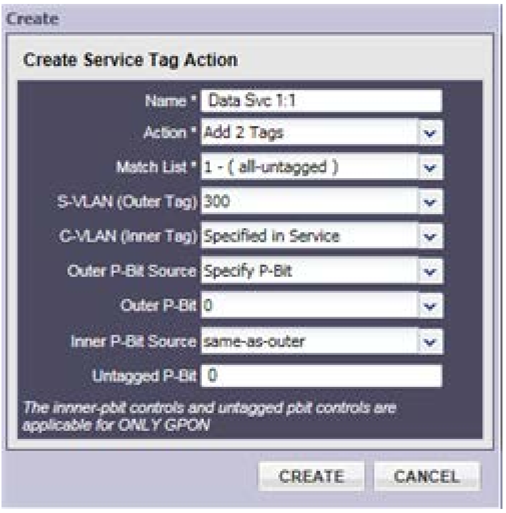

Create Service Tag Action

Name *	Data Svc 1:1
Action *	Add 2 Tags
Match List *	1 - (all-untagged)
S-VLAN (Outer Tag)	300
C-VLAN (Inner Tag)	Specified in Service
Outer P-Bit Source	Specify P-Bit
Outer P-Bit	0
Inner P-Bit Source	same-as-outer
Untagged P-Bit	0

The innner-pbit controls and untagged pbit controls are applicable for ONLY GPON

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.

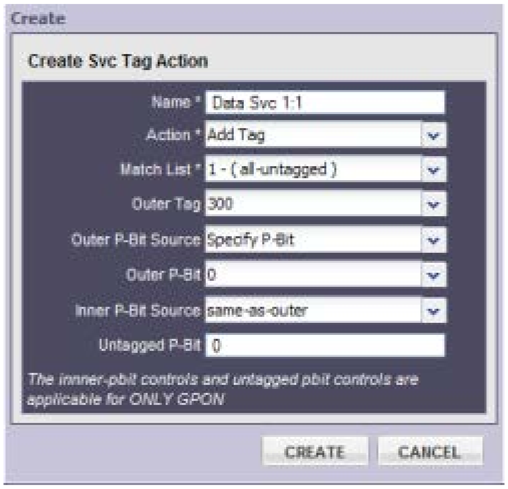

Create Svc Tag Action

Name *	Data Svc 1:1
Action *	Add Tag
Match List *	1 - (all-untagged)
Outer Tag	300
Outer P-Bit Source	Specify P-Bit
Outer P-Bit	0
Inner P-Bit Source	same-as-outer
Untagged P-Bit	0

The innner-pbit controls and untagged pbit controls are applicable for ONLY GPON

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

Parameter	Description	Valid Options
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 ‡, 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 E-Series honors the P-bit value of the incoming traffic and passes the existing P-bit value upstream.	0-7, copy 0 ‡
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
‡ Default

To create a service tag action

1. If you have not already done so, create a match list. See *Creating Match Lists and Rules* (on page 78).
2. On the Navigation Tree, double-click the unit.
3. Click **Profiles** > **Service** > **Tagging** > **Tag Actions** > **Profiles**.
4. In the menu, click **Create**.
5. In the Create Service Tag Action dialog box, do the following:
6. In the Name box, enter a descriptive name for the tag action.

7. In the Action list, select the action to perform if the match criteria is met.
8. In the Match Criteria List, select the service match list to associate match criteria with the service tag action.
9. In the Outer Tag list, select whether to use the tag value specified in the Ethernet service provisioning or to enter a tag value for use in the tag action.
10. In the Inner Tag list, select whether to use the Inner Tag, to use the tag value specified in the Ethernet service provisioning, or to enter a tag value for use in the tag action.

    Note: The Inner Tag parameter only applies if the selected action is "Add 2 Tags" or "Add and Change Tag."

11. In the P-Bit Source list, select where to derive the P-bit value for marking the outer tag for services provisioned on xDSL ports.
    - **Specify P-Bit** requires that you enter the P-bit value for marking the outer tag, or select "copy" to pass the existing P-bit value upstream.
    - **Map a layer-3 priority** references the system default DSCP and IP Precedence values, unless you specify another existing custom map from the xDSL interface parameter.

    Note: The upstream traffic on a service VLAN must have the same priority as is set for downstream traffic.

12. Click **Create**.
13. Associate the tag action to an xDSL port when provisioning a service.

For CLI:

```
create svc-tag-action <name> type <add-2-tags|add-and-change> outer
<VLAN-ID> inner <VLAN-ID> svc-match-list <l-name> use-p-bit|derive-
p-bit

create svc-tag-action <name> type <add-tag|change-tag> outer
<VLAN_ID|use-svc-vlan> svc-match-list <l-name> use-p-bit|derive-p-
bit

show svc-tag-action [name]
delete svc-tag-action <name>
```

Creating a Class Map and Rules

Policy maps are lists of QoS-related actions to perform on packets that match certain criteria. The matching criteria are specified by an associated classification map that lists rules to identify packets, such as VoIP traffic.

This topic describes how to create a class map and class rules that specify certain selection criteria, such as P-bit or VLAN values, to classify traffic from the network that is arriving at the E-Series. 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 E-Series 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 E-Series 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:

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:

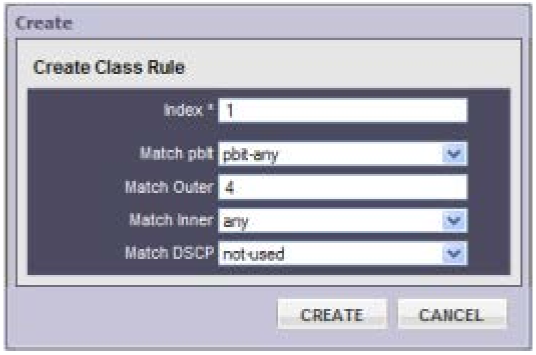

Index * 1
Match pbit pbit-any
Match Outer 4
Match Inner any
Match DSCP not-used

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 ‡
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 ‡
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 ‡ Enter Value
Match Inner	Inner VLAN ID to match. VLANs can be specified by name or by numeric VLAN ID (range 1-4093).	Ignore ‡, any, untagged, Enter Value
Match DSCP	DSCP value to match.	Not used ‡, 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
‡ Default

To create a class map

1. On the Navigation Tree, click **E-Series**.
2. Click **Policies** > **Class Map.**
3. In the menu, click **Create**.
4. In the Create Class Map dialog box, do the following:
   a. In the Name box, enter a descriptive name of the classification map that you are creating.
   b. In the Map Type list, select whether the packets must match any or all of the criteria listed in the class map.
5. Click **Create**.

For CLI:

- `create class-map <c-map name>`
- `create class-map <c-map name> match-type all`
- `create class-map <c-map name> match-type any`

To create a class rule

1. If you have not already done so, create a class map. See "To create a class map," above.
2. On the Navigation Tree, click the unit.
3. Click **Policies** > **Class Map**.

   The table of the previously created class maps appears.
4. Double-click the name of the class map in which you want to add a class rule.
5. Click **Create**.
6. Reference the table above to configure the parameters.
7. Click **Create**.
8. Associate the class map with a policy map that is associated to an E-Series port interface to use on ingress traffic packets. See *Creating a Policy Map and Policies* (on page 93).

For CLI:

- `add class-rule <rule index> to-map <c-map name> match-all`
- `add class-rule <rule index> to-map <c-map name> match-pbit <P-bit value>`
- `add class-rule <rule index> to-map <c-map name> match-tag <outer vlan ID> [match-pbit]`
- `add class-rule <rule index> to-map <c-map name> match-2-tags outer <outer vlan ID> inner <inner vlan ID> [match-pbit]`
- `add class-rule <rule index> to-map <c-map name> match-dscp <dscp value>`
- `add class-rule <rule index> to-map <c-map name> match-tag <outer vlan ID> [match-dscp]`

Creating a Policy Map and Policies

This topic describes how to create a policy map and policies that specify 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 *Creating a Policy Map for L3 Priority Mapping* (on page 102) for instructions on using a class map to match the layer-3 priority value of incoming 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.

See the *Calix E-Series Engineering and Planning Guide* 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 E-Series; the edge device that shapes traffic greatly reduces the buffer requirements in the infrastructure device. Also see *Creating Ethernet Port Class of Service* (on page 71) for information on maximum rate and burst size on a per-port basis.

Configuration Considerations

- The E-Series 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

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-it, only the policy rule with the lowest sequence number ID is carried out.	1-1500
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 = 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 E-Series Ethernet port interface are forwarded. If you select "Set a Rate Limit" you must also enter a limit value.	none ‡ 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 ‡, Set a Burst Size (4-16000 kilobytes)

*Required fields
‡ Default

To create a policy map

1. On the Navigation Tree, click the unit.
2. Click **Policies** > **Policy Map**.
3. In the menu, click **Create**.
4. In the Create Policy Map dialog box, enter a descriptive name for the policy map that you are creating.
5. Click **Create**.

For CLI:

```
create policy-map <p-map name>
```

To create a policy

1. If you have not already done so, create a policy map. See "To create a policy map," above.
2. On the Navigation Tree, click **E-Series**.
3. Click **Policies** > **Policy Map**.

   The table of the previously created policy maps appears.
4. Double-click the name of the policy map in which you want to add a policy.
5. In the menu, click **Create**.
6. Reference the table above to configure the parameters.
7. Click **Create**.
8. Associate the policy map with an E-Series Ethernet port interface to use on ingress traffic packets. See *Configuring an Ethernet or LAG Interface* (on page 33).

For CLI:

```
add policy <policy index> to-map <p-map name> class-map <c-map name> [set-
pbit|rate-limit|max-burst-size]
```

Mapping Layer 3 Priority Values to P-Bits

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.
  - Guidelines:
    - Only one service per port can use DSCP-Bit mapping function.
    - Works with Add-Tag and Add-2-Tags tag actions.
    - Either DSCP or IP Precedence can be used for a service, not both.
  - Use the following provisioning sequence:

    a. Create DSCP to P-Bit map or Create IP Precedence to P-Bit Mask.
    Creating and Modifying a DSCP Map for Ingress Traffic (on page 98)
    Creating and Modifying an IP Precedence Map (on page 100)

    b. Assign DSCP to P-Bit map to xDSL port associated interface.
    Configuring an xDSL Port Associated Interface (on page 182)

    c. Create tag action with an Outer P-Bit source of “Map a layer-3 Priority.”
    Creating Service-Tag Actions (on page 82)

d. Create service using the defined tag action.
Creating Data Services (on page 183)
Creating IP Video Services (on page 196)

Note: The upstream traffic on a service VLAN must have the same priority value as is set for the downstream traffic.

- **Traffic coming from the network going downstream 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).

  See *Creating a Policy Map for L3 Priority Mapping* (on page 102).

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
5	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.

Creating and Modifying a DSCP Map for Ingress Traffic

This topic shows you how to create and modify an DSCP Map profile that is applied to xDSL port interfaces to enable and/or specify a mapping of the incoming frame DSCP values to IEEE 802.1p priority bits.

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 E-Series 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 E-Series 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

Note: Calix recommends NOT assigning P-bit value 7 to any DSCP value.

To create a DSCP map

1. In the Navigation Tree, click the unit.
2. In the Work Area, click **Profiles** > **DSCP** > **Create**.
3. In the Create DSCP Map dialog, do the following:
   a. In the Name box, enter a name that describes the use of this custom profile.
   b. In each of the designated DSCP value lists, select the correlating P-bit value, following site requirements and company policies and procedures to assign 802.1p priority bit values to the values listed alongside each list box.
4. Click **Create**.

For CLI:

```
create dscp-map <name>
[default|cs0|af11|af12|af13|af21|af22|af23|cs3|af31|af41|ef|cs6|cs7]
```

To modify a DSCP map

1. In the Navigation Tree, click the unit.
2. In the Work Area, click **Profiles** > **DSCP**.
3. Double-click the name of the DSCP map that you want to modify.
4. In each of the designated DSCP value lists, select the correlating P-bit value, following site requirements and company policies and procedures to assign 802.1p priority bit values to the values listed alongside each list box.
5. Click **Apply** to save the changes to the map.

For CLI:

```
set dscp-map <name>
[default|cs0|af11|af12|af13|af21|af22|af23|cs3|af31|af41|ef|cs6|cs7]
show dscp-map
show dscp-map <name>
```

Creating and Modifying an IP Precedence Map

This topic shows you how to 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	Decimal	802.1p Priority-Bit
000 xxx	Routine/Best Effort	0	0
001 xxx	Priority	1	1
010 xxx	Immediate	2	2
011 xxx	Flash/Voice/Video	3	3
100 xxx	Flash Override	4	4
101 xxx	Critical/Voice RTP	5	5
110 xxx	Internet	6	0
111 xxx	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.

To create an IP precedence profile

1. In the Navigation Tree, click the unit.
2. In the Work Area, click **Profiles** > **IP Pred** > **Create**.
3. In the Create IP Precedence Map dialog, do the following:
    a. In the Name box, enter a name that describes the use of this custom profile.
    b. In each of the designated IP Precedence value lists, select the correlating P-bit value, following site requirements and company policies and procedures to assign 802.1p priority bit values to the values listed alongside each list box.
4. Click **Create**.

For CLI:

```
create ip-precedence-map <name> [ip-precedence-0|ip-precedence-1|ip-
precedence-2|ip-precedence-3|ip-precedence-4|ip-precedence-5|ip-
precedence-6|ip-precedence-7]
```

To modify IP precedence mapping

1. In the Navigation Tree, click the unit.
2. In the **Work Area**, click **Profiles** > **IP Pred**.
3. Double-click the name of the IP precedence map that you want to modify.
4. In each of the designated IP Precedence value lists, select the correlating P-bit value, following site requirements and company policies and procedures to assign 802.1p priority bit values to the values listed alongside each list box.
5. Click **Apply** to save the changes to the map.

For CLI:

- `set ip-precedence-map <name> [ip-precedence-0|ip-precedence-1|ip-precedence-2|ip-precedence-3|ip-precedence-4|ip-precedence-5|ip-precedence-6|ip-precedence-7]`
- `show ip-precedence-map`
- `show ip-precedence-map <name>`

Creating a Policy Map for L3 Priority Mapping

For ingress traffic on Trunk and Edge Ethernet (GE/10GE) ports, you can use a class map to match the layer-3 priority value of incoming frames, and then use a policy map to assign a corresponding P-bit (priority) value into the classified traffic.

This topic shows example outlines of single, specific DSCP value and multiple values used to define mapping onto P-bits. This assumes that the upstream routers or aggregation switch are not P-bit aware and do not manage traffic based on P-bit marking. Instead, DSCP is used for this purpose.

For detailed instructions on how to create policy maps, see the following topics:

- *Creating a Class Map and Rules* (on page 90)
- *Creating a Policy Map and Policies* (on page 93)

Note: The upstream traffic on a service VLAN must have the same priority value as is set for downstream traffic.

For further information on layer-3 priority mapping, see:

- Mapping Layer-3 Priority Values to P-Bits

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.

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.

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.

Example 1: Single DSCP value mapping to CoS 5

1. Create class map (CoS 5).
    - Create class rule:
        - Match pbit pbit-any
        - Match Outer 100
        - Match Inner ignore
        - Match DSCP 40 (equates to CS5)
2. Create policy map (CoS 5)
    - Create Policy: class map (CoS 5), Out pbit 5
3. Associate the policy map to the Ethernet port interface.

Example 2: Multiple DSCP value mapping

1. Create class maps and add rules.

    a. Create class map (mix1)
    - create rule 1: match tag 101, match-dscp af11

    b. Create class map (mix2)
    - create rule 1: match tag 102, match-dscp cs2

    c. Create class map (mix3)
    - create rule 1: match-dscp af33

    d. Create class map (mix4)
    - create rule 1: match tag 201, match-dscp cs4

    e. Create class map (mix5)
    - create rule 1: match tag 202, match-dscp ef

    f. Create class map (mix6)
    - create rule 1: match tag 203, match-dscp cs6

    g. Create class map (mix7)
    - create rule 1: match tag 207, match-dscp 56

h. Create class map (op4)

   - create rule 1: match p-bit 4, match-dscp af33

2. Create policy maps and add policies.

a. Create policy map (mix)

   - create policy 1, class map (mix1), set P-bit 1
   - create policy 2, class map (mix2), set P-bit 2, rate-limit 15, max burst size 4
   - create policy 3, class map (mix3), set P-bit 3
   - create policy 4, class map (mix4), set P-bit 4
   - create policy 5, class map (mix5), set P-bit 5
   - create policy 6, class map (mix6), set P-bit 6
   - create policy 7, class map (mix7), set P-bit 7, rate-limit 10, max-burst size 4

b. Create policy map (0p4)

   - create policy 1, class map (0p4), set P-bit 3

3. Add policy maps to Ethernet interfaces.

- 1/x3 policy map (mix)
- 2/x3 policy map (mix)
- 2/g4 policy map (0p4)

Related topics

- *Creating a Class Map and Rules* (on page 90)
- *Creating a Policy Map and Policies* (on page 93)

Creating Data and Video Service Profiles

This section describes how to create profiles and objects that will be associated with data and video services provisioned on a VDSL2 card xDSL port.

Creation of following profiles and objects are described:

- Ethernet Bandwidth Profile for xDSL Services
- Multicast Profile
- DSL Vectoring Group
- DSL Port Template
- Ethernet Security Profile
- PPPoE Profile
- xDSL Bonded Interface

Next steps:

After completing the creation of profiles for data and video services, see the following sections to continue configuring data and video services:

- *Configuring Data Services* (on page 177)
- *Configuring IP Video Services* (on page 189)

Creating an 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.

Traffic Class Type	E7 P-bits (default)	Service Category	E7 BW Profile
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 simplist 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:

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.	0-2048kpbs, 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.	0-2048kpbs, 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.	0-2048kpbs, 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 16000kbps) • 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 16000kbps) • Use "m" suffix for Mbps (0 to 16 mbps) • Use whole number increments	4-16000 Kbps

*Required field

To create a bandwidth profile

1. On the Navigation Tree, click the unit.
2. Click **Profiles** > **Service** > **Ethernet Bandwidth** > **Profiles.**
3. On the menu, click **Create**.
4. In the Create Bandwidth Profile dialog box, do the following:
   a. In the Name box, enter a descriptive name for the profile.

   b. In the Committed Rate for Upstream box, enter a value that specifies the minimum rate to allow traffic to flow upstream.

   c. In the Peak Rate for Upstream box, enter a value that specifies the un-guaranteed maximum bandwidth for upstream traffic.

   d. In the Peak Rate for Downstream box, enter a value that specifies the un-guaranteed maximum bandwidth for downstream traffic.

   e. In the Downstream Peak Burst Size box, enter a value that specifies the un-guaranteed maximum bandwidth for downstream peak burst size.

   f. In the Upstream Committed Peak Burst Size box, enter a value that specifies the minimum bandwidth for upstream peak burst size.

   g. In the Upstream Peak Burst Size box, enter a value that specifies the un-guaranteed maximum bandwidth for upstream peak burst size.

   h. Click **Create**.
5. Associate the bandwidth profile to a video or data service being created on an xDSL port. See *Configuring IP Video Services* (on page 189) and *Configuring Data Services* (on page 177).

For CLI:

```
create bw-profile <p-name> [upstream-cir|upstream-pir|downstream-
pir|upstream-cbs|upstream-pbs|downstream-pbs]
```

Creating a Multicast Profile

This topic describes how to create a multicast profile that enables provisioning of the IPTV service limits that can be applied when adding an Ethernet service to an xDSL port. A multicast profile is defined by the parameters listed below and references two previously defined profiles:

- **Multicast Address Map**

  Identifies the optional global allowable multicast IP ranges. See Creating a Multicast Address Map and Ranges.

- **Multicast VLAN Registration (MVR) Profile**

  Identifies the optional MVR address ranges associated with specified multicast VLANs. See *Creating a MVR profile and MVR VLAN Addresses* (on page 111).

Parameters

You can provision the following parameters for multicast profiles:

Parameter	Description	Valid Options
Name*	Name of the multicast profile.	32 character string
Max Streams*	Maximum number of multicast streams allowed at the subscriber port. This 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.	1-128
Multicast Maps	Identifies the map of optional global multicast IP addresses that a subscriber service can join when this multicast profile is used. The multicast map is an optional parameter of the multicast profile. With no multicast address map identified, multicast group destination addresses are unconstrained.	Any established multicast map
MVR Profile	Identifies the optional MVR profile to be used by this multicast profile. **Note:** An MVR profile cannot be paired with more than one multicast map.	Any established MVR profile
General Query Interval	The value in seconds that matches the upstream router's general query interval. **Note:** General Query Interval is NOT applicable for VDSL services. In the absence of a channel join, the ONT keeps the channel active for the query interval before cutting it off. The assumption is the router sends a general query every interval which causes the set-top box to join the channel which also keeps the channel alive.	10-3600 240 ‡

Convert to Unicast	Whether to convert multicast packets to unicast before sending them out the subscriber Ethernet port.
	The Convert to Unicast parameter should ONLY be enabled if switches are deployed on the subscriber side that cannot manage heavy multicast traffic.
	• **Note:** Convert to Unicast feature is only applicable to the Calix GX ONTs.
		selected=enabled
		unselected=disabled

* Required
‡ Default

To create a multicast profile

1. If you have not already done so, create an MVR Profile and Multicast Address Map, if necessary.
2. On the Navigation Tree, click the unit.
3. Click **Profiles** > **Service** > **Multicast** > **Multicast** > **Profiles**.
4. In the menu, click **Create** to open the Create Multicast Profile dialog box.
5. Reference the table above to configure the parameters.
6. Click **Create**.
7. Associate the multicast profile when creating video service on an xDSL port.

For CLI:

- `create mcast-profile <p-name> [name|max-strms|query-interval|convert-mcast|mcast-map|mvr-profile]`
- `delete mcast-profile <p-name>`
- `set mcast-profile <p-name> [max-strms|query-interval|convert-mcast]`
- `show mcast-profile [p-name]`
- `show mcast [ip *|on-interface *|on-interface * ip *|on-dsl-port *|on-dsl-port * ip *|on-gpon-port *|on-gpon-port * ip *|on-vlan *|on-vlan * ip *]`

Creating an 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).

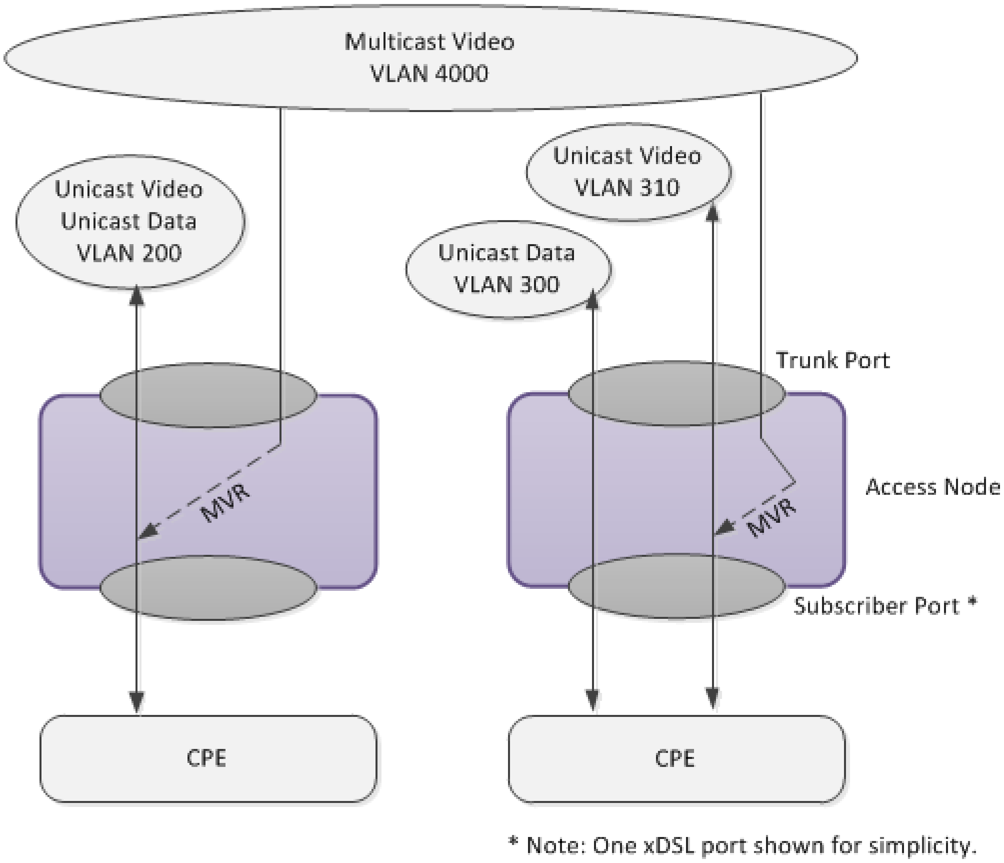

* Note: One xDSL port shown for simplicity.

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 E-Series 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).

Before starting

Before you start the procedure to create a MVR profile and include an MVR VLAN address range, the VLAN must exist on the system.

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".
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

To create an MVR profile and MVR VLAN addresses

1. Access the profile page:
   - From CMS:
     - On the Navigation Tree, click **CMS**.
     - In the Work Area, click **Profile** > **E3-48C/E5-48/E7/ONT** > **Profile** > **Service** > **Multicast** > **MVR**.
   - Locally on the node:
     - On the Navigation Tree, click the unit.
     - In the Work Area, click **Profiles** > **Service** > **Multicast** > **MVR** > **Profiles**.
2. In the menu, click **Create** to open the Create MVR Profile dialog box.
3. In the Name box, enter a descriptive name for the profile, and then click **Create**.
4. In the list of existing MVR profiles, double-click the profile name to select it.
5. In the menu, click **Create** to open the Create MVR VLAN dialog box.
6. Reference the table above to configure the parameters.
7. Click **Create** to add the MVR VLAN to the MVR profile.
8. Associate the MVR profile when creating a multicast profile.

For CLI:

- `create mvr-profile <p-name>`
- `add vlan <vlan-id> to-mvr-profile <p-name> [mcast-range-1|mcast-range-2|mcast-range-3|mcast-range-4]`
- `delete mvr-profile <p-name>`

- `remove vlan <vlan-id> from-mvr-profile <p-name>`
- `set mvr-profile <p-name> name <new-name>`
- `set mvr-profile <p-name> vlan <vlan-id> [mcast-range-1|mcast-range-2|mcast-range-3|mcast-range-4]`
- `show mvr-profile [p-name]`

Creating Multicast Address Map and Ranges

This topic describes how to create a multicast address map and IP address ranges that are assigned to the map. The map is referenced by a multicast profile that is applied when adding a service to an xDSL port. See *Creating a Multicast Profile* (on page 110).

Configuration guidelines

- The multicast address map is an optional parameter of the multicast profile.
- With no multicast address map identified, multicast group destination addresses are unconstrained.
- Each multicast address map can contain up to eight distinct multicast IP address ranges, not overlapping.
- 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 address map is ignored.

Parameters

You can provision the following parameters for multicast maps and ranges:

Parameter	Description	Valid Options
Name*	Name of the multicast map.	32 character string
ID*	Index value designating the multicast range in the multicast map. You can leave the value, or select another (1-8).	1-8
Start IP Address*	First address in the multicast range.	Any multicast address, lower than the end IP address (224.0.0.0 - 239.255.255.255)
End IP Address*	Last address in the multicast range.	Any multicast address, higher than the end IP address (224.0.0.0 - 239.255.255.255)

*Required fields

To create a multicast address map and range

1. Access the profile page:
   - From CMS:
     - On the Navigation Tree, click **CMS**.
     - In the Work Area, click **Profile** > **E3-48C/E5-48/E7/ONT > Profile** > **Service** > **Multicast** > **Multicast Maps**.

- Locally on the E-Series:
  - On the Navigation Tree, click the unit.
  - In the Work Area, click **Profiles** > **Service** > **Multicast** > **Multicast Maps** > **Profiles**.

2. In the menu, click **Create** to open the Create Multicast Maps dialog box.
3. In the Name box, enter a descriptive name for the map, and then click **Create**.
4. In the list of existing multicast maps, click the map name to select it.
5. In the menu, click **Create** to open the Create Multicast Address Range dialog box.
6. Reference the table above to configure the parameters.
7. Click **Create** to add the multicast address range to the multicast map.
8. Repeat Steps 5 through 7 to include another multicast range in the multicast map.
9. Associate the multicast map when creating a multicast profile.

For CLI:

- `create mcast-map <m-name>`
- `add range to-mcast-map <m-name> mcast <m-range>`
- `delete mcast-map <index>`
- `remove range <index> from-mcast-map <m-name>`
- `set mcast-map <old-name> name <new-name>`
- `set mcast-map <m-name> range <index> mcast <m-range>`
- `show mcast-map [m-name]`

Creating a Multicast Whitelist Map and Ranges

This topic describes how to create a multicast white list and IP address ranges that are assigned to the list. You can use multicast white lists to create complex channel line-ups on a per subscriber basis without the use of Middleware. The map is then applied when adding a video service to an ONT or xDSL Ethernet port.

Define IPTV packages by specifying multicast channel ranges, organizing diverse channel line-ups into more manageable groups, known as maps:

- Each system supports:
  - up to 128 multicast white list maps where each supports up to 32 channel sequences
- Each video service supports:
  - up to 16 multicast white list maps per subscriber service

For example, there can be a Multicast White list that defines all the local television channels at standard definition and another Multicast White list for local television channels at high definition. Other examples of Multicast White list uses is for defining channel bundles like HBO, Cinemax, ESPN, one for each individual bundle offered.

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 exiting 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

Parameters

You can provision the following parameters for multicast white lists and ranges:

Parameter	Description	Valid Options
Name*	Name of the multicast white list.	32 character string
ID*	Index value designating the multicast range in the multicast white list map. You can leave the value, or select another (1-32).	1-32
Start IP Address*	First address in the multicast range.	Any multicast address, lower than the end IP address (224.0.0.0 - 239.255.255.255)
End IP Address*	Last address in the multicast range. May be 0 (for single address), if non-0 must be in Mcast range and greater than start IP address value.	Any multicast address, higher than the end IP address (>239.255.255.255)

*Required fields

To create a multicast white list and address ranges

1. Access the profile page:
   - From CMS:
     - On the Navigation Tree, click **CMS**.
     - In the Work Area, click **Profile** > **E3-48C/E5-48/E7/ONT > Profile** > **Service** > **Multicast White List**.
   - Locally on the E-Series:
     - On the Navigation Tree, click the unit.
     - In the Work Area, click **Profiles** > **Service** > **Multicast White List**.
2. In the menu, click **Create** to open the Create Multicast White List dialog box.
3. In the Name box, enter a descriptive name for the list, and then click **Create**.
4. In the list of existing multicast white lists, double-click the name to select it.
5. In the menu, click **Create** to open the Create Multicast White List Range dialog box.
6. Reference the table above to configure the parameters.
7. Click **Create** to add the multicast address range to the multicast white list.
8. Repeat Steps 5 through 7 to include another multicast range in the multicast white list.
9. Associate the multicast white list when creating a video service.

   See Configuring IP video Services for instructions.

For CLI:

Managing the map object

- **`create mcast-white-list <list-name>`**
- **`set mcast-white-list <list-name> name <new list-name>`**
- **`show mcast-white-list <list-name>`**
- **`delete mcast-white-list <list-name>`**

Managing ranges within the map object

- **`add range to-mcast-white-list <list-name> mcast <a.b.c.d[-w.x.y.z]>`**
- **`set mcast-white-list <list-name> range <index> mcast <a.b.c.d[-w.x.y.z]>`**
- **`remove range <index> from-mcast-white-list <list-name>`**

Managing Associations of maps to eth-srv objects

- **`add mcast-white-list <list-name> to-interface <vdsl interface> eth-svc <service>`**
- **`add mcast-white-list <list-name>  to-ont-port <ont port> eth-svc <service>`**
- **`add mcast-white-list <list-name> to-dsl-bond-interface <vdsl interface> eth-svc <service>`**
- **`remove mcast-white-list <list-name> from-interface <vdsl-interface> eth-svc <service>`**
- **`remove mcast-white-list <list-name> from-ont-port <ont port> eth-svc <service>`**
- **`remove mcast-white-list <list-name> from-dsl-bond-interface <vdsl-interface> eth-svc <service>`**
- **`show interface <vdsl-interface> eth-svc <service> detail`**
- **`show ont-port <ont-port> eth-svc <service> detail`**

Creating an IGMP Profile

This topic describes how to create a profile for VLAN association that sets configuration attributes of the Internet Group Management Protocol (IGMP) snoop used to establish membership in a multicast video services group. 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 mrounter 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 = 130sec). 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 = 130sec). 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.	Y (enable) N (disable) ‡
Robustness	Number of IGMP group-specific queries sent per query interval when proxy is configured.	1-10 1 ‡
Last Member Query Count	Number of group-specific queries sent when a Leave is received.	1-8 2 ‡
Last Member Query Interval	Time to wait in milliseconds for a response to a group-specific query.	100-5000 1000 ‡
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-and-dynamic ‡ static-only
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.	Y (enable) ‡ N (disable)

Parameter	Description	Valid Options
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 ‡
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 ‡
Startup Query Interval	Time interval between general queries during startup in seconds.	2-250 15 ‡
Startup Query Count	Time to wait for responses to general queries during startup.	1-10 2 ‡
Host Purge Time	Purge time in seconds for host ports.	130-3600 260 ‡
Router Port Purge Time	Purge time in seconds for router ports.	60-600 260 ‡
Proxy VLAN IGMP Version	The VLAN IGMP operating version.	v2 ‡, 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 ‡
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 ‡

*Required field
‡Default

To create an IGMP profile

1. On the Navigation Tree, click **E-Series**.
2. Click **Profiles** > **IGMP** > **Create**.
3. Reference the table above to configure the parameters.
4. Click **Create**.
5. Associate the IGMP profile to a specific VLAN, if the VLAN IGMP mode is set to Proxy.

For CLI:

```
create igmp-profile <p-name> [name|immediate-leave|robustness|last-memb-
query-count|last-memb-query-intrvl|router-learning-mode|router-solicit-top-
chg|version|query-interval|query-resp-interval|startup-query-
interval|startup-query-count|proxy-ip|host-port-purge-time|router_port-
purge-time|source-specific-mcast-range]
```

Configuring VDSL Vectoring

E-Series supports unit level vectoring on the VDSL2 lines. Vectoring eliminates cross talk between VDSL2 lines and thus recovers 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.

To achieve cross talk 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 can't share binder groups with non-vectored VDSL2 lines or the benefits of cross talk noise recovery are lost.

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.

Configuration guidelines

- The E-Series support board/unit level vectoring, where all ports on a card work together as a vector group. Using vectoring on some lines but not all causes un-vectored lines to interfere with vectored lines.
- Both single and bonded VDSL2 ports are supported.
- 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.
- The vector group is configured on the system and individual lines are added to the vector group. Vectoring is disabled by default.
- VDSL vectoring is not supported on E7 modular chassis systems.
- Since there is only one vectoring group per VDSL2 card or VDSL2 unit, it is always "1".
  - On an E7-2, the syntax is "<card>/1", e.g. "1/1" or "2/1".
  - On an E3-48C, E5-48, and E5-48C the syntax is just "1".

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-vectored)	Low

Planning Considerations when Enabling Vectoring

When planning a unit level 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 section titled "Unit Level Vectoring" in the *Calix E7 Platform Release Product Planning Guide* for more information about planning considerations when enabling vectoring.

Creating a VDSL2 Vectoring Group

This topic describes how to create a DSL Vector Group. Only one vector group may be configured per VDSL2 line card or VDSL2 unit.

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, 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.

Parameters

You can provision the PSD mask parameters for a DSL Vector Group:

Parameter	Description
PSD Mask*	Power spectral density mask. Valid values: • a-nus0 (VDSL2, Annex A, POTS compatibility, do not use band US0) • a-eu-32 ‡, a-eu-36, a-eu-40, a-eu-44, a-eu-48, a-eu-52, a-eu-56, a-eu-60, a-eu-64, a-eu-128, (VDSL2, Annex A, POTS compatibility, end US0 on subcarrier specified) • a-adlu-32, a-adlu-36, a-adlu-40, a-adlu-44, a-adlu-48, a-adlu-52, a-adlu-56, a-adlu-60, a-adlu-64, a-adlu-128, (VDSL2, Annex A, All Digital, end US0 on subcarrier specified) • b8-1, b8-4 (VDSL2, Annex B, 12 MHz, US0 as in ADSL2+ annex A) • b8-2, b8-6 (VDSL2, Annex B, 12 MHz, US0 as in ADSL2+ annex B) • b8-3, b8-7 (VDSL2, Annex B, 12 MHz, does not use US0) • b8-5 (VDSL2, Annex B, 12 MHz, US0 as in ADSL2/2+ annex M) • b8-6 (VDSL2, Annex B, 12 MHz, US0 as in ADSL2+ annex B) • b8-8, b8-9, b8-10 (VDSL2, Annex B, 17 MHz, does not use US0) • b8-11 (VDSL2, Annex B, 17 MHz, US0 as in ADSL2+ annex A) • b8-12 (VDSL2, Annex B, 17 MHz, US0 as in ADSL2+ annex B) • b7-1 (VDSL2, Annex B, 7 MHz, US0 as in ADSL2+ annex A) • b7-2, b7-4 (VDSL2, Annex B, 8.8 MHz, US0 as in ADSL2/2+ annex M) • b7-3 (VDSL2, Annex B, 12 MHz, US0 as in ADSL2/2+ annex M) • b7-5 (VDSL2, Annex B, 12 MHz, US0 as in ADSL2+ annex A) • b7-6 (VDSL2, Annex B, 12 MHz, US0 as in ADSL2/2+ annex M) • b7-7 (VDSL2, Annex B, 17.6 MHz, does not use US0) • b7-9 (VDSL2, Annex B, 17.6 MHz, US0 as in ADSL2+ annex A) • c-138-co (VDSL2, Annex C, 12 MHz, DS1 breaks at 138 kHz, type: co) • c-276-co (VDSL2, Annex C, 12 MHz, DS1 breaks at 276 kHz, type: co) • c-vdsl1-qam-compatible (Calix-specific mask. Allows VDSL2 to coexist in the same binder as VDSL1 QAM.)

*Required field
‡ Default

To create a DSL Vectoring Group

1. On the Navigation Tree, click the E-Series unit.
2. In the Work Area, click **DSL Vector Group**.
3. DSL Vector Group 1 has been created by default.
4. Double click the DSL Vector Group box to modify the PSD mask for the group from the default value of a-eu-32.
5. Click **Apply** to save the DSL vector group PSD setting.
6. Add all of the individual lines in the VDSL2 card or unit to the vector group from the Power Spectral Density (PSD) settings on a xDSL port or template. In the Work area, click **Port** > **Provisioning** > **PSD**.

7. Scroll down to the bottom of the provisioning form, and then select the Vectoring Group from the list. Since there is only one vectoring group per card or system, it is always “1”.

   - On an E7-2, the syntax is “<card>/1”, for example “1/1” or “2/1”.
   - On an E3-48C, E5-48, and E5-48C the syntax is just “1”.

   All VDSL2 ports in a card or VDSL2 unit work together as a vector group. Using vectoring on some lines but not all causes un-vectored lines to interfere with vectored lines.

   Both single and bonded VDSL2 ports are supported.

8. Click **Apply**. Vectoring is activated when the vectoring group is assigned to the port.
9. To verify the port vectoring status, do the following:

   a. In the Navigation Tree, click the xDSL port of interest.

   b. In the Work Area, click **Port** > **Status** > **Status**.

   c. View the following fields:

      - Actual Vectoring Mode = G.993.5
      - Line Vectoring State = steady
      - Operation Status fields = Showtime

For CLI:

- `set dsl-vectoring-group <number> [psd-mask <mask>]`
- `show dsl-vectoring-group`
- `set dsl-port <port-id> psd ds-vectoring [enabled|disabled]`
- `show dsl-port <port-id> [psd|status]`

Creating and Applying a DSL Port Template

This topic shows you how to create a DSL port template that is used to determine the modem train rate and the operating characteristics of the xDSL ports. You then apply a DSL template to an xDSL port, applying the template attributes to the DSL port. The parameter values from the applied template automatically display as the port settings.

You can then adjust the values supplied by the template on a per-parameter basis, as needed by accessing the parameter settings in the following levels:

- Basic
- Advanced
- Power Spectral Density (PSD)

Note: See *Calix xDSL Best Practices* for a description of the physical layer factors that may directly influence the quality of data and video services delivery, and recommendations for the best practices necessary to achieve optimal results.

Applying multiple DSL templates to a single port

Each DSL template may specify some or all of the xDSL port parameters. When a DSL template is applied to an xDSL port, the parameter values that are specified in that template are copied into the port object.

If another template is subsequently applied to the same xDSL port, the parameter values specified in that template are copied into the port object.

- If the same parameters are specified in both templates, the values defined in the last-applied template are applied to the xDSL port.
- If parameters specified in the first-applied template are not specified in the last-applied template, the values are retained from the first-applied template for those parameters.

To create a DSL port template

1. On the Navigation Tree, click **E-Series**.
2. Click **Template** > **DSL**.
3. In the menu, click **Create**.
4. In the Create DSL Port Template dialog box, enter a descriptive name for the template (40 characters).
5. Click Create.

For CLI:

```
create dsl-template <name>
```

To apply a DSL template to a port

1. On the Navigation Tree, click the xDSL port on which you want to apply the DSL template.
2. Click **Provisioning** > **Basic** > **Action** > **Apply Template**.
3. In the Template list, select from the list of previously-created DSL templates.
4. Click **Apply Template**.

For CLI:

```
apply dsl-template <t-name> to-dsl-port <port>
clear dsl-template <name> [basic|advanced|psd]
```

Related topics

- *Modifying the Basic DSL Port Template Parameters* (on page 131)
- *Modifying the Advanced DSL Port Template Parameters* (on page 135)
- *Modifying the Power Spectral Density (PSD) DSL Port Template Parameters* (on page 139)

Modifying the Basic DSL Port Template Parameters

Modifies the basic parameters for a DSL Port template that can be applied to a xDSL port, replacing the xDSL port settings.

Parameters

You can provision the following for the basic parameters of an xDSL port template:

Parameter	Description	Valid Options
Admin State	Service state of the xDSL port.	enabled-no-alarms ‡, enabled, disabled
Interface	Name of DSL bonded interface of which to add this port as a member.	EthIntf ‡, any existing bonded interface
Description	Description of DSL port.	text string
DSL GoS Profile	Index of DSL Grade of Service (GoS) profile to use for performance monitoring (PM) on the port. The ID number (and description) of a defined GoS profile. GoS profiles define performance monitoring counter thresholds. To apply a Grade of Service profile, select a profile ID from the list. You can apply the default Grade of Service profile defined by the Standards, apply a custom profile, or disable reporting of threshold crossing alerts. To take effect, the GOS feature must also be enabled.	DslPortGos: 1 ‡, any existing GoS profile
Ethernet GoS Profile	Index of Ethernet GOS profile to use.	EthPortGOS: 1 ‡

Parameter	Description	Valid Options
Service Type	Specifies the operating mode that dictates the handshaking protocol, channel capacity, and other physical line characteristics based on xDSL specifications. • **auto** – Enables the port to automatically detect and train up to the service type supported by the CPE. Enables G.DMT, T1.413, G.Lite, ADSL2, READSL2, ADSL2+,ADSL2 Annex M, ADSL2+ Annex M, and VDSL2. Note: If you want to constrain the modem from using a service type (e.g. vdsl2) that does not support the service type you want (e.g. ATM-TC), set the service type that only applies to the desired service type (e.g. mm2+) rather that using auto. • **mm** (multi-mode) – Enables the port to automatically detect and train up to the service type supported by the CPE; supports G.DMT, T1-413, and G.Lite standards • **mm2+** (multi-mode 2+) – Allows the port to automatically detect and train up to the service type supported by the CPE; supports ADSL2+, ADSL2, and READSL2 standards as well as G.DMT, T1-413, and G.Lite standards • **t1.413** – ANSI standards specification • **g.dmt** – ITU-T G.992.1 standards specification • **g.lite** – G.Lite standards specification • **adsl2** – ITU-T G.992.3 standards specification • **readsl2** (Reach Extended ADSL2) – READSL2 standards specification; supports both READSL2 and ADSL2 standards • **adsl2+** – ITU-T G.992.5 and ITU-T G.992.3 standards specification; supports ADSL2+, ADSL2, and READSL2 (G.992.3, annex-L) standards • **annexm** – ADSL2/ADSL2+ Annex M standards specification • **vdsl2** – ITU-T G.993.2 standards specification for profiles 8a, 8b, 8c, 8d, 12a, 12b, and 17a • **vdsl2mm** – (VDSL multi-mode) – Enables the port to automatically detect and train up to the service type supported by the CPE; enables ADSL2, READSL2, ADSL2+, and VDSL2 Each of the following service types use the Annex B frequency spectrum. Annex B is an addendum to the G.992.1 (G.dmt), G.992.3 (ADSL2), and G.992.5 (ADSL2+) specifications that specifies power shaping (as a function of frequency) to be used in conjunction with ISDN services. Annex B should be used when ISDN lines are contained in the same wiring bundle as ADSL/2/2+ services. • **etsi** – ETSI TS 101 388 technical specification • **gdmt-isdn** – ITU-T G.992.1 standards specification • mm-isdn (multi-mode) – Enables the port to automatically detect and train up to the service type supported by the CPE; supports ETSI and G.DMT standards • **adsl2-isdn** – ITU-T G.992.3 standards specification • **adsl2+-isdn** – ITU-T G.992.5 standards specification • **mm2+-isdn** (multi-mode 2+) – Allows the port to automatically detect and train up to the service type supported by the CPE; supports ADSL2+, ADSL2 standards as well as G.DMT and ETSI standards • **vdsl2mm-isdn** – (VDSL multi-mode) – Enables the port to automatically detect and train up to the service type supported by the CPE; enables ADSL2, ADSL2+ and VDSL2 • **auto-isdn** – Enables the port to automatically detect and train up to the service type supported by the CPE. Enables ETSI, G.DMT, T1.413, G.Lite, ADSL2, READSL2, ADSL2+,ADSL2 Annex M, ADSL2+ Annex M, and VDSL2.	auto ‡, mm, mm2+ t1.413, g.dmt, g.lite, adsl2 readsl2, adsl2+, annexm, vdsl2, vdsl2mm, etsi, gdmt-isdn, mm-isdn, adsl2-isdn, adsl2+-isdn, mm2+-isdn, vdsl2mm-isdn, auto-isdn

<table>
<tr><th>Parameter</th><th>Description</th><th>Valid Options</th></tr>
<tr><td>Path Latency</td><td>Path latency for DSL port specifies the operating mode of the channel.
• <b>fast</b> – for delay sensitive applications like voice and online gaming.
• <b>interleaved</b> – interleaves DSL frames to optimize error protection in the presence of impulse noise sources that are common to DSL.
• For IPTV, Calix recommends the "interleaved" setting in the downstream direction. Latency is tunable when using the interleaved path, Calix recommends maximizing the downstream delay of 8 ms with MS Mediaroom, or 20 ms without.
• For HSI with no IPTV, Calix recommends the "fast" setting in both upstream and downstream directions.
Interleaving leverages Reed-Solomon forward-error correction with the cost of added latency. It is useful for applications that are not too delay-sensitive and require very low bit error rates. The Fast setting provides a low-latency transmission path for delay-sensitive applications, such as Internet gaming.
The terms “fast path” and “interleaved path” are pertinent to G.dmt. In newer xDSL standards, an interleaver is always used, and that interleaver is controlled by the “Min Impulse Noise Protection” and “Interleave Max Latency” parameters. The “Path Latency” parameter allows the operator to select fast or interleaved path in G.dmt, and in other modes, to configure the interleaver to behave similarly to fast path (minimal delay, little to no impulse noise protection). For standards other than G.dmt, setting Path Latency to “fast” is equivalent to a setting of S1, as described in G.997.1, paragraph 7.3.2.2, “Maximum Interleaving Delay”.</td><td>fast, interleaved ‡</td></tr>
<tr><td>Fallback VPI</td><td>The VPI number of the fallback PVC when the port is operating in ADSL mode with ATM encapsulation.</td><td>0-255
0 ‡</td></tr>
<tr><td>Fallback VCI</td><td>The VCI number of the fallback PVC when the port is operating in ADSL mode with ATM encapsulation.</td><td>32-65535
35 ‡</td></tr>
<tr><td>VDSL2 Profile</td><td>VDSL profile for xDSL port. Applicable for VDSL2 and VDSLMM service types only.
<b>auto</b> = enables all of the profiles and allows the xDSL card to negotiate which profile to use.
<table>
<tr><th>Profile</th><th>Max Freq DS/US (MHz)</th><th>Max‡ US Train Rate</th><th>Max DS Power (dBm)</th></tr>
<tr><td>8a</td><td>8.5/5.2</td><td>18 Mbps</td><td>+17.5</td></tr>
<tr><td>8b</td><td>8.5/5.2</td><td>18 Mbps</td><td>+20.5</td></tr>
<tr><td>8c</td><td>8.5/5.2</td><td>18 Mbps</td><td>+11.5</td></tr>
<tr><td>8d</td><td>8.5/5.2</td><td>18 Mbps</td><td>+14.5</td></tr>
<tr><td>12a</td><td>8.5/12</td><td>60 Mbps</td><td>+14.5</td></tr>
<tr><td>12b</td><td>8.5/12</td><td>60 Mbps</td><td>+14.5</td></tr>
<tr><td>17a</td><td>17.7/12</td><td>60 Mbps</td><td>+14.5</td></tr>
</table></td><td>auto, 8a, 8b, 8c, 8d ‡, 12a, 12b, 17a</td></tr>
<tr><td>Report Events</td><td>Report events for DSL port.
Enables/disables reporting of port remove/restore events (including modem retrains) in the system event logs. Reporting of these events can congest the system event logs.</td><td>enabled ‡, disabled</td></tr>
<tr><td>Power Save</td><td>Enables the xDSL port power save feature.</td><td>enabled ‡, disabled</td></tr>
<tr><td>Power Save Timeout</td><td>Configures the amount of time in minutes before the xDSL port enters power save mode after the AC power is interrupted.</td><td>0-480
5 ‡</td></tr>
</table>

Parameter	Description	Valid Options
	Downstream and Upstream	
Min Rate	Minimum downstream or upstream rate for DSL port (Kb/s, or use "m" suffix for Mb/s).	0-512000 downstream: 384 ‡ upstream: 0 ‡
Max Rate	Maximum downstream or upstream of rate for DSL port (Kb/s, or use "m" suffix for Mb/s). Calix recommends a value equal to or slightly greater than the minimal bitrate required to support the corresponding service. **Note:** The maximum rate depends on the xDSL service type used. For example, for MM2+, ADSL2+, and Annex M service types, the maximum downstream rate is 32736 Kbps. For VDSL2 and VDSLMM service types, the maximum provisionable downstream rate is 512000 Kbps; however, the maximum achievable downstream rate is limited by the physical loop.	64-512000 100000 ‡
Interleave Min INP	The level of impulse noise (burst) protection for a slow (or interleaved) channel related to upstream or downstream transmissions. This parameter is defined as the number of consecutive DMT symbols or fractions thereof. The number of symbols decides how long in one period errors can be completely corrected. A higher symbol value provides higher error correction capability, but it causes overhead and higher delay which may impact multimedia data receiving quality. • For bonded VDSL2 services, Calix recommends a setting of 2 for the upstream direction. • For IPTV, Calix recommends a value of 3 for the downstream direction and a value of 1 for the upstream direction. • For HSI with no IPTV, Calix recommends a value of 0.5 for both the downstream and upstream direction.	0, 0.5, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16
Interleave Max Latency	Downstream or upstream interleave maximum latency (msec). The tolerable delay of the data transmission in milliseconds (ms) for the upstream direction on an interleaved channel (if the Channel Latency parameter value is INTLV). Select AUTO, or a number within the min/max range based on the following guidelines: • Enter a lower number to prioritize delay over performance (less delay). • Enter a higher number to prioritize performance over delay (more delay). • Select **auto** to optimize latency between delay and error performance. • For IPTV, Calix recommends a value of 8ms for applications with Media Room, or 20ms without in the downstream direction, and a value of 5ms in the upstream direction. • For HSI with no IPTV, Calix recommends a value of 5ms for both upstream and downstream directions. • For bonded VDSL2 services, Calix recommends a value of 5 ms for the upstream direction.	1-63, auto ‡
Min SNR	Minimum downstream or upstream signal-to-noise ratio (SNR) margin defines the minimum acceptable level of operation. • For IPTV, Calix recommends a value of 2 dB in both downstream and upstream directions.	0.0 ‡ - 31.0 (dB, in 0.1 dB increments)

Parameter	Description	Valid Options
Max SNR	Maximum downstream or upstream signal-to-noise ratio (SNR) margin defines the amount of margin above target SNR that must be present before power cutback occurs. • For IPTV, Calix recommends enabling power cutback by setting maximum SNR margin to 16 dB for both upstream and downstream directions. • For HSI, Calix recommends an additional 5 dB of margin above target SNR before power cutback is allowed. This corresponds to a maximum SNR margin of 10 dB for supporting HSI.	0.0-31.0 16 ‡ (dB, in 0.1 dB increments)
Target SNR	Target downstream or upstream signal-to-noise ratio (SNR) margin defines the SNR margin that must be available when the handshake process is determining the capability of each subcarrier. This has a direct impact on the attainable bitrate as a higher margin forces few symbols per constellation. • For standard DSL deployments, Calix recommends a value of 6 dB for both upstream and downstream directions. • For IPTV, Calix recommends a value of at least 8 dB for both upstream and downstream directions. • For bonded VDSL2 services, Calix recommends a value of 10 dB for the upstream direction.	0.0-31.0 6 ‡ (dB, in 0.1 dB increments)

‡Default

To modify the basic parameters in a DSL template

1. On the Navigation Tree, click **E-Series**.
2. Click **Template** > **DSL**.
3. In the Work area, double-click the DSL template that you want to modify.
4. Click **Basic** to access the basic level of parameters for the DSL template.
5. Refer to the table above for the parameter descriptions and options and set the values as necessary.
6. Click **Apply**.

For CLI:

```
set dsl-port <name> basic
set dsl-template <name> basic [dsl-bond-interface|service-type|path-
latency|fallback-vpi|fallback-vci|vdsl-profile|report-events|power-
save|power-save-timeout|description|ds-min-rate|ds-max-rate|us-min-
rate|us-max-rate|dsl-gos|eth-gos|ds-min-inp|us-min-inp|ds-intrlv-
max-latency|us-intrlv-max-latency|ds-min-snr|ds-max-snr|ds-target-
snr|us-min-snr|us-max-snr|us-target-snr|admin-state]
```

Modifying the Advanced DSL Port Template Parameters

Modifies the advanced parameters for a DSL Port template that can be applied to a DSL port, replacing the DSL port settings.

The following additional parameters apply to ADSL2+ capable ports only.

Parameters

You can provision the following for the advanced parameters of an xDSL port:

Parameter	Description	Valid Options
PTM Override	PTM override mode dictates the TC (Transmission Convergence) encapsulation on the line for the xDSL port. • **auto** – The VDSL2 card automatically chooses PTM-TC or ATM-TC if the customer deploys a multimode CPE that is capable of training up in VDSL2 or ADSL mode. Note: This option cannot be used on an xDSL port that is to be added to a bonded interface. Note: When using any service-type that includes adsl2/adsl2+, you must avoid using PTM-Override=auto. • **atm** - Asynchronous Transfer Mode (ATM-TC) is traditionally used on ADSL lines and is supported for all ADSL Service Types, but is not supported when the line trains up in a VDSL2 Service Type. ATM cells are being transmitted on the line. The DSL port MUST be set to PTM-Override=ATM in order for Multi-VC support to work. • **ptm** - Packet Transfer Mode (PTM-TC) is always used on VDSL2 lines. Ethernet frames are being transmitted on the line. An xDSL port can be forced to use PTM-TC for ADSL and ADSL2+ service types by setting PTM Override to ptm. ADSL1 service types do not support PTM-TC and will not work if PTM Override is set to ptm. **Note:** Requires corresponding feature support on the CPE.	auto, atm, ptm ‡

ATM Header Compression	Whether ATM header compression is implemented, leaving more bits available for data.
	Note: An increase in the downstream data rate is only realized with this feature if Broadcom-based CPE also implements Broadcom-proprietary ATM header compression (Nitro).
		selected=enabled,
		unselected=disabled ‡

	Downstream and Upstream	
Rate Adaption Mode	Downstream or upstream rate adaptation mode for the xDSL port. Specifies whether the data rate for the downstream or upstream signal is allowed to vary dynamically as the noise level varies. There are two values: • **init** – the rate is negotiated during handshaking and remains at its initial rate upon train-up and does not change. • **dynamic** – the rate is allowed to change automatically in response to changes in the noise level (Seamless Rate Adaptation). (Note that this may not be supported by all CPE.) **Note:** To achieve a fixed rate in a particular direction, set the min and max rates to the same value.	init ‡, dynamic
Downshift Rate Adaption Margin	Downstream or upstream downshift rate adaptation margin defines an SNR margin below the Target SNR Margin that triggers a bitrate "downshift." For the decrease to occur, the noise margin must stay below this value for the time specified in upstream or downstream Downshift Rate Adaptation Time. Applies only when dynamic rate adaptation is specified. • For ADSL2 applications, Calix recommends a value of 3 dB below target for both upstream and downstream directions. • For VDSL2 applications, Calix recommends a value of 5 dB below target for both upstream and downstream directions.	0.0-31.0 9.0 ‡ (dB, in 0.1 dB increments)

Parameter	Description	Valid Options
Upshift Rate Adaption Margin	Downstream or upstream upshift rate adaption margin defines an SNR margin above the Target SNR Margin which, when exceeded, initiates a bitrate "upshift." For the increase to occur, the noise margin must stay above this value for the time specified in upstream or downstream Upshift Rate Adaptation Time. Applies only when dynamic rate adaptation is specified. • For IPTV, Calix recommends a value of 3 dB above target for both upstream and downstream directions. • For HSI, Calix recommends a value of 1 dB above target for both upstream and downstream directions.	0.0-31.0 3.0 ‡ (dB, in 0.1 dB increments)
Downshift Rate Adaption Time	Downstream or upstream downshift rate adaptation time (sec). Applies only when dynamic rate adaptation is specified. Specifies the duration in seconds that the downstream or upstream noise margin must stay below the downstream or upstream Downshift Rate Adaptation Margin before the modem decreases the downstream or upstream data rate. The valid range is 0 to 16383 seconds. 0 disables the feature.	0-16383 60 ‡
Upshift Rate Adaption Time	Downstream or upstream upshift rate adaptation time (sec). Applies only when dynamic rate adaptation is specified. Specifies the duration in seconds that the downstream or upstream noise margin must stay above the Rate Adaptation Upshift Margin Downstream before the modem increases the downstream data rate. The valid range is 0 to 16383 seconds. 0 disables the feature.	0-16383 60 ‡
Enhanced Impulse Noise Protection	Downstream or upstream enhanced impulse noise protection mode that used in conjunction with Retransmission which is a proprietary method for retransmitting data and requires a DSL modem that is Broadcom based and PhyR enabled, or G.INP capable CPE. • For IPTV, Calix recommends using the "g.inp" or "phyr" setting in the downstream and upstream direction. **Note:** The "PhyR" setting is a special mode that gives precedence to rate. Yet, PhyR is only enabled as long as rate is not compromised. Therefore, the "phyr-fixed" setting allows you to force PhyR to be enabled, regardless of impact to rate. • Impulse Noise Immunity (G.INP) enhances techniques to protect against impulse noise for ADSL2, ADSL2+, and VDSL2. • Impulse noise is a noise event of limited duration that can degrade one or more transmitted symbols. Unlike the various types of continuous noise found on DSL, impulse noise has a short duration and may repeat, either randomly or periodically. • Impulse noise that does not appear to repeat periodically, but occur as unpredictable events is termed SHINE (Single high impulse noise event). • Impulse noise caused by noise from electrical mains and thus repeats at a constant period related to the local AC power frequency is termed REIN (Repetitive electrical impulse noise). G.INP specifies a physical layer retransmission method for enhancing INP.	none ‡, g.inp, phyr, phyr-fixed

Parameter	Description	Valid Options
Min Expected Throughput Rate	With G.INP, we use downstream/upstream Min Expected Throughput Rate (MINETR) and Max Net Data Rate (MAXNDR). MAXNDR must be configured higher than MINETR to account for retransmission bandwidth overhead. Downstream or upstream minimum effective throughput (Kbps, in 1 Kbps increments). Roughly equivalent to minimum train rate when using interleaving. • Variable rate, but fixed at train-up: For G.INP, set MINETR < MAXNDR and set rate adapt mode to init. This is similar to the fixed rate case. • Variable, dynamic rate: For G.INP, set MINETR < MAXNDR and set rate adapt mode to dynamic	32-512000 32 ‡
Max Net Data Rate	Downstream or upstream maximum net data rate (Kbps, in 1 Kbps increments). Same as maximum train rate when using interleaving.	32-512000 128000 ‡
Max Delay	Downstream or upstream maximum allowed delay for retransmission. (ms, in 1 ms increments).	1-63 20 ‡
Min Single High Impulse Noise Event (symbols)	Downstream or upstream minimum impulse noise protection against SHINE (symbols).	0-63 4 ‡
Single High Impulse Noise Ratio (symbols)	Downstream or upstream Single High Impulse Noise Event (SHINE) Ratio (NDR, in .001 increments). As stated in G.998.4, this is expected to be set by the operator using empirical methods.	0.000-0.100 0.010 ‡
Min Repetitive Impulse Noise (symbols)	Downstream or upstream minimum impulse noise protection against REIN (symbols).	0-7 0 ‡
Inter-Arrival Time Repetitive Impulse Noise (Hz)	Downstream or upstream Repetitive Impulse Noise (REIN) inter-arrival time (Hz). This is how you indicate that REIN is caused by 60Hz AC or 50Hz AC.	100, 120-hz ‡

‡Default

To modify the advanced parameters in a DSL template

1. On the Navigation Tree, click **E7**, **E3-48C**, **E5-48**, **or E5-48C**.
2. Click **Template** > **DSL**.
3. In the Workarea, double-click the DSL template that you want to modify.
4. Click **Advanced** to access the basic level of parameters for the DSL template.
5. Reference the table above to configure the parameters.
6. Click **Apply**.

For CLI:

```
set dsl-template <name> advanced [ptm-override|ds-rate-adapt-
mode|us-rate-adapt-mode|ds-downshift-adapt-margin|ds-upshift-adapt-
margin|us-downshift-adapt-margin|us-upshift-adapt-margin|ds-
downshift-adapt-time|ds-upshift-adapt-time|us-downshift-adapt-
time|us-upshift-adapt-time|ds-enhanced-inp|us-enhanced-inp|atm-
header-compress]
```

Modifying the Power Spectral Density (PSD) DSL Port Template Parameters

Modifies the Power Spectral Density (PSD) parameters for a DSL Port template that can be applied to a DSL port, replacing the DSL port settings.

PSD settings only apply to VDSL2 and VDSL2MM Service Types, and are only used when a line trains up in VDSL2 mode.

Parameters for the PSD parameters of an xDSL port

Parameter	Description	Valid Options
Limit mask	A PSD limit mask describes how the xDSL port and modem (CPE) will cap power across the range of transmit frequencies. By changing the PSD limit mask, a particular transmission scheme can be shaped so as to coexist with different underlying technologies (e.g. POTS or ISDN) on the same wire pair. In addition, you can select a VDSL QAM PSD mask. Developed by Calix in cooperation with Broadcom, this limit mask is compatible with QAM-based VDSL, which was widely deployed by older NextLevel Communications (NLC) systems and known as Classic VDSL. With this limit mask, both VDSL2 and VDSL1 QAM technologies can coexist in the same copper loop binder group. • **a-nus0** (VDSL2, Annex A, POTS compatibility, do not use band US0) • **a-eu-32 ‡, a-eu-36, a-eu-40, a-eu-44, a-eu-48, a-eu-52, a-eu-56, a-eu-60, a-eu-64, a-eu-128** (VDSL2, Annex A, POTS compatibility, end US0 on subcarrier specified) • **a-adlu-32, a-adlu-36, a-adlu-40, a-adlu-44, a-adlu-48, a-adlu-52, a-adlu-56, a-adlu-60, a-adlu-64, a-adlu-128** (VDSL2, Annex A, All Digital, end US0 on subcarrier specified) • **b8-1, b8-4** (VDSL2, Annex B, 12 MHz, US0 as in ADSL2+ annex A) • **b8-2, b8-6** (VDSL2, Annex B, 12 MHz, US0 as in ADSL2+ annex B) • **b8-3, b8-7** (VDSL2, Annex B, 12 MHz, does not use US0) • **b8-5** (VDSL2, Annex B, 12 MHz, US0 as in ADSL2/2+ annex M) • **b8-8, b8-9, b8-10** (VDSL2, Annex B, 17 MHz, does not use US0) • **b8-11** (VDSL2, Annex B, 17 MHz, US0 as in ADSL2+ annex A) • **b8-12** (VDSL2, Annex B, 17 MHz, US0 as in ADSL2+ annex B) • **b7-1** (VDSL2, Annex B, 7 MHz, US0 as in ADSL2+ annex A) • **b7-2, b7-4** (VDSL2, Annex B, 8.8 MHz, US0 as in ADSL2/2+ annex M) • **b7-3** (VDSL2, Annex B, 12 MHz, US0 as in ADSL2/2+ annex M) • **b7-5** (VDSL2, Annex B, 12 MHz, US0 as in ADSL2+ annex A) • **b7-6** (VDSL2, Annex B, 12 MHz, US0 as in ADSL2/2+ annex M) • **b7-7** (VDSL2, Annex B, 17.6 MHz, does not use US0) • **b7-9** (VDSL2, Annex B, 17.6 MHz, US0 as in ADSL2+ annex A) • **c-138-co** (VDSL2, Annex C, 12 MHz, DS1 breaks at 138 kHz, type: co) • **c-276-co** (VDSL2, Annex C, 12 MHz, DS1 breaks at 276 kHz, type: co) • **vdsl1-qam** (Calix-specific mask. Allows VDSL2 to coexist in the same binder as VDSL1 QAM.	see column left

Parameter	Description	Valid Options
	Upstream Power Back-off (UPBO)**	
Band 1-4 parameter A-B	Downstream or upstream rate adaptation mode for DSL port.	40.00-80.95 40.00 ‡
Kl_0	Upstream PBO electrical length, in dB.	0.0-128.0 no-force ‡
	Downstream Power Back-off (DPBO)***	
ESEL	Exchange-to cabinet electrical length, in dB.	0.0-255.5 0 ‡
ESCM A-C	Exchange-side cable model parameter A-C, in dB.	-1.0-1.5
MUS	Minimum usable receive PSD mask, in dBm/Hz.	-127.5 to 0.0 0 ‡ value is negative
FMIN	Downstream Power Back Off (PBO) minimum subcarrier base frequency in kHz.	0-8832 0 ‡
FMAX	Downstream PBO maximum subcarrier base frequency in kHz.	138-29997.75 29997.75 ‡
	Downstream PBO Breakpoints (DPBOEPSD)	
Frequency 1-16	Specifies a list of up to 16 breakpoints that define the Power Spectral Density Limit Mask being used at the exchange site (reference G.997.1, section 7.3.1.2.13, DPBOEPSD). For example, if the exchange site is using ADSL2+, Annex-A, non-overlapped, these breakpoints should be set to match the diagram in G.992.5, section A.1.3, figure A.2/G.992.5. Each breakpoint is defined by a frequency:psd pair, where frequency is in kHz and psd is in dBm/Hz. The system rounds the input values to the nearest subcarrier (spaced at 4.3125 kHz). The power will be rounded to the nearest 0.5 dBm/Hz.	4.3125 – 29997.75
PSD mask level 1-16	PSD mask level, in dBm/Hz.	-127.5 to 0.0 0 ‡ values are negative
	RFI Bands****	
Band 1-16 Start	Specifies the start frequency (in KHz) for RFI bands 1 to 16 where each frequency is an integer and has a resolution of 1 kHz. The start and stop frequencies define the limits of a low-power band. The system converts the input values to the closest encompassing DMT subcarrier pair (i.e. multiples of 4.3125 kHz). The default value is 0 (unused).	0 ‡ - 1000000
Band 1-16 End	Specifies the stop frequency (in KHz) for RFI bands 1 to 16. An RFI band is entered as a start-stop frequency pair where each frequency is an integer and has a resolution of 1 kHz. These frequencies define the limits of a low-power band. The system converts the input values to the closest encompassing DMT subcarrier pair (i.e. multiples of 4.3125 kHz). The default value is 0 (unused).	0 ‡ - 1000000
	Gap Bands*******	
Band 1-4 Start	Specifies the start frequencies (in kHz) for gap bands 1 to 4. A GAP band is entered as a start-stop frequency pair where each frequency is an integer and has a resolution of 1 kHz. These frequencies define the limits of a no-power band. The system converts the input values to the closest encompassing DMT subcarrier pair (i.e. multiples of 4.3125 kHz). The default value is 0 (unused).	0 ‡ - 1000000
Band 1-4 End	Specifies the stop frequencies (in kHz) for gap bands 1 to 4. A GAP band is entered as a start-stop frequency pair where each frequency is an integer and has a resolution of 1 kHz. These frequencies define the limits of a no-power band. The system converts the input values to the closest encompassing DMT subcarrier pair (i.e. multiples of 4.3125 kHz). The default value is 0 (unused).	0 ‡ - 1000000

Parameter	Description	Valid Options
	DSL Vectoring	
Downstream Vectoring Upstream Vectoring	Calix recommends that vectoring be enabled or disabled for both upstream and downstream.	enable ‡, disable
Vectoring Group	Vectoring is activated when the vectoring group is assigned to the port. **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.	Vectoring group ID.

‡Default
**For a detailed description of these parameters, see ITU-T G.997.1, section 7.3.1.2.14 Upstream power back-off shaped.
***For a detailed description of these parameters, see ITU-T G.997.1, section 7.3.1.2.13 Downstream power back-off – Shaped.
****These are bands in the frequency spectrum that are transmitted at -80dBm/Hz to reduce radio frequency interference that may be caused by the VDSL2 line.
*******These are bands in the frequency spectrum that carry no power to eliminate radio frequency interference that may be caused by the VDSL2 line.

To modify the PSD parameters in a DSL template

1. On the Navigation Tree, click **E7**, **E3-48C**, **E5-48**, **or E5-48C**.
2. Click **Template** > **DSL**.
3. In the Workarea, double-click the DSL template that you want to modify.
4. Click **PSD** to access the basic level of parameters for the DSL template.
5. Reference the table above to configure the parameters.
6. Click **Apply**.

For CLI:

```
set dsl-template <name> psd [mask|upbo-band-*-*|upbo-k10|dpbo-bp-
*|dpbo-esel|dpbo-escm-*|dpbo-mus|dpbo-fmin|dpbo-fmax|rfi-band-*|gap-
band-*]
```

Example DSL Port Template Values for Data and Video Services

This topic shows example DSL templates for VDSL2 and ADSL2+ service types.

Note: See *Calix xDSL Best Practices* for a description of the physical layer factors that may directly influence the quality of data and video services delivery, and recommendations for the best practices necessary to achieve optimal results.

Applying multiple DSL templates to a single port

Each DSL template may specify some or all of the xDSL port parameters. When a DSL template is applied to an xDSL port, the parameter values that are specified in that template are copied into the port object.

If another template is subsequently applied to the same xDSL port, the parameter values specified in that template are copied into the port object.

- If the same parameters are specified in both templates, the values defined in the last-applied template are applied to the xDSL port.

- If parameters specified in the first-applied template are not specified in the last-applied template, the values are retained from the first-applied template for those parameters.

DSL port parameters for data and video services

DSL Port Category	Parameter	Data Service	Video Service
Basic	Service Type	vdsl2 (or adsl2+)	vdsl2 (or adsl2+)
	VDSL2 Profile	17a (auto for adsl2+)	17a (auto for adsl2+)
	Downstream Min Rate	384	384
	Downstream Max Rate	100000	100000
	Upstream Min Rate	384	384
	Upstream Max Rate	100000	100000
	Downstream Min Impulse Noise Protection*	0.5	2
	Upstream Min Impulse Noise Protection	0.5	0.5
	Downstream Interleave Max Latency*	5	8◊
	Upstream Interleave Max Latency	5	5
	Downstream Min SNR*	0	2
	Downstream Max SNR*	10	16
	Downstream Target SNR*	5	8
	Upstream Min SNR*	0	2
	Upstream Max SNR*	10	16
	Upstream Target SNR*	5	8
Advanced	PTM Override	auto	auto
	ATM Header Compression	enabled	enabled
	Downstream Rate Adaption Mode	dynamic	dynamic
	Upstream Rate Adaption Mode	dynamic	dynamic
	Downstream Downshift Rate Adaption Margin*	2	5
	Downstream Upshift Rate Adaption Margin*	6	9
	Upstream Downshift Rate Adaption Margin*	2	5
	Upstream Upshift Rate Adaption Margin*	6	9
	Downstream Downshift Rate Adaption Time	30	30
	Downstream Upshift Rate Adaption Time	30	30
	Upstream Downshift Rate Adaption Time	30	30
	Upstream Upshift Rate Adaption Time	30	30
	Downstream Enhance Impulse Noise Protection	phyr	phyr
	Upstream Enhance Impulse Noise Protection	phyr	phyr

*Different DSL template parameter values are recommended for data and video services
◊ 8 ms with MMR, or 20 ms without MMR.

For CLI:

```
set dsl-template bp_data basic service-type vdsl2 vdsl-profile 17a
description bp_data ds-min-rate 384 ds-max-rate 100000 us-min-rate
384 us-max-rate 100000 ds-min-inp 0.5 us-min-inp 0.5 ds-intrlv-max-
latency 5 us-intrlv-max-latency 5 ds-min-snr 0 ds-max-snr 10 ds-
target-snr 5 us-min-snr 0 us-max-snr 10 us-target-snr 5 admin-state
enabled
```

```
set dsl-template bp_data advanced ptm-override auto ds-rate-adapt-
mode dynamic us-rate-adapt-mode dynamic ds-downshift-adapt-margin 2
ds-upshift-adapt-margin 6 us-downshift-adapt-margin 2 us-upshift-
adapt-margin 6 ds-downshift-adapt-time 30 ds-upshift-adapt-time 30
us-downshift-adapt-time 30 us-upshift-adapt-time 30 ds-enhanced-inp
phyr us-enhanced-inp phyr atm-header-compress enabled

set dsl-template bp_video basic service-type vdsl2 vdsl-profile 17a
description bp_video ds-min-rate 384 ds-max-rate 100000 us-min-rate
384 us-max-rate 100000 ds-min-inp 2 us-min-inp 0.5 ds-intrlv-max-
latency 8 us-intrlv-max-latency 5 ds-min-snr 2 ds-max-snr 16 ds-
target-snr 8 us-min-snr 2 us-max-snr  16 us-target-snr 8 admin-state
enabled

set dsl-template bp_video advanced ptm-override auto ds-rate-adapt-
mode dynamic us-rate-adapt-mode dynamic ds-downshift-adapt-margin 5
ds-upshift-adapt-margin 9 us-downshift-adapt-margin 5 us-upshift-
adapt-margin 9 ds-downshift-adapt-time 30 ds-upshift-adapt-time 30
us-downshift-adapt-time 30 us-upshift-adapt-time 30 ds-enhanced-inp
phyr us-enhanced-inp phyr atm-header-compress enabled
```

Creating an Ethernet Security Profile

This topic shows you how to create an Ethernet security profile that is applied to xDSL port interfaces to enable and/or specify security attributes of the port.

Note: The E-Series implementation of security profiles applies to non-TLAN services only. For TLAN services, the L2CP Filter parameter must be set to all-tunnel.

Layer 2 Control Protocol (L2CP) handling

The E-Series supports provisionable Layer 2 Control Protocol (L2CP) handling (tunnel, discard) on xDSL port 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:

- Bridge Block of protocol frames with destination MAC addresses 0x01:80:c2:00:00:00 through 0x01:80:c2:00:00:0f.
- GARP Block of protocol frames with destination MAC addresses 0x01:80:c2:00:00:20 through 0x01:80:c2:00:00:2f.
- All LANs Bridge Management Group protocol frames with destination MAC address 0x01:80:c2:00:00:10.

Parameters

You can provision the following parameters for Ethernet security profiles:

Parameter	Description	Valid Options
Name*	Name of the security profile.	up to 32 character string

Parameter	Description	Valid Options
Source MAC Limit	Number of unique MAC addresses allowed on a subscriber port. **Note:** A limit of 0 means “do not enforce a limit” (that is, unlimited).	0 ‡-255
Source MAC Age (s)	Maximum age for source MAC addresses, in seconds.	60-86400 300 ‡
DHCP Lease Limit	Specifies the maximum number of DHCP leases allowed on the xDSL port (aggregate for all services). 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 and applies to IPv4 and IPv6 addresses, independently. A security profile used by an xDSL port must have a DHCP lease limit value of 10 or less. The E7 does not support infinite lease times on DHCP snooping.	1-16 8 ‡
Upstream Broadcast/Multicast Limit (Kb/sec)	Specifies the maximum rate of Layer 2 broadcast traffic per second allowed on the xDSL port.	0 – 10240 Kbps (10 Mbps) 24 ‡
L2CP Filter	Layer 2 Control Protocol filter sets whether to pass or discard the L2CP protocol frames. Besides selecting from the system default filters (all-discard or all-tunnel), you can create additional filters to specify which L2CP ranges to discard or tunnel. Note: For TLAN services, this parameter must be set to all-tunnel.	all-discard ‡ all-tunnel, or any previously-created filter

DOS Attack Detection	Whether to enable or disable Denial Of Service (DOS) attack detection against PPPoE control flows. DOS attack is implemented for PPPoE discovery phase packets only, where it meters the arrival of PPPoE control packets from a particular subscriber, and uses that rate to detect an arrival rate being above a threshold for a time. When the threshold is detected, the condition is considered "yellow". When the threshold continues for a specified time, the condition is considered "red" and a DOS attack condition exists. Counters are kept for packets arriving in the "yellow" and "red" conditions.
	Threshold rates:
	• Yellow - 10 pkts/second
	• Red - 10 pkts/second for 5 consecutive seconds
	• When the Red condition occurs, the interface is disabled for 300 seconds.
		selected = enabled ‡
		unselected = disabled

Allow IPv6	IPv6 traffic flows by default. If the setting in this profile is changed to "Disabled", IPv6 unicast, multicast and broadcast traffic ingressing the xDSL interface is blocked.	selected = enabled ‡
		unselected = disabled

802.1x Profile	Specifies the previously created 802.1x profile to apply to the security profile.	1-10

*Required field
‡ Default

To create a Ethernet security profile

1. On the Navigation Tree, click the unit.
2. Click **Profiles** > **Security > Ethernet > Profiles.**
3. In the menu, click **Create** to open the Create Security Profile dialog box.
4. Configure the Ethernet Security Profile, as described in the table above.
5. Apply the Ethernet Security Profile to an xDSL port associated interface by selecting the profile by name under the **Security Profile** option.

For CLI

- **`create eth-sec-profile <p-name> [src-mac-limit|src-mac-age|dhcp-lease-limit|upstrm-bcast-mcast-limit|l2cp-filter|dos-attack-detection|allow-ip-v6|dot1x-profile]`**
- **`create l2cp-filter <name> [range-1-action|range-2-action|range-3-action]`**
- **`delete eth-sec-profile`**
- **`set eth-sec-profile <p-name> [src-mac-limit|src-mac-age|dhcp-lease-limit|upstrm-bcast-mcast-limit|l2cp-filter|dos-attack-detection|allow-ip-v6|dot1x-profile]`**
- **`set l2cp-filter <name> [range-1-action|range-2-action|range-3-action]`**
- **`delete l2cp-filter <name>`**
- **`create dot1x-profile <name>`**
- **`set dot1x-profile <name> [reauth-period|quiet-period|reauth-timer|max-retries|svr-timeout|retrans-timer]`**
- **`show dot1x-profile`**
- **`show eth-sec-profile [<p-name>]`**

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 E-Series 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 an xDSL port interface to enable and/or specify security attributes of the xDSL port.

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

Layer 2 Control Protocol (L2CP) handling

The E-Series supports provisionable Layer 2 Control Protocol (L2CP) handling (tunnel, discard) on xDSL 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.

Parameters

You can provision the following parameters for L2CP Filters:

Parameter	Description	Valid Options
Name*	Name of the L2CP filter.	up to 32 character string
BPDU	Specifies the action for layer-2 control protocol packets in the BPDU range.	discard ‡ tunnel
GARP	Specifies the action for layer-2 control protocol packets in the GARP range.	discard ‡ tunnel
ALL-LANS	Specifies the action for layer-2 control protocol packets in the ALL-LANS range.	discard ‡ tunnel

*Required field
‡ Default

To create an L2CP filter

1. Access the profile page:
   - From CMS:
     - On the Navigation Tree, click **CMS**.
     - In the Work Area, click **Profile > E3-48C/E5-48/E7/ONT > Profile** > **Security > L2CP**.
   - Locally on the E-Series:
     - On the Navigation Tree, click the unit.
     - In the Work Area, click **Profiles** > **Security > L2CP** > **Filters**.
2. In the menu, click **Create** to open the Create L2CP Filter dialog box.
3. Reference the table above to configure the parameters.
4. Click **Create**.

For CLI

```
create l2cp-filter <name> [range-1-action|range-2-action|range-3-action]
delete l2cp-filter <name>
```

Creating a PPPoE Profile

This topic shows you how to create a PPPoE profile that is referenced by the service VLAN that is configured for PPPoE operation.

See *Configuring PPPoA/PPPoE Operation for a Data Service* (on page 184) for configuration guidelines and instructions.

Configuration guidelines:

- The system 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 system, the PPPoE traffic must adhere to a specific VLAN membership on the Trunk interface.
- The PPPoE profile is applied to a service VLAN where simultaneous operation of DHCP Snooping and PPPoE are not supported. Both PPPoE and DHCP are mechanisms for a subscriber host to acquire an IP address with which to communicate. If a service is using PPP, then it is not using DHCP, and vice versa.
- The PPPoE profile must be set to Auto mode for the InterWorking function (PPPoA-to-PPPoE conversion).
- 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).
- Up to one service on each xDSL interface (including bonded-links) may resolve to a VLAN with an associated PPPoE profile.
- To disable downstream broadcast traffic on the E-Series, enable a PPPoE profile on the service VLAN to only forward PADI broadcast packets upstream to the BRAS port once it detects which one it is.

Parameters

You can provision the following parameters for PPPoE profiles:

Parameter	Description	Valid Options
Name*	Descriptive name for the PPPoE profile.	text string
Mode	The mode to use for the operation: • **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.	10-300 300 ‡
Discovery Timeout (s)	Discovery timeout (seconds).	1-30 3 ‡

Parameter	Description	Valid Options
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

To create a PPPoE profile

1. On the Navigation Tree, click **E7/E5-48/E3-48C**.
2. Click **Profiles** > **PPPOE**.
3. In the menu, click **Create**.
4. Reference the table above to configure the parameters.
5. Click **Create**.

For CLI:

```
create pppoe-profile <name> [mode|stale-timeout|disc-
timeout|allowed-bng-1|allowed-bng-2|allowed-bng-3|allowed-bng-4]
```

Modifying an Access-Identifier Profile for xDSL Ports

This topic shows you how to modify the Access-Identifier profile "eth-system-default" that is applied to xDSL and Ethernet ports for PPPoE operation.

Also, the E-Series provides a default Calix-format syntax and a TR-101-format that complies to TR-101 R-124 requirements.

- 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>
      - 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

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 63 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').

To configure the global access-identifier profile for xDSL ports

1. On the Navigation Tree, select **E-Series**.
2. In the work area, click **Profiles > Access Identifier** to view the table of default access-identifier profiles.
3. Double-click the name of the profile that you want to configure:
   - **eth-system-default** is used for xDSL and GE ports.
4. In the Access Identifier Profile form, select the parameter from the attribute list:
   - Circuit ID list parameters:
     - **calix-format**
     - **tr101-format**

- Remote ID list parameters:
  - **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 63 characters of the Subscriber ID text field are inserted.
  - **none**

5. In the toolbar, click **Apply**.

For CLI:

```
set access-identifier-profile <eth-system-default> remote-id
[subscriber-id|none]

set access-identifier-profile <eth-system-default> circuit-id
[calix-format|tr101-format]
```

Configuring a Bonded Interface

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 two individual lines. Calix E5 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 single port. 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.

Configuration guidelines

- A maximum of two ports are supported per bonding group.
- A maximum of 24 Bonding Groups are supported on a single VDSL2 card.
- Maximum throughput for a bonding group:
  - 100 Mbps downstream/60 Mbps upstream using a 12a DSL profile (20 Mbps per DSL port)
  - 100 Mbps downstream/12 Mbps upstream using an 8a DSL profile

  Note: Although bonded ports can be up to 60Mb/s downstream, bonded together the group has a 100Mb/s limit.

- The member ports must have the same Service Type, either all VDSL2 or all ADSL service types.

- All ports must have the same encapsulation (ATM-TC or PTM-TC).
    - For ports with PTM encapsulation, bonding of any xDSL port pairs is supported on the same card.
    - For ports with ATM encapsulation, bonding is restricted to pairs of odd/even adjacent ports.
- 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 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
- The xDSL ports that you want to add to a bonding group at the time of creation must first be configured with the valid combination of service type and PTM Override attributes. For VDSL bonding, use PTM Override = PTM (not Auto).
- 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 E5 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 segmentation of the larger Ethernet frames carrying IP traffic serviced all the way from the subscriber end devices, through the access network and into the core.
- Calix recommends using the same DSL port settings that you would use for a single-pair.
- 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.
- It is important to match the transport mode desired with the modem’s ability to support such a mode, as not all ADSL2+ capable modems support PTM mode.

- Once the bonding group has been configured, the subscriber service provisioning is on the group interface.

  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.
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.

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 ‡, disabled
Subscriber ID	Subscriber ID information, such as phone number, or account number.	text string (blank) ‡
Description	Description of the DSL bonded interface, such as the subscriber address, name, or service.	text string (blank) ‡
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.	text string

Parameter	Description	Valid Options
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 ‡, IP PrecMap:access, any previously created profile
Member 1	The xDSL port to assign as the first member of a bonding group. The xDSL ports that you want to add to a bonding group at the time of creation must first be configured with the valid combination of service type and PTM Override attributes.	any xDSL port
Member 2	The xDSL port to assign as the first member of a bonding group. The xDSL ports that you want to add to a bonding group at the time of creation must first be configured with the valid combination of service type and PTM Override attributes.	any xDSL port

* Required
‡ Default

To create a DSL Bonding Group

1. Choose ports on a VDSL2 card intended for DSL pair bonded services, considering the following:
   - The member ports must have the same Service Type, either all VDSL2 or all ADSL service types.
   - All ports must have the same encapsulation (ATM-TC or PTM-TC).
     - For ports with PTM encapsulation, bonding of *any* xDSL port pairs is supported on the same card.
     - For ports with 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].
2. Configure the xDSL ports that will be members of the bonding group, identified in Step 1.

   a. On the Navigation Tree, click the first xDSL port that you identified in Step 1.

   b. In the Work Area, click **Port** > **Provisioning** > **Basic**.

   c. Configure the Basic settings on the bonding group member's xDSL port:

   - Service Type (all VDSL2 or ADSL service types and avoid auto)
   - Path Latency
   - VDSL2 Profile
   - Min Rate (Downstream and Upstream)
   - Max Rate (Downstream and Upstream)
   - SNR margins

   d. In the menu, click **Apply** for the changes to take effect.

e. Click **Advanced**, and then configure the Advanced setting on the bonding group member's xDSL port:

- PTM Override (auto is not allowed)

Note: PhyR is not supported on bonded lines. If PhyR is enabled by the user, it will be automatically disabled by the hardware.

f. In the menu, click **Apply** for the changes to take effect.

g. Repeat Step 2 to configure the second xDSL port that will be a member of the bonding group.

3. In the Navigation Tree, click the VDSL2 card which is intended to have the bonding group.

4. In the Work Area, click **Provisioining** > **Create** > **XDSL Group**.

5. In the Create DSL Bonding Group dialog, do the following:

a. In the Name box, enter a descriptive name for the bonding group.

b. In the Admin State list, select whether to enable the bonding group.

c. In the Subscriber ID box, enter information that is unique to the subscriber, such as an account number or address.

d. In the Description box, enter further information that would be helpful to identify the bonding group.

e. In the IGMP Immediate Leave list, select whether a multicast stream is dropped as soon as a Leave is received.

Note: This 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.

f. In the DS Min Rate and US Min Rate boxes, enter values that are used to generate low-rate alarms.

g. In the Security Profile list, select the name of the existing security profile to use for the bonding group.

h. In the DSCP/IP Precedence Map list, select the name of the DSCP or IP Precedence profile to use for mapping Layer 3 priority values to P-bits.

i. In the Member 1 list, select the xDSL port to assign as the first member of a bonding group.

j. In the Member 2 list, select the xDSL port to assign as the second member of a bonding group.

Note: The xDSL ports that you want to add to a bonding group at the time of creation must first be configured with the valid combination of service type and PTM Override attributes.

6. Click **Create**. When the Operation Status in the Task Progress window shows Success, click **OK**.
7. To add a service to an xDSL Bonding Group, see *Configuring Data Services* (on page 177) and *Configuring Video Services* (on page 189).

For CLI:

- `create dsl-bond-interface <intfc-name> [description|subscriber-id|dscp-p-bit-map|ip-prec-p-bit-map|eth-sec-profile|immediate-leave|ds-min-rate|us-min-rate|admin-state]`
- `set dsl-port <id> <intfc-name>`
- `set dsl-bond-interface <intfc-name> [eth-svc|description|subscriber-id|dscp-p-bit-map|ip-prec-p-bit-map|eth-sec-profile|immediate-leave|ds-min-rate|us-min-rate|admin-state]`
- `show dsl-bond-interface [detail|vlans]`
- `show dsl-bond-interface <intfc-name> [detail|eth-svc|members|vlans]`
- `show dsl-port [port] [advanced|all|basic|inventory|psd|status|subcarriers]`
- `clear dsl-bond-interface igmp-counters`
- `clear dsl-bond-interface <intfc-name> [igmp-counters|pppoe]`

Creating Voice Service Profiles

This section describes how to create profiles and objects that will be associated with voice services provisioned on a VDSL2 card Voice (POTS) port.

Creation of following profiles and objects are described:

- IP Host for Voice Services
- Dial Plan
- SIP Gateway Profile
- TDM Gateway Profile
- TDM Gateway Service Group
- H.248 Gateway Profile
- H.248 Gateway

Next steps:

After completing the creation of profiles for voice services, see the following section to continue configuring voice services:

Configuring Voice Services (on page 200)

Creating an IP Host for Voice Services on a VDSL2 Card

This topic describes how to configure an IP host that defines how a VDSL2 card obtains an IP address for communication and specifies the voice service VLAN. (The definition references a system default tag action that specifies the classifying and marking of packets from the subscriber port into the service VLAN specified in the IP host.)

Each VDSL2 card supports one voice service type, either SIP or TDM gateway or H.248 gateway. Therefore, a VDSL2 card supports a maximum of one IP host for voice services.

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.

Note: When the IP Host protocol is changed from Static to Dynamic at the VDSL line card level, the applied SIP Gateway Profile DNS entries must be populated with zero entries (0.0.0.0).

- 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

Parameter	Description	Valid Options
S-VLAN (Outer Tag)	Indicates the customer-specific tag for the VoIP service VLAN. The definition references a system default tag action that specifies the classifying and marking of packets from the subscriber port into the service VLAN specified in the IP host.	2-4093 (Except for 1002-1005 which are reserved for E-Series 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 ‡, dhcp-v6
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
Ping	Whether to respond to ping messages.	enabled ‡, disabled
Traceroute	Whether to respond to traceroute messages.	enabled ‡, disabled
Configuration File Instance	A configuration name instance is mapped to a retrieved DSL Configuration file, whose contents will be applied to all DSL ports on the VDSL2 card. The configuration name specified in the IP Host must match the configuration name specified in the Retrieve/Apply DSL Configuration File operation. A configuration name is specified by instance number alone or by “voip-<1 through 4>.” • voip-1 - Instance for VoIP (1) • voip-2 - Instance for VoIP (2) • voip-3 - Instance for VoIP (3) • voip-4 - Instance for VoIP (4)	None ‡, voip1-4, enter value

To create an IP Host for voice services

1. From the Navigation Tree, click the VDSL card where you want to create an IP Host for voice services.
2. In the Work Area, click **IP Hosts** > **Create**.
3. Reference the table above to configure the parameters.
4. Click **Create**.

For CLI:

```
create ip-host <[shelf]card/name>
delete ip-host <[shelf]card/name>
set ip-host <[shelf]card/name>
show ip-host [[shelf]card/name]
```

Example:

```
create ip-host 1/tdmgw outer-vlan 300 host-config static static-ip
192.168.32.100 static-gw 192.168.32.1 static-netmask 255.255.240.0
```

Creating a Dial Plan and Rules

This topic describes how to create and delete a number plan table that is assigned to a SIP gateway service provisioned on a Voice port.

Number plans are used to identify specific types of phone numbers dialed by a subscriber, and to process the number before transmission by deleting, replacing, or adding digits according to the relevant rule. The rule can also automatically add the country code and national destination (region) code, or deny the number pattern entirely.

If a custom numbering plan is not applied when you create a SIP gateway service, the default numbering plan table (Access) is applied which contains a generic digit entry rule and a rule for 911 calling. Up to twenty number plans may be created per E-series node or modular chassis system.

See the "Dial Plan Example" section in the *Calix Application Note: Using the ONT VoIP Configuration Files* for an example of a dial plan and an explanation of each code and its function.

Configuration considerations

- On legacy systems, the dial plan is embedded in the SIP gateway object.
- On current systems using the SIP Gateway with per-port provisioning method, a separate SIP dial plan profile can be created and attached to the SIP port service.
- When an applied dial plan is modified, the modification is pushed down to the configured ONTs with no further input.
- T-Series ONTs only accept H.248 standard dial plan rules via OMCI. However, enhanced dial plan symbols are supported through the XML voice configuration files.
- The following restrictions apply when creating dial plan rules:
  - A “|” rule separating character is required at the end of each rule, which limits the rule to 27 characters per the G.984.4 OMCI table entry length.
  - When creating the rule through CLI, the pattern must begin with a “^” to indicate the start of a rule.
  - In general, dial plan rules are wrapped in parentheses with an open parenthesis before the first rule and a closing paren following the last rule within the dial plan.
  - Each table row allows for 28 bytes (27 for the rule plus one for the separator/terminating character); if multiple rules can fit within a single row, they will be separated by the “|” indicator.
  - Repeated ranges (a range followed by {n}) get expanded out before being sent to T-Series ONTs, when the range is not [0-9].

    For example, [0-9] {5} will not expand larger (it will be xxxxx) but [1-9] will expand larger as [1-9][1-9][1-9][1-9][1-9], and the total length must not exceed 28 bytes.

Digit collection timeout and calling

The default digit collection timeout (referred to as the inter digit timeout) is 10 seconds for defined numbering plan tables. If at any point during the 10-second timeout sufficient digits are collected, the call is made immediately.

Note: If you apply a custom numbering plan with no rule entries, a digit collection timeout of three seconds is enforced, regardless of how many digits are collected.

System-default dial plan rules

Rule ID	Pattern
system-default-1	^911n
system-default-2	^411
system-default-3	^[2-9][0-9]{6}
system-default-4	^1[2-9][0-9]{9}
system-default-5	^011[0-9]*T
system-default-6	^S[0-9]{2}

Parameters

You can provision the following parameters for creating a dial plan and adding rules:

Parameter	Description	Valid Options
Name*	Descriptive name for the dial plan.	text string up to 31 characters
Digit Short Timer(s)	Value to use that resolves an overlapping dial plan, where you can enter additional digits if needed, causing a much faster timeout and sending the digit string / INVITE much quicker.	1-16 4 ‡
Digit Long Timer(s)	Value = 16 to use as the first digit timer (where dial tone will time out) as well as in conditions where no dial plan match has been found so more digits are required. Once a partial match is found the long timer is no longer used.	4-20 16 ‡
ID	Index of dial plan rule.	1-100

Parameter	Description	Valid Options
Pattern	Dial plan rule pattern. • **^** Required to match from the start of the dial string • **|** Required vertical bar (pipe key) as a rule-separating character at the end of each rule in the dial plan • **[a-b]** Square brackets are used to define options or sub-ranges of allowable digits • **{n}** Curly brackets are used to define the number of of allowed digits in a string (range match length). Applies only to variable directly preceding it. • ***** Wild card match - Matches on a variable number of digits • **T** Variable digit timeout • **S** Star key on the handset - applies to Vertical Service Codes • **c** Confirmation tone is played after star code is executed • **r** Recall tone is played during call forwarding sequence • **d** Dial tone is played during a call forwarding sequence • **,** (Comma) Outside dial tone is played if preceded by a 9 • **n** no local disconnect • # Pound indicator • **b** /*Must immediately follow '#'*/ **Note:** The Maximum Network Dial Plan Table size is 100 rows x 28 (2800 bytes). A 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.	^ | [a-b] {n} * T S c r d , n #

* Required
‡ Default value

To create a dial plan

1. On the Navigation Tree, click the unit.
2. Click **Profiles** > **Service** > **Dial Plan**.
3. In the menu, click **Create**.
4. Reference the table above to configure the parameters.
5. Click **Create**.
6. Apply the custom dial numbering plan when you create a SIP gateway service.

For CLI:

```
create dial-plan <name> [digit-short-timer|digit-long-timer]
delete dial-plan <name>
set dial-plan <p-name> [rule|name|digit-short-timer|digit-long-
timer]
show dial-plan [<name>]
```

To create dial plan rule

1. On the Navigation Tree, click **E-Series**.
2. Click **Profiles** > **Service** > **Dial Plan**.
3. In the Work area, double-click the dial plan of which you want to add a rule.
4. In the menu, click **Create** > **Dial Plan Rule**.
5. In the Create Dial Plan Rule dialog box, do the following:
   a. In the ID list, select the index value to use if different from the default rule of "1."
   b. In the Pattern box, type the pattern string and rule string.

      The pattern must start with `^` and end with `|`.

      For allowed characters, refer to the above table.
6. Click **Create**.

For CLI:

```
add rule <p-index> to-dial-plan <p-name> pattern <token>
delete dial-plan <name>
set dial-plan <p-name> [rule|name|digit-short-timer|digit-long-
timer]
show dial-plan [<name>]
```

Creating a SIP Gateway Profile

This topic describes how to create a SIP gateway profile that is assigned to a SIP gateway service provisioned on a Voice port.

You can define a unique SIP Profile ONTs and E-series xDSL nodes, entering attributes for both the Primary Proxy Server settings, Primary and Secondary DNS servers, as well as Secondary Proxy Server settings if enabling SIP Server Redundancy using static host provisioning.

Parameters

You can provision the following parameters for SIP gateway profiles:

Parameter	Description	Valid Options
Name*	Name of the SIP profile.	Any established SIP service profile

Parameter	Description	Valid Options
Proxy Server	IP address or hostname of the SIP proxy server of the SIP server or outbound proxy SIP server. If the primary path or server is disrupted, the ONT or VDSL2 card will resolve to a pre-provisioned secondary server without the need for DNS. **Note:** A DNS server is required if this parameter value is a Fully-Qualified Domain Name (FQDN) of the SIP server. In R2.4 and higher, the FQDN is used in both the DNS and SIP request. Customers using DNS must set domain=IP to continue using an IP address for SIP requests. **Note:** It should be noted that T-series ONTs requires the proxy server and register server to be set to the same value. When the proxy server is specified, OMCI will set both the proxy server and register server. The P-series ONTs will ignore attributes it does not understand.	IP address, hostname
Proxy Server Port	UDP port for proxy server.	0-65535 5060 ‡
Secondary Proxy Server	IP address of the secondary SIP proxy server, or outbound proxy SIP server. If the primary path or server is disrupted, the SIP client remains connected, as long as the secondary proxy functions correctly. When the secondary proxy detects a failure or is disabled, the SIP client will again try to switch back to the primary SIP proxy. • The secondary proxy server can only be configured *if* the proxy server is configured as an IP address (not a hostname). • The secondary proxy server can only be configured as an IP address (not a hostname). **Note:** Not supported on the T-Series ONTs.	IP address
Secondary Proxy Server Port	UDP port for secondary proxy server. **Note:** Not supported on the T-Series ONTs.	0-65535 5060 ‡
Primary DNS Server*	IP address or hostname of the primary DNS server. If you prefer DNS from the IP Host, leave the IP address as 0.0.0.0, and then instead the DNS from the DHCP lease will be handed down for host resolution.	IP address
Secondary DNS Server*	IP address or hostname of the secondary DNS server. If you prefer DNS from the IP Host, leave the IP address as 0.0.0.0, and then instead the DNS from the DHCP lease will be handed down for host resolution.	IP address
RTP Codec First Order	The Realtime Transport Protocol (RTP) code to use. • **a-law** algorithm is commonly used in Europe. • **u-law** (mu-law or µ-law) algorithm is commonly used in the USA and Japan. • **g723** (T-series only) G723 encoding • **g729** (T-series only) G729 encoding	u-law ‡, a-law, g723, g729
Packet Rate First Order	(xDSL or T-series only) The expected RTP packet rate sent by the ONT or E-series (packets/msec).	10ms ‡, 20ms, 30ms
Silence Suppression First Order	Whether to enable Silence Suppression first-order priority codec.	selected (enabled), unselected (disabled) ‡

Parameter	Description	Valid Options
RTP Codec Second Order	The Realtime Transport Protocol (RTP) code to use. • **a-law** algorithm is commonly used in Europe. • **u-law** (mu-law or µ-law) algorithm is commonly used in the USA and Japan. • **g723** (T-series only) G723 encoding • **g729** (T-series only) G729 encoding • **none** all settings for this order are copied from the First Order values.	none ‡, u-law, a-law, g723, g729
Packet Rate Second Order	(xDSL or T-series only) The expected second-order RTP packet rate sent by the ONT or E-series (packets/msec).	10ms ‡, 20ms, 30 ms
Silence Suppression Second Order	Whether to enable Silence Suppression second-order priority codec.	selected, unselected ‡
RTP Codec Third Order	The Realtime Transport Protocol (RTP) code to use. • **a-law** algorithm is commonly used in Europe. • **u-law** (mu-law or µ-law) algorithm is commonly used in the USA and Japan. • **g723** (T-series only) G723 encoding • **g729** (T-series only) G729 encoding • **none** all settings for this order are copied from the First Order values.	none ‡, u-law, a-law, g723, g729
Packet Rate Third Order	(xDSL or T-series only) The expected third-order RTP packet rate sent by the ONT or E-series (packets/msec).	10ms ‡, 20ms, 30ms
Silence Suppression Third Order	Whether to enable Silence Suppression third-order priority codec.	selected (enabled), unselected (disabled) ‡
T1 Timer (ms)	T1 and T2 are SIP timers. T1 is an estimate of the round trip time, the client will start to retransmit an INVITE transaction at T1 and then double the time for each subsequent retransmission.	100-1500 500 ‡
T2 Timer (s)	T1 and T2 are SIP timers. T2 is the maximum retransmit interval for non-INVITE requests and INVITE responses. These will rarely change.	1-5 4 ‡
Registration Period(s)	Duration of the SIP registration request.	60-86400 3600 ‡
Distinctive Ring Prefix	Distinctive ring prefix is an identifier used in the Alert-Info header field (up to 36 characters). The E-Series acts as the media gateway to support different ring cadences generated from the softswitch in real time. Assigning distinctive ring patterns to specific incoming numbers for a specific VoIP port must be configured on the softswitch. Important: The case-sensitive name must match the Distinctive Ringing text string used by the softswitch. The softswitch sends an INVITE to the E-Series with the header Alter-Info field, for example: AlertInfo = <xxx://xxx.xx.xx/xxx/Bellcore-dr1> The string after last forward slash (/) is parsed as the ring ID. When a line is in ring state, the ring cadence is used based on the Alert-Info header in INVITE message. The cycle that is defined by the ring cadence index is repeated until the line leaves the ring state.	text string Bellcore-dr ‡
Call Waiting Prefix	Call-waiting ring prefix.	text string CallWaitingTone ‡

Parameter	Description	Valid Options
Out Of Band DTMF	Out-of-band Dual-Tone Multi-Frequency (DTMF) mode. • Select **Info** to relay DTMF tones as SIP INFO messages. • Select **rfc2833** to relay DTMF tones according to RFC 2833. (Only supported on GigaCenter ONTs.) • Select **none** to not relay DTMF tones.	none ‡, rfc2833, info
Local Hook Flash	Defines where hook-flash control resides. • When enabled (selected), the local User Agent will consume the hook-flash and provide the service locally. • When disabled (unselected), then the hook-flash is passed to the Softswitch for processing.	selected (enabled) ‡, unselected (disabled)
RTP DSCP	The DiffServ Code Point (DSCP) value for traffic using this SIP profile.	• **0-63** - DSCP for RTP packets • cs0 - DSCP CS0 (0) • cs1 - DSCP CS1 (8) • af11 - DSCP AF11 (10) • af12 - DSCP AF12 (12) • af13 - DSCP AF13 (14) • cs2 - DSCP CS2 (16) • af21 - DSCP AF21 (18) • af22 - DSCP AF22 (20) • af23 - DSCP AF23 (22) • cs3 - DSCP CS3 (24) • af31 - DSCP AF31 (26) • af32 - DSCP AF32 (28) • af33 - DSCP AF32 (30) • cs4 - DSCP CS4 (32) • af41 - DSCP AF41 (34) • af42 - DSCP AF41 (36) • af43 - DSCP AF43 (38) • cs5 - DSCP CS5 (40) • ef - DSCP EF (46) ‡ • cs6 - DSCP CS6 (48) • cs7 - DSCP CS7 (56)
RTP Ethernet QoS	Ethernet QoS for RTP packets override.	0-7 6 ‡

Parameter	Description	Valid Options
Domain	Allows you to specify an internet type domain address. Alternatively, "none" can be used for no domain name. If present, the domain is used to populate the SIP destination addresses; the “to” and “request uri” fields. To use IP in SIP request, set and IP address. If not present, the proxy server IP is used in these fields. **Example:** • Domain: “empty” • Proxy-server: 10.0.20.10 Generated by software: • SIP to: sip:7663339@10.0.20.10 • SIP request URI: sip:7663339@10.0.20.10:5060 **Example:** • Domain: mytelco.com • Proxy-server: 10.0.20.10 Generated by software: • SIP to: sip:7663339@mytelco.com • SIP request URI: sip:7663339@mytelco.com:5060	1-63 character text string, none ‡
Country Code	E.164 Country code designator (Protocol Country Variant profile). This attribute specifies the country code where the service is being deployed. This code selects country specific tone settings, line interfaces, line levels and line frequencies. Currently supported values include the following: • North America: 1 (GPON and xDSL) • Italy: 39 (xDSL only) • Switzerland: 41 (GPON only) • United Kingdom: 44 (GPON only) • Sweden: 46 (GPON only) • Poland: 48 (GPON only) • Brazil: 55 (GPON only) • Australia: 61 (GPON only) • New Zealand: 64 (GPON only) • Ukraine: 380 (GPON only) • ETSI: 9000 (GPON only) **Note:** All ONTs using this profile will reset if country-code is modified. **Note:** Country codes are not currently supported on T-series ONTs or P-series 700GX ONTs.	1‡ -9999
Release Timer (s)	Specifies the amount of time it takes to terminate a call after an on-hook is detected.	1-20 10‡
RTP Port	Identifies the starting RTP Port range for the SIP RTP path.	49152 ‡ - 65535

Parameter	Description	Valid Options
Switch Type	Populates the Softswitch attribute of the SIP agent configuration data ME with the value provided by the Switch Type attribute of the SIP gateway profile. **zte** • Unreserved URI's will be escaped. • Only request URI is used to find the SIP line when request is received by the SIP User Agent. **huaw** (Huawei) • The # shall be transmitted in the SIP Invite instead of escape quoted with %23. • Request URI and URI in the header can be used to find the sip line. • UA profile on subscribe message (RFC 6080) is supported. **syla, ERIC, CS2K, BELL** • The # is escape-quote (%23) in the SIP Invite. • Request URI and URI in the header can be used to find the SIP line. • **None** for supported switch types that are not in the list.	none ‡, zte, huaw, syla, eric, cs2k, bell

*Required fields
‡ Default

To create a SIP gateway profile

1. On the Navigation Tree, click the unit.
2. Click **Profiles** > **Service > SIP GW** > **Profiles**.
3. In the menu, click **Create**.
4. Reference the table above to configure the parameters.
5. Click **Create** to save the profile.

Apply a SIP profile when you create a SIP voice service on a voice port.

For CLI:

- `create sip-gw-profile <p-name>`

Creating and Applying a VoIP Configuration File

The voice configuration characteristics of P-Series ONT devices and E-Series access devices have a common set of configuration capabilities that can be provisioned through a VoIP configuration file.

Creating VoIP configuration files

The VoIP 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 system.

Retrieving and applying VoIP configuration files

This topic describes how to retrieve, and apply VoIP configuration files. This topic also shows how to remove VoIP configuration files from the system.

Using VoIP Configuration Files Overview

1. **Retrieve** the configuration file from an external server that is reachable by the system and assign it a particular configuration name (voip-1 through voip-4), 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 configuration file to the VDSL2 card, causing the card to reset and apply the configuration from the file to the provisioning of all xDSL ports on that card.
3. **Remove** the old VoIP configuration file from the VDSL2 flash memory.

Note: The configuration name specified in the Retrieve/Apply DSL Configuration File operation must match the configuration file instance specified in the IP Host. A configuration name is specified by instance number alone or by “voip-<1 through 4>.”

If you retrieve another VoIP configuration file version to the same VDSL2, there will be two configuration files present in the VDSL2 card’s flash:

- One that is currently in use by the VDSL2 card (applied file)
- One that is a newly retrieved file (not-applied file).
- Once you apply the newly retrieved file to a VDSL2 card, the previously applied version of the configuration file (now obsolete) will be deleted from flash.

Note: From CMS, you can schedule the above tasks to be performed on multiple nodes that span multiple regions or network groups. The Retrieve scheduled task has an option of "Apply" (default is unselected), to push the retrieved configuration file to VDLS2 cards and trigger a reset for the cards to start using the provisioning indicated in the file.

Configuration guidelines

- The Configuration File Instance specified in the VDSL2 card IP Host for voice services, indicates which of the four configuration names will be used by that card “voip-<1 through 4>.”
- When a VoIP configuration file is applied to a VDSL2 card, the file contents are applied to all xDSL ports on that card.
- A given node will accept up to four different VoIP configuration files.
- The configuration file has a version number string embedded into is as an XML comment.
- A VoIP configuration file is not bundled with a system software release, so upon a system upgrade, the existing configuration file is carried over.
- The same configuration file can be used for VoIP on both ONTs and VDSL cards.

Parameters when managing configuration files

Parameter	Description
Source FTP Server*	The IP address of your FTP server.
Source User*	Username required by your FTP server.
Source Password*	Password required by your FTP server.
Source File Path*	File path of the configuration file on your FTP server.
Config Name*	A configuration name instance is mapped to a retrieved DSL Configuration file, whose contents will be applied to all DSL ports on the VDSL2 card. Each VDSL2 card supports 2 VoIP configuration files (1 applied and 1 not-applied) for a total of 4 VoIP configuration files per system. The configuration name specified in the Retrieve/Apply DSL Configuration File operation must match the configuration file instance specified in the IP Host. A configuration name is specified by instance number alone or by "voip-<1 through 4>." • voip-1 - Instance for VOIP (1) • voip-2 - Instance for VOIP (2) • voip-3 - Instance for VOIP (3) • voip-4 - Instance for VOIP (4)
Version*	The version number of the DSL configuration file <x.x.x.x>. Typically, the version string is the first line in the configuration file. **Note:** The VDSL2 card will not download a new configuration file, unless the version number in the file is different from what is in the file currently applied.
Force	Select the Force box to force the system to retrieve and name the file as specified.

*Required field

To manage configuration files

1. On the Navigation Tree, click a VDSL2 card.
2. In the Work Area, click **DSL Configuration** > **Action**.
3. Select one of the actions, described in the table above:
   - **Retrieve**
   - **Apply**
   - **Remove**
4. Click the appropriate confirmation to invoke the action.

For CLI:

```
retrieve dsl-config
apply dsl-config
remove dsl-config
cancel dsl-config
reset card
```

Creating a TDM Gateway Profile

This topic describes how to create a TDM voice gateway profile used by the TDM voice gateway service on a VDSL2 voice port. The profile is referenced when you provision the service on a voice (POTS) port.

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.

Parameters

You can provision the following parameters for TDM Gateway voice profiles:

Parameter	Description	Valid Options
Name*	Name of the TDM gateway profile.	Any established TDM Gateway profile
Server IP Address*	The C7 SIPVCG IP address for call signaling.	IP address in dot notation
Packetization Rate	Packetization rate in msec.	10ms ‡, 20ms

Parameter	Description	Valid Options
RTP DSCP	DSCP value for RTP packets.	• **0-63** - DSCP for RTP packets • **cs0** - DSCP CS0 (0) • **cs1** - DSCP CS1 (8) • **af11** - DSCP AF11 (10) • **af12** - DSCP AF12 (12) • **af13** - DSCP AF13 (14) • **cs2** - DSCP CS2 (16) • **af21** - DSCP AF21 (18) • **af22** - DSCP AF22 (20) • **af23** - DSCP AF23 (22) • **cs3** - DSCP CS3 (24) • **af31** - DSCP AF31 (26) • **af32** - DSCP AF32 (28) • **af33** - DSCP AF32 (30) • **cs4** - DSCP CS4 (32) • **af41** - DSCP AF41 (34) • **af42** - DSCP AF41 (36) • **af43** - DSCP AF43 (38) • **cs5** - DSCP CS5 (40) • **ef** - DSCP EF (46) ‡ • **cs6** - DSCP CS6 (48) • **cs7** - DSCP CS7 (56)
RTP Ethernet QoS	Ethernet QoS for RTP packets.	0-7 6 ‡
DHCP Filter	Whether to use the server IP address as the DHCP offer filter. When enabled, if the associated IP-host is using DHCP, then the DHCP offer from the specific server will be accepted.	enabled (selected), disabled (unselected) ‡

*Required field
‡ Default

To create a TDM voice gateway profile

1. Access the profile page:
   - From CMS:
     - On the Navigation Tree, click **CMS**.
     - In the Work Area, click **Profile** > **E3-48C/E5-48/E7/ONT** > **Profile** > **Service** > **TDM GW**.
   - Locally on the E-Series:
     - On theNavigation Tree, click the unit.
     - In the Work Area, click **Profiles** > **Service** > **TDM GW** > **Profiles**.
2. In the menu, click **Create**.
3. Reference the table above to configure the parameters.
4. Click **Create** to save the profile.

Apply a TDM gateway profile when you create a TDM voice service on an ONT POTS port.

For CLI:

- `create tdm-gw-profile <p-name> [server-ip]`
- `delete tdm-gw-profile [p-name]`
- `set tdm-gw-profile <p-name> [server-ip]`
- `show tdm-gw-profile [p-name] [services]`

Creating a TDM Gateway Service Group

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 E-Series. 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.

This topic describes how to create a TDM Gateway Service Group that defines the C7 network and interface groups used for delivering the TDM Gateway service, allowing CMS to automatically create and validate C7 gateway cross-connects.

Note: The E7 GPON only supports a single VoIP Interface Group for 700GE, 700GX and 760GX, although the GX ONTs support multiple VoIP IP Hosts.

Parameters

You can provision the following parameters for TDM Gateway service groups:

Parameter	Description	Valid Options
ID*	Index of the TDM Gateway service group.	Any value between 1-1000, inclusively.
Description	Unique name that indicates the entire description of the C7 network and Calix C7 gateway interface group used for this TDM service group. This value is auto-filled once the network name and interface groups are selected.	String text If left blank, the system automatically enters the combination of names for the selected C7 network and interface groups.
C7 Network Name*	Name of the C7 network where the interface groups have been created for C7 TDM gateway.	Any available configured C7 network where the required interface groups exist.
GR303/GR8 IG*	GR-303/GR-8 Interface Group (IG) for CMS to use for automatically creating the cross-connects.	Any available GR-303/GR-8 IG that is established on the selected C7 Network.
SIPVCG IG*	SIP Voice Concentration Group Interface Group (IG) for CMS to use for automatically creating the cross-connects.	Any available SIP VCG IG that is established on the selected C7 Network.
OSMINE-Compliant Northbound Interface use only Parameters applicable for E7 GPON and xDSL ports when the OSMINE-compliant Northbound Interface has been installed with CMS. If you are defining end-to-end E7 TDM voice service provisioning, enter the applicable information in the fields.		
Enabled	Whether the service group is enabled for use.	Y ‡ = enabled, N = disabled

*Required field
‡ Default

To create a CMS TDM gateway service group

1. On the Navigation Tree, click **CMS**.
2. Click **System** > **TDM Service Group**.
3. In the menu, click **Create**.
4. Reference the table above to configure the parameters.
5. Click **Create** to save the service group.

Apply a TDM gateway service group when you create a TDM voice service on an E7 system via CMS.

Creating an H.248 Gateway Profile

This topic describes how to create a profile that specifies the H.248 gateway properties for the VDSL2 H.248 gateway services.

Note: The 'Primary Gateway Controller' and 'Secondary Gateway Controller' parameters must be configured with a fully qualified domain name, when the IP host for H.248 services is provisioned using Fully Qualify Domain Name (FQDN).

Parameters

You can provision the following parameters for H.248 gateway profiles:

Parameter	Description	Valid Options
Name*	Name of the H.248 gateway profile.	A 31-character text string
Base RTP Port	Base port number for RTP packets.	49152 ‡ to 65535
Primary Gateway Controller	IP address or hostname of the primary H.248 gateway controller, that is, the softswitch.	text string
Primary Switch Type	Type of voice soft switch. • h248-ansi-generic - softswitch type set to H.248 ANSI Generic • none - softswitch type not specified • cs-2000 - Nortel Communication Server 2000 • cs-1500 - Nortel Communication Server 1500 • metaswitch - Metaswitch softswitches • sonus - Sonus softswitches • genband-g2 - GENBAND G2 Compact Gateway • genband-g9 - GENBAND G9 Converged Gateway • taqua - Taqua softswitches • tss - Ericsson softswitches (xDSL only)	cs-2000, cs-1500, metaswitch, sonus, genband-g2, genband-g9, taqua, tss, none ‡
Secondary Gateway Controller	IP address or hostname of the secondary H.248 gateway controller, that is, the softswitch.	text string

Parameter	Description	Valid Options
Secondary Switch Type	Type of voice soft switch. • h248-ansi-generic - softswitch type set to H.248 ANSI Generic • none - softswitch type not specified • cs-2000 - Nortel Communication Server 2000 • cs-1500 - Nortel Communication Server 1500 • metaswitch - Metaswitch softswitches • sonus - Sonus softswitches • genband-g2 - GENBAND G2 Compact Gateway • genband-g9 - GENBAND G9 Converged Gateway • taqua - Taqua softswitches • tss - Ericsson softswitches (xDSL only)	cs-2000, cs-1500, metaswitch, sonus, genband-g2, genband-g9, taqua, tss, none ‡
Termination Prefix	Prefix string to use for terminations.	TP ‡, text string
Ephemeral Termination ID	ID to use for ephemeral terminations.	RTP ‡, text string

ESA Mode	Whether to enable the ESA mode.	selected = enabled
		unselected = disabled ‡

RTP Codec	RTP code to use.	u-law ‡, a-law
Packet Rate	Packet rate (packets/msec).	10ms ‡, 20ms
Country Code (xDSL only)	E.164 Country code designator (Protocol Country Variant profile). This attribute specifies the country code where the service is being deployed. This code selects country specific tone settings, line interfaces, line levels and line frequencies. Currently supported values include the following: • North America: 1 (xDSL only) • Italy: 39 (xDSL only) • Switzerland: 41 (xDSL only) • United Kingdom: 44 (xDSL only) • Sweden: 46 (xDSL only) • Poland: 48 (xDSL only) • Brazil: 55 (xDSL only) • Australia: 61 (xDSL only) • New Zealand: 64 (xDSL only) • Algeria: 213 (xDSL only) • Ukraine: 380 (xDSL only) • ETSI: 9000 (xDSL only)	1‡ -9999

*Required field
‡Default

To create an H.248 Gateway profile

1. Access the profile page:
   - From CMS:
     - On the Navigation Tree, click **CMS**.
     - In the Work Area, click **Profile** > **E5-48/E3-48C/E7/ONT** > **Profile** > **Service** > **H.248 GW**.

- Locally via the EWI:
  - On the Navigation Tree, click the **E7, E5-48 or E3-48C**.
  - In the Work Area, click **Profiles** > **Service** > **H.248 GW** > **Profiles**.

2. In the menu, click **Create**.
3. Reference the table above to configure the parameters.
4. Click **Create**.

For CLI:

```
create h248-gw-profile <p-name> [rtp-base-port|pri-gw-controller|pri-switch-
type|sec-gw-controller|sec-switch-type|term-prefix|ephemeral-term-id|esa-
mode|rtp-codec|packet-rate|country-code]
```

Creating an H.248 Gateway

This topic describes how to create a profile that specifies the H.248 gateway properties for the VDSL2 H.248 gateway services.

Parameters

You can provision the following parameters for H.248 gateway:

Parameter	Description	Valid Options
Name*	Name of the H.248 gateway.	A 31-character text string
Admin State	Administrative status for the gateway.	enabled ‡, disabled
H.248 Gateway Profile*	Name of the H.248 gateway profile.	text string
IP Host*	Name of the line card IP Host.	text string

*Required field
‡Default

To create an H.248 Gateway

1. In the Navigation Tree, click the **E-Series** on which to create an H.248 Gateway.
2. In the Work Area, click **H.248 Gateway**, and then click **Create**.
3. Reference the table above to configure the parameters.
4. Click **Create**.

For CLI:

```
create h248-gw <shelf/card/gw-name> h248-gw-profile <p-name> ip-host <h-
name> [admin-state]
remove h248-gw-svc from-pots-port <vdsl-port>
```

Retrieving and Applying a VoIP Coefficient File

This topic describes how to retrieve, and apply VoIP coefficient files for international POTS support. This topic also shows how to remove VoIP coefficient files from the system.

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.

Configuration guidelines

- When a coefficient file is applied to a VDSL2 card, the file contents are applied to all xDSL ports on that card.
- A given node will accept up to two different coefficient files.
- The coefficient file has a unique differentiator, therefore it is not necessary to have a version number associated with the file.
- A coefficient file is bundled with a system software release, so upon a system upgrade, any retrieved coefficient files on the VDSL2 cards will NOT be carried over. Assuming that the updated system software has all of the necessary coefficient settings built-in, the previously applied coefficient files would be obsolete.

Parameters

You can provision the following parameters when managing coefficient files:

Parameter	Description
Source FTP Server*	The IP address of your FTP server.
Source User*	Username required by your FTP server.
Source Password*	Password required by your FTP server.
Source File Path*	File path of the configuration file on your FTP server.
Force	Select the Force box to force the system to retrieve the file as specified.

*Required field

To manage coefficient files

1. On the Navigation Tree, click a VDSL2 card.
2. In the Work Area, click **Action** > **DSL Coefficient**.
3. Select one of the actions, described in the table above:
    - **Retrieve**
    - **Apply**
    - **Remove**
4. Click the appropriate confirmation to invoke the action.

For CLI:

```
retrieve dsl-coefficient
apply dsl-coefficient
remove dsl-coefficient
cancel dsl-coefficient
show dsl-coefficient
reset card
```

Step 3. Configure Subscriber Services

This section describes how to configure various VDSL2 services, using the previously-created profiles and configured uplinks.

Topics Covered

This section covers the following **topics in bold** that are part of the overall VDSL2 services configuration process:

1. Configure network uplinks for VDSL2 services
2. Creating system profiles that support VDSL2 applications
3. **Configure subscriber services**
    - **Configuring Data Services**
    - **Configuring Video Services**
    - **Configuring Voice Services**
    - **Configuring T1/E1 PWE3 Services**

Configuring Data Services

This section describes how to create data services on xDSL ports.

Topics Covered

This section covers the following **topics in bold** that are part of the overall E-Series VDSL2 services configuration process:

1. Configure network uplinks for VDSL2 services.
2. Create profiles that support VDSL2 applications.
3. **Configure data subscriber services.**
    - **An overview of the VDSL2 data services provisioning**
    - **Configuring an xDSL port and interface for service**
    - **Configuring a data service**

Overview: Configuring xDSL Data Services

The E-Series provides data interconnection between Internet service providers and subscribers via the VDSL2 card ports. Service profiles and service tag actions in the E-Series help provide tiered service offerings by specifying the priority and marking of packets from the subscriber port into the service VLAN.

The Calix E-Series supports the following data service models over IP on VDSL2 card 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.
- **VLAN-per-Port:** where each subscriber port is assigned a dedicated VLAN. This model is often referred to as 1:1.
- **VLAN Stacking, Q-in-Q:** where a transparent LAN service on each port is configured using VLAN double tagging.

Before starting

The port service provisioning options allow you to activate services on an xDSL port and apply previously-created profiles. Before starting the data services configuration process, ensure the following steps of the turn-up process have been completed:

1. **The network uplinks for VDSL2 services are configured.**
   - Ethernet port interfaces
   - *Ethernet ports* (on page 21)
   - Service VLANs
   - *VLAN memberships* (on page 64)

     Uplink interfaces (or ERPS domain, if the uplink resides on a different shelf) must be added to the VLAN membership.
2. **The necessary system profiles that support VDSL2 data service applications are created.**
   - *Ethernet bandwidth profile* (on page 105)
   - *Rules to classify traffic* (on page 71)
     - *Service match list* (on page 78)
     - *Service tag action* (on page 82)

     The data traffic should be set where the service tag action specifies a priority of a P-bit set to 0, 1, 2, 3 or the correlating DSCP value set to a value of 0, CS1 & AF11-13 (8-14), CS2 & AF21-23 (16-22), CS3 & AF31-33 (24-30).

   Optional profiles:

   - *DSL Port Template* (on page 129)
   - *Layer 3 priority map* (on page 96)
   - *Security Profile* (on page 143)
   - *Bonding groups, if required* (on page 150)

Note: See *Calix xDSL Best Practices* for a description of the physical layer factors that may directly influence the quality of data and video services delivery, and recommendations for the best practices necessary to achieve optimal results.

Step 3. Configuring an xDSL Data Service

After completing the steps outlined above, continue with the following topics in this section to configure an xDSL data service:

- *Configure the xDSL port.* (on page 179)
- *Configure the xDSL port associated interface.* (on page 182)
- *Add the data service to the xDSL port.* (on page 183)

Configuring an xDSL Port for Data Services

This topic describes how to configure various parameters that are more likely to be modified on an xDSL port for data service. You can also apply a previously-created DSL template to multiple ports for ease of configuration. See *Creating and Applying a DSL Port Template* (on page 129).

Applying multiple DSL templates to a single port

Each DSL template may specify some or all of the xDSL port parameters. When a DSL template is applied to an xDSL port, the parameter values that are specified in that template are copied into the port object.

If another template is subsequently applied to the same xDSL port, the parameter values specified in that template are copied into the port object.

- If the same parameters are specified in both templates, the values defined in the last-applied template are applied to the xDSL port.
- If parameters specified in the first-applied template are not specified in the last-applied template, the values are retained from the first-applied template for those parameters.

Recommended parameter values for HSI:

Basic DSL Port Parameters	
Path Latency	Specifies the operating mode of the primary channel. • fast – for delay sensitive applications like voice and online gaming. Specifies a minimum of 4 ms delay. • interleaved – interleaves DSL frames to optimize error protection in the presence of impulse noise sources that are common to DSL. Specifies a delay greater than or equal to 5 ms. For HSI with no IPTV, Calix recommends the "fast" setting in both upstream and downstream directions.
Max Rate	Defines maximum downstream or upstream of rate for xDSL port (Kb/s, or use "m" suffix for Mb/s). For HSI, Calix recommends a value equal to or slightly greater than the minimal bitrate required to support the corresponding service.
Max SNR	Defines the maximum downstream or upstream signal-to-noise ratio (SNR) margin (dB, in 0.1 dB increments) defines the amount of margin above target SNR that must be present before power cutback occurs. For HSI, Calix recommends an additional 5 dB of margin above target SNR before power cutback is allowed. This corresponds to a maximum SNR margin of 10 dB for supporting HSI.
Advanced DSL Port Parameters	
DS or US downshift rate adaptation margin	Defines an SNR margin below the Target SNR Margin that triggers a bitrate "downshift" (dB, in 0.1 dB increments). For the decrease to occur, the noise margin must stay below this value for the time specified in upstream or downstream Downshift Rate Adaptation Time. Applies only when dynamic rate adaptation is specified. • For ADSL2 applications, Calix recommends a value of 3 dB below target for both upstream and downstream directions. • For VDSL2 applications, Calix recommends a value of 5 dB below target for both upstream and downstream directions.
DS or US upshift rate adaption margin	Defines an SNR margin above the Target SNR Margin which, when exceeded, initiates a bitrate "upshift" (dB, in 0.1 dB increments). For the increase to occur, the noise margin must stay above this value for the time specified in upstream or downstream Upshift Rate Adaptation Time. Applies only when dynamic rate adaptation is specified. • For IPTV, Calix recommends a value of 3 dB above target for both upstream and downstream directions. • For HSI, Calix recommends a value of 1 dB above target for both upstream and downstream directions.

If necessary, see the complete set of parameter descriptions in the following topics:

- *Basic xDSL Port Parameters* (on page 131)
- *Advanced xDSL Port Parameters* (on page 135)
- *Power Spectral Density (PSD) xDSL Port Parameters* (on page 139)

To apply a DSL Template to a port

You can modify multiple parameters with a single action of applying a DSL template to the port, replacing the DSL port settings.

1. On the Navigation Tree, click the xDSL port on which you want to apply the DSL template.
2. Click **Provisioning** > **Basic** > **Action** > **Apply Template**.
3. In the Template list, select from the list of previously-created DSL templates.
4. Click **Apply Template**.

For CLI:

```
apply dsl-template <t-name> to-dsl-port <port>
```

To modify the xDSL parameters to recommended values

1. On the Navigation Tree, click the xDSL port of which you want to modify parameters.
2. In the Workarea, click **Port** > **Provisioning** > **Basic**.
3. Refer to the table above for the parameter descriptions and options and set the values as necessary.
4. Click **Apply**.
5. In the Workarea, click **Port** > **Provisioning** > **Advanced**.
6. Refer to the table above for the parameter descriptions and options and set the values as necessary.
7. Click **Apply**.

For CLI:

```
set dsl-port <port-id> advanced
set dsl-port <port-id> basic
set dsl-port <port-id> psd
show dsl-port [port]
[advanced|all|basic|inventory|psd|status|subcarriers]
```

Configuring an xDSL Port Associated Interface

This topic describes how to configure an xDSL port associated interface for service.

xDSL port associated interface parameters

You can provision the following parameters for an xDSL port associated interface:

Parameter	Description	Valid Options
Admin State	Admin state of the port, select whether to enable the interface.	enabled ‡, disabled
Description	Description of the interface for easy identification later in a search	31 character text string
EtherType	The Ethertype indicates the protocol being transported in the Ethernet frame. The VLAN tagged frames are identified as having a tag by utilizing the Ethertype field. • 0x8100 - IEEE 802.1Q-tagged (default) • 0x88a8 - IEEE 802.3ad provider bridging • 0x9100 - Q-in-Q (double tagged)	0x8100 ‡ 0x88a8 0x9100
LACP Tunnel	When set to enabled, LACP protocol packets are forwarded to the their destination as ordinary multicast packets. This function should only be used when configuring TLAN service.	enabled, disabled ‡
IGMP Immediate Leave	Enable or disable IGMP immediate leave. When enabled, queries are omitted that would identify if there are other hosts interested in the multicast group. If the interface is connected to a single subscriber device then the system does not need to query the interface and it can terminate the delivery of the Multicast stream right away. This mode of operation is desired when the interface is connected to a Residential Gateway that provides IGMP proxy functionality for all devices behind it.	enabled, disabled, use-vlan-setting ‡
Subscriber ID	Subscriber ID information, such as phone number, address, or account number.	47 character text string
Security Profile*	Name of Security Profile to apply to the interface that can limit DHCP leases, limit the rate of Layer 2 broadcast traffic allowed on the interface, and indicate whether to pass the L2CP protocol frames.	any existing security profile
DSCP/IP Precedence Profile	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.	DscpMap: access ‡, IpPrecMap: access, any previously-created profile
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 ‡ authorized unauthorized

*Required field
‡ Default

To configure an xDSL port associated interface

1. On the Navigation Tree, click the xDSL port of interest.
2. Click **Associated Interface > Provisioning.**
3. Reference the table above to configure the parameters.
4. In the menu, click **Apply**.

For CLI:

```
set interface <dsl-port> [eth-svc|description|subscriber-
id|immediate-leave|igmp-max-rate|igmp-max-groups|dscp-p-bit-map|ip-
prec-p-bit-map|lacp-tunnel|eth-sec-profile|force-dot1x|admin-state]
```

Related topics

- *Mapping Layer 3 Priority Values to P-Bits* (on page 96)
- *Creating an Ethernet Security Profile* (on page 143)
- Creating VLANs
- *Creating an IGMP profile* (on page 119)

Creating a Data Service

This section describes how to create a data (Internet) service on an xDSL port, using the Form option on the xDSL port services menu. You can create six (6) Ethernet Services on a single xDSL port.

- The Form option allows you to initially provision multiple services on the xDSL port in one view.
- The Table option allows you to provision additional single services, delete single services, or update existing services.

Profiles referenced in service provisioning

When provisioning subscriber services on a VDSL2 card xDSL port, you set up information that is common to multiple subscribers and maintained in profiles. Repeating the similar task of subscriber service activation then consists of referencing the existing profiles and sometimes providing subscriber-specific information to complement the profiles.

Both service provision options allow you to activate services on an xDSL port and apply the following previously-created profiles:

- Ethernet Bandwidth Profile
- Service Tag Actions

Note: When a profile is changed, all subscribers using that profile are affected. Profiles cannot be deleted while they are in use. Profiles can also be reloaded for a specific VDSL2 xDSL port from the Port Services Table by clicking **Refresh**.

Information you need

At a minimum, you should have the following information on hand to configure xDSL data service:

- Ethernet bandwidth profile to use on the xDSL port
- Service tag action ID to use for the data service
- Tag ID(s) to apply, if "specify in service" was selected in the service tag action

Parameters

You can provision the following parameters for a data service on an xDSL port:

Parameter	Description	Valid Options
Subscriber ID	Subscriber ID information, such as phone number, or account number.	text string (blank) ‡
Description	Optional description field for the port, subscriber address, name, or service.	text string (blank) ‡
BW Profile*	Name of bandwidth profile to apply to the service.	text string
Service Tag Action*	Name of service tag action to apply to the service. If the "specified in service" option was selected for either the Outer Tag or the Inner Tag in the service tag action, enter the applicable VLAN IDs.	text string
Outer Tag	Specifies the service VLAN ID, if the Service Tag Action references "Specified in Service" for the Outer Tag.	none, any VLAN created on the system
Inner Tag	Specifies the customer VLAN ID, if the Service Tag Action references "Specified in Service" for the Inner Tag.	none, any VLAN created on the system

*Required fields
‡ Default

To create a data service on an xDSL port

1. On the Navigation Tree, click the xDSL port or bonding group on which you want to add a data service.
2. In the Work Area, click **Port** > **Services** > **Form**.

   Note: To create a service on an xDSL Bonded interface, on the Navigation Tree, click the xDSL Group under the VDSL2 card, and then in the Work Area, click **Services** > **Form**.

3. Reference the table above to configure the parameters.
4. In the menu, click **Apply**.
5. Verify your modem train rate using the **Port** > **Status** tab.

For CLI:

```
add eth-svc <service name> to-interface <dsl-port> bw-profile <bw-profile-name> svc-tag-action [outer-vlan <vlan ID> inner-vlan <vlan ID> mcast-profile <profile name> description <service> admin-state [enabled|disabled]]

add eth-svc <service name> to-dsl-bond-interface <b-intfc> bw-profile <bw-profile-name> svc-tag-action [outer-vlan <vlan ID> inner-vlan <vlan ID> mcast-profile <profile name> description <service> admin-state [enabled|disabled]]
```

Configuring PPPoA/PPPoE Operation for a Data Service

The E-Series supports 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 E-Series 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 E-Series 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 E-Series 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.

This topic shows you how to configure PPPoE operation on an xDSL port that has a provisioned data service.

See *Creating a PPPoE Profile* (on page 146) for instructions, if necessary.

Configuration guidelines:

- The E-Series uplink toward the router or PPPoE Server must be on an interface set to the mode of "Trunk."
- 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 port interface to either a common service VLAN, or to a 1:1 VLAN associated with that customer. At the aggregating Trunk interface northbound of the E-Series, the PPPoE traffic must adhere to a specific VLAN membership on the Trunk interface.
- Sessions may be manually terminated.
- In ATM mode on the VDSL2 card, the E-Series supports PPPoA on the client side.

- 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 port at once.
- The PPPoE profile must be set to Auto mode for the InterWorking function (PPPoA-to-PPPoE conversion), and the match rule of the Ethernet service must match on untagged packets.
- Only one PPPoE session is supported per xDSL port. If a second new session is initiated, the original session tears down and the new session is allowed.
- Up to one service on each xDSL interface (including bonded-links) may resolve to a VLAN with an associated PPPoE profile.
- The PPPoE profile is applied to a service VLAN where simultaneous operation of DHCP Snooping and PPPoE are not supported. When a PPPoE profile is selected for a data service VLAN, the DHCP features are disabled.
- 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).
- When a PPPoE snoop is NOT applied to a VLAN in the E7, PPP traffic is forwarded as part of the traffic stream, but, the Intermediate Agent and security aspects of the feature are not enabled in this mode.
- Each xDSL subscriber can only have up to 1 Ethernet service with PPPoE enabled.
- To disable downstream broadcast traffic on the E-Series, enable a PPPoE profile on the service VLAN to only forward PADI broadcast packets upstream to the BRAS port once it detects which one it is.

To configure PPPoA/PPPoE operation

1. Create a service VLAN and do the following:
    - Set the VLAN to reference the PPPoE profile.
    - Ensure that the DHCP Snoop is disabled (unselected).
    - Add the VLAN Member to the uplink interface configured in Step 1.
2. Create a match list and service tag action to match the data traffic at the subscriber port and mark it with the service VLAN ID.
3. Create a data service on the subscriber port, associating the service tag action created in Step 2.

To view the PPPoE sessions

1. On the Navigation Tree, click the xDSL port or bonding group of interest.
2. In the Workarea, click **PPPoE** > **Sessions**.

For CLI:

```
show pppoe sessions [detail]
show pppoe sessions id <ses-id> [detail]
show pppoe sessions mac <m-add> [detail|id]
show pppoe sessions interface <intfc-name> [detail|vlan <vlan-id>]
show pppoe sessions dsl-bond-interface <intfc-name> [detail|vlan
<vlan-i>]
```

To delete a PPPoE session

1. On the Navigation Tree, click the xDSL port or bonding group of interest.
2. In the Workarea, click **PPPoE** > **Sessions**.
3. Click on the table row that indicates a PPPoE session to select it.
4. In the menu, click **Delete**.

For CLI:

```
delete pppoe sessions interface <intfc-name> id <ses-id>
delete pppoe sessions dsl-bond-interface <intfc-name> id <ses-id>
```

Adding Static IP Host Addresses and Subnets to xDSL Services

This topic shows you how to configure a static Ethernet service IP address or subnet, to associate with an xDSL port service or an xDSL bonding group service.

Configuration guidelines

- To identify a specific IP host for the purposes of securing the subscriber's traffic, the IP address must be provided, and optionally, the MAC address information.
- To identify an IP subnet, a MAC address is not required. The smallest subnet that can be provisioned in the system is 255.255.255.252, or /30.
- A given data service on a VDSL2 card represents exactly one subnet. To support a multi-homed router off an xDSL port, multiple data services must be provisioned to represent multiple subnets.
- The gateway address and subtending IP addresses must belong to the same subnet as indicated by the mask.
- The static IP address must not be the same as the gateway address.
- Duplicate static IP addresses are not allowed in the system.
- IP Source Verification for Static IP hosts requires MAC FF be enabled on the VLAN.
- If a provisioned static IP address conflicts with a DHCP learned IP address, the system will accept the static IP address and reject/delete the learned IP address.

- 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).
- 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 an xDSL Port
  - Static IP Subnet: Checks individual host IP addresses to the subnet and binds the subnet to an xDSL port

Before starting

Before starting this procedure, you will need to provision a service on the xDSL port where you want to configure the static IP host address or subnet.

Parameters

You can provision the following parameters for a static IP host address and subnet:

Parameter	Description	Valid Options
Service*	Name of the service to associate with an IP address.	Any established service name
Type*	IP address or hostname of the primary SIP configuration server.	IP address, hostname
Static IP Host Address		
IP Address*	Static IP address for the ONT port service.	IP address in dotted quad" format: "192.168.1.100". Alternately, "none" can be used to reset the value to "0.0.0.0"
MAC Address	L2 MAC address associated with the IP address.	Six hexadecimal digits in the range 0-FF, optionally separated by colons
Gateway*	Address of the default gateway for subtending static IP address objects.	IP address in dotted quad" format: "192.168.1.100". Alternately, "none" can be used to reset the value to "0.0.0.0"
Net Mask*	IP Netmask for subscriber host.	IP address in dotted quad" format: "192.168.1.100". Alternately, "none" can be used to reset the value to "0.0.0.0"

Parameter	Description	Valid Options
Static Subnet		
Subnet Address*	Subscriber IP address subnet.	IP address in dotted quad" format: "192.168.1.100". Alternately, "none" can be used to reset the value to "0.0.0.0"
Gateway*	Address of the default gateway for subtending static IP address objects.	IP address in dotted quad" format: "192.168.1.100". Alternately, "none" can be used to reset the value to "0.0.0.0"
Subnet Mask*	Subnet mask for subtending static IP address objects.	IP address in dotted quad" format: "192.168.1.100". Alternately, "none" can be used to reset the value to "0.0.0.0"

*Required fields

To add a static IP host address

1. On the Navigation Tree, click the xDSL port or xDSL bonding group on which the service is provisioned where you want to add a static IP host address or subnet.
2. Click **Static IP/Subnet**.
3. If there are multiple Ethernet services provisioned on the selected port or bonding group, at the top of the Work Area, click the drop-down list to select the service where you want to add a static IP host or subnet.
4. Click **Create** to open the Create Ethernet Service IP/Subnet dialog box.
5. Reference the table above to configure the parameters.
6. If you chose to add a subnet address, fill in the required addresses and subnet mask.
7. Click **Create**.

For CLI:

```
add static-ip-entry to-dsl-bond-interface <dsl-intfc> eth-svc <s-
name> type host ip <h-ip> netmask <n-ip> default-gw <gw-ip> [mac <m-
address>]

add static-ip-entry to-dsl-bond-interface <dsl-intfc> eth-svc <s-
name> type subnet ip <sub-ip> netmask <n-ip> default-gw <gw-ip>

add static-ip-entry to-interface <dsl-intfc> eth-svc <s-name> type
host ip <h-ip> netmask <n-ip> default-gw <gw-ip> [mac <m-address>]

add static-ip-entry to-interface <dsl-intfc> eth-svc <s-name> type
subnet ip <sub-ip> netmask <n-ip> default-gw <gw-ip>
```

Configuring IP Video Services

This section describes how to create video services on xDSL ports.

Topics Covered

This section covers the following **topics in bold** that are part of the overall VDSL2 services configuration process:

1. Configure network uplinks for VDSL2 services.
2. Create profiles that support VDSL2 applications.
3. **Configure video subscriber services.**
   - **An overview of the VDSL2 video services provisioning**
   - **Configuring an xDSL port and interface for service**
   - **Configuring a video service**

Overview: Configuring xDSL Video Services

The E-Series provides interconnection between video service providers and subscribers via the VDSL2 card ports. Service profiles and service tag actions in the E-Series help provide tiered service offerings by specifying the priority and marking of packets from the subscriber port into the service VLAN.

The Calix E-Series supports the following data service models over IP on 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.
- **VLAN-per-Port:** where each subscriber port is assigned a dedicated VLAN. This model is often referred to as 1:1.
- **VLAN Stacking, Q-in-Q:** where a transparent LAN service on each port is configured using VLAN double tagging.

Before starting

The port service provisioning options allow you to activate services on an xDSL port and apply previously-created profiles. Before starting the video services configuration process, ensure the following steps of the turn-up process have been completed:

1. **The network uplinks for VDSL2 services are configured.**
   - Ethernet port interfaces
   - *Ethernet ports* (on page 21)
   - Service VLANs
   - *VLAN memberships* (on page 64)

     Uplink interfaces (or ERPS domain, if the uplink resides on a different shelf) must be added to the VLAN membership.

2. **The necessary system profiles that support VDSL2 video service applications are created.**

   - *Ethernet bandwidth profile* (on page 105)
   - *Multicast profile* (on page 110)
   - *Rules to classify traffic* (on page 71)
     - *Service match list* (on page 78)
     - *Service tag action* (on page 82)

     The match list rules are for the Unicast VLAN, when configuring an MVR video service. The Multicast VLAN will follow the same match rule as setup for the Unicast stream when the service is provisioned on the ONT Ethernet port. The video traffic should be set where the service tag action specifies a priority of a P-bit set to 4, or the correlating DSCP value set to a value of CS4 & AF41-43 (32-38).

   Optional profiles:

   - *DSL Port Template* (on page 129)
   - *Layer 3 priority map* (on page 96)
   - *Security Profile* (on page 143)
   - *Bonding groups, if required* (on page 150)

Note: See *Calix xDSL Best Practices* for a description of the physical layer factors that may directly influence the quality of data and video services delivery, and recommendations for the best practices necessary to achieve optimal results.

Step 3. Configuring an xDSL Video Service

After completing the steps outlined above, continue with the following topics in this section to configure an xDSL video service:

- *Configure the xDSL port.* (on page 191)
- *Configure the xDSL port associated interface.* (on page 182)
- *Add the video service to the xDSL port.* (on page 196)

Configuring an xDSL Port for Video Services

This topic describes how to configure various parameters that are more likely to be modified on an xDSL port for video service. You can also apply a previously-created DSL template to multiple ports for ease of configuration. See *Creating and Applying a DSL Port Template* (on page 129).

Applying multiple DSL templates to a single port

Each DSL template may specify some or all of the xDSL port parameters. When a DSL template is applied to an xDSL port, the parameter values that are specified in that template are copied into the port object.

If another template is subsequently applied to the same xDSL port, the parameter values specified in that template are copied into the port object.

- If the same parameters are specified in both templates, the values defined in the last-applied template are applied to the xDSL port.
- If parameters specified in the first-applied template are not specified in the last-applied template, the values are retained from the first-applied template for those parameters.

Recommended parameter values

For IPTV, Calix recommends the following settings.

Basic DSL Port Parameters	
Path Latency	Specifies the operating mode of the primary channel. • fast – for delay sensitive applications like voice and online gaming. Specifies a minimum of 4 ms delay. • interleaved – interleaves DSL frames to optimize error protection in the presence of impulse noise sources that are common to DSL. Specifies a delay greater than or equal to 5 ms. For IPTV, Calix recommends the "interleaved" setting in the downstream direction. Latency is tunable when using the interleaved path, Calix recommends maximizing the downstream delay of 8 ms with MS Mediaroom, or 20 ms without. The default value of "interleaved" is recommended for supporting video services on the xDSL port.
Max Rate	Defines maximum downstream or upstream of rate for xDSL port (Kb/s, or use "m" suffix for Mb/s). The downstream rate should be sized to accommodate the maximum number of simultaneous multicast (live channels) and unicast (VOD, OTTV) video streams, as well as any guaranteed minimum rate for data (HSI).
Max SNR	Defines the maximum downstream or upstream signal-to-noise ratio (SNR) margin (dB, in 0.1 dB increments) defines the amount of margin above target SNR that must be present before power cutback occurs. For IPTV, Calix recommends enabling power cutback by setting maximum SNR margin to 16 dB for both upstream and downstream directions.

Target SNR	Defines the target downstream or upstream signal-to-noise ratio (SNR) margin (dB, in 0.1 dB increments) defines the SNR margin that must be available when the handshake process is determining the capability of each subcarrier. This has a direct impact on the attainable bitrate as ta higher margin forces few symbols per constellation. For IPTV, Calix recommends a downstream and upstream target SNR margin of at least 8 dB for DSL profiles supporting IPTV.
Advanced DSL Port Parameters	
DS or US downshift rate adaptation margin	Defines an SNR margin below the Target SNR Margin that triggers a bitrate "downshift" (dB, in 0.1 dB increments). For the decrease to occur, the noise margin must stay below this value for the time specified in upstream or downstream Downshift Rate Adaptation Time. Applies only when dynamic rate adaptation is specified. • For ADSL2 applications, Calix recommends a value of 3 dB below target for both upstream and downstream directions. • For VDSL2 applications, Calix recommends a value of 5 dB below target for both upstream and downstream directions.
DS or US upshift rate adaption margin	Defines an SNR margin above the Target SNR Margin which, when exceeded, initiates a bitrate "upshift" (dB, in 0.1 dB increments). For the increase to occur, the noise margin must stay above this value for the time specified in upstream or downstream Upshift Rate Adaptation Time. Applies only when dynamic rate adaptation is specified. • For IPTV, Calix recommends a value of 3 dB above target for both upstream and downstream directions. • For HSI, Calix recommends a value of 1 dB above target for both upstream and downstream directions.
Ds or US enhanced impulse noise protection mode	Used in conjunction with Retransmission which is a proprietary method for retransmitting data and requires a DSL modem that is Broadcom based and PhyR enabled. For IPTV, Calix recommends an Enhanced Impulse Noise Protection setting of "phyr" in the downstream direction.

If necessary, see the complete set of parameter descriptions in the following topics:

- *Basic xDSL Port Parameters* (on page 131)
- *Advanced xDSL Port Parameters* (on page 135)
- *Power Spectral Density (PSD) xDSL Port Parameters* (on page 139)

To apply a DSL Template to a port

You can modify multiple parameters with a single action of applying a DSL template to the port, replacing the DSL port settings.

1. On the Navigation Tree, click the xDSL port on which you want to apply the DSL template.
2. Click **Provisioning** > **Basic** > **Action** > **Apply Template**.
3. In the Template list, select from the list of previously-created DSL templates.
4. Click **Apply Template**.

For CLI:

```
apply dsl-template <t-name> to-dsl-port <port>
```

To modify the xDSL parameters to recommended values

1. On the Navigation Tree, click the xDSL port of which you want to modify parameters.
2. In the Workarea, click **Port** > **Provisioning** > **Basic**.
3. Refer to the table above for the parameter descriptions and options and set the values as necessary.
4. Click **Apply**.
5. In the Workarea, click **Port** > **Provisioning** > **Advanced**.
6. Refer to the table above for the parameter descriptions and options and set the values as necessary.
7. Click **Apply**.

For CLI:

```
set dsl-port <port-id> advanced
set dsl-port <port-id> basic
set dsl-port <port-id> psd
show dsl-port [port]
[advanced|all|basic|inventory|psd|status|subcarriers]
```

Configuring an xDSL Port Associated Interface

This topic describes how to configure an xDSL port associated interface for service.

xDSL port associated interface parameters

You can provision the following parameters for an xDSL port associated interface:

Parameter	Description	Valid Options
Admin State	Admin state of the port, select whether to enable the interface.	enabled ‡, disabled
Description	Description of the interface for easy identification later in a search	31 character text string
EtherType	The Ethertype indicates the protocol being transported in the Ethernet frame. The VLAN tagged frames are identified as having a tag by utilizing the Ethertype field. • 0x8100 - IEEE 802.1Q-tagged (default) • 0x88a8 - IEEE 802.3ad provider bridging • 0x9100 - Q-in-Q (double tagged)	0x8100 ‡ 0x88a8 0x9100
LACP Tunnel	When set to enabled, LACP protocol packets are forwarded to the their destination as ordinary multicast packets. This function should only be used when configuring TLAN service.	enabled, disabled ‡
IGMP Immediate Leave	Enable or disable IGMP immediate leave. When enabled, queries are omitted that would identify if there are other hosts interested in the multicast group. If the interface is connected to a single subscriber device then the system does not need to query the interface and it can terminate the delivery of the Multicast stream right away. This mode of operation is desired when the interface is connected to a Residential Gateway that provides IGMP proxy functionality for all devices behind it.	enabled, disabled, use-vlan-setting ‡
Subscriber ID	Subscriber ID information, such as phone number, address, or account number.	47 character text string
Security Profile*	Name of Security Profile to apply to the interface that can limit DHCP leases, limit the rate of Layer 2 broadcast traffic allowed on the interface, and indicate whether to pass the L2CP protocol frames.	any existing security profile
DSCP/IP Precedence Profile	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.	DscpMap: access ‡, IpPrecMap: access, any previously-created profile
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 ‡ authorized unauthorized

*Required field
‡ Default

To configure an xDSL port associated interface

1. On the Navigation Tree, click the xDSL port of interest.
2. Click **Associated Interface > Provisioning.**
3. Reference the table above to configure the parameters.
4. In the menu, click **Apply**.

For CLI:

```
set interface <dsl-port> [eth-svc|description|subscriber-
id|immediate-leave|igmp-max-rate|igmp-max-groups|dscp-p-bit-map|ip-
prec-p-bit-map|lacp-tunnel|eth-sec-profile|force-dot1x|admin-state]
```

Related topics

- *Mapping Layer 3 Priority Values to P-Bits* (on page 96)
- *Creating an Ethernet Security Profile* (on page 143)

- Creating VLANs
- *Creating an IGMP profile* (on page 119)

Creating a Video Service

This section describes how to create a video service on an xDSL port, using the Form option on the xDSL port services menu. You can create six (6) Ethernet Services on a single xDSL port.

- The Form option allows you to initially provision multiple services on the xDSL port in one view.
- The Table option allows you to provision additional single services, delete single services, or update existing services.

Profiles referenced in service provisioning

When provisioning subscriber services on a VDSL2 card xDSL port, you set up information that is common to multiple subscribers and maintained in profiles. Repeating the similar task of subscriber service activation then consists of referencing the existing profiles and sometimes providing subscriber-specific information to complement the profiles.

Both service provision options allow you to activate services on an xDSL port and apply the following previously-created profiles:

- Multicast Profile
- Ethernet Bandwidth Profile
- Service Tag Action

Note: When a profile is changed, all subscribers using that profile are affected. Profiles cannot be deleted while they are in use. Profiles can also be reloaded for a specific VDSL2 xDSL port from the Port Services Form or Table by clicking **Refresh**.

Information you need

At a minimum, you should have the following information on hand to configure xDSL video service:

- Multicast profile to use for the video service
- Ethernet bandwidth profile to use on the xDSL port
- Service tag action ID to use for the video service
- Tag ID(s) to apply, if "specify in service" was selected in the service tag action
- Multicast white list (optionally)

Parameters

You can provision the following parameters for a video service on an xDSL port:

Parameter	Description	Valid Options
Subscriber ID	Subscriber ID information, such as phone number, or account number.	text string (blank) ‡
Description	Optional description field for the port, subscriber address, name, or service.	text string (blank) ‡
BW Profile*	Name of bandwidth profile to apply to the service.	text string Any established Ethernet bandwidth profile
Multicast Profile*	Name of Ethernet multicast profile to apply to the port. **Note:** To create video service with MVR, this profile must reference an MVR profile where the MVR VLAN ID is that of the Multicast VLAN. The MVR Profile must be enabled for correct configuration.	text string Any established multicast profile
Multicast White List	Name of optional multicast white list(s) to apply to the video service that defines global allowable multicast IP ranges. Select any single name of a previously created list, or hold down the Ctrl key to select or deselect multiple items. **Note:** A multicast white list can be applied at the time of creating the service, or later to an existing service as follows: • In the Navigation Tree, click **XDSL#**. • In the Work Area, click **Port** > **Multicast White List** > **Create**, and then select the service and multicast white list(s).	Any existing multicast white list
Service Tag Action*	Name of service tag action to use for the service. • For video service with MVR, select the service tag action that references the Unicast VLAN and has been associated with the matched traffic rule of the Unicast VLAN stream. • If the "specified in service" option was selected for either the Outer Tag or the Inner Tag in the service tag action, enter the applicable VLAN IDs.	text string Any established service tag action
S-VLAN (Outer Tag)	Specifies the service VLAN ID, if the Service Tag Action references "Specified in Service" for the Outer Tag.	none, any VLAN created on the system
C-VLAN (Inner Tag)	Specifies the customer VLAN ID, if the Service Tag Action references "Specified in Service" for the Inner Tag.	none, any VLAN created on the system

*Required fields
‡ Default

To create a video service on an xDSL port

1. On the Navigation Tree, click the xDSL port or bonding group on which you want to add a video service.
2. In the Work Area, click **Port** > **Services** > **Form**.

   Note: To create a service on an xDSL Bonded interface, on the Navigation Tree, click the xDSL Group under the VDSL2 card, and then in the Work Area, click **Services** > **Form**.

3. Reference the table above to configure the parameters.
4. In the menu, click **Apply**.

5. Verify your modem train rate using the **Port** > **Status** tab.

For CLI:

```
add eth-svc <service name> to-interface <dsl-port> bw-profile <bw-
profile-name> svc-tag-action [outer-vlan <vlan ID> inner-vlan <vlan
ID> mcast-profile <profile name> description <service> admin-state
[enabled|disabled]]

add eth-svc <service name> to-dsl-bond-interface <b-intfc> bw-
profile <bw-profile-name> svc-tag-action [outer-vlan <vlan ID>
inner-vlan <vlan ID> mcast-profile <profile name> description
<service> admin-state [enabled|disabled]]

add mcast-white-list <list-name> to-dsl-bond-interface <intfc-id>
eth-svc <svc-name>

add mcast-white-list <list-name> to-interface <intfc-id> eth-svc
<svc-name>
```

Adding Static IP Host Addresses and Subnets to xDSL Services

This topic shows you how to configure a static Ethernet service IP address or subnet, to associate with an xDSL port service or an xDSL bonding group service.

Configuration guidelines

- To identify a specific IP host for the purposes of securing the subscriber's traffic, the IP address must be provided, and optionally, the MAC address information.
- To identify an IP subnet, a MAC address is not required. The smallest subnet that can be provisioned in the system is 255.255.255.252, or /30.
- A given data service on a VDSL2 card represents exactly one subnet. To support a multi-homed router off an xDSL port, multiple data services must be provisioned to represent multiple subnets.
- The gateway address and subtending IP addresses must belong to the same subnet as indicated by the mask.
- The static IP address must not be the same as the gateway address.
- Duplicate static IP addresses are not allowed in the system.
- IP Source Verification for Static IP hosts requires MAC FF be enabled on the VLAN.
- If a provisioned static IP address conflicts with a DHCP learned IP address, the system will accept the static IP address and reject/delete the learned IP address.
- 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).
- 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 an xDSL Port
  - Static IP Subnet: Checks individual host IP addresses to the subnet and binds the subnet to an xDSL port

Before starting

Before starting this procedure, you will need to provision a service on the xDSL port where you want to configure the static IP host address or subnet.

Parameters for a static IP host address and subnet

Parameter	Description	Valid Options
Service*	Name of the service to associate with an IP address.	Any established service name
Type*	IP address or hostname of the primary SIP configuration server.	IP address, hostname
Static IP Host Address		
IP Address*	Static IP address for the ONT port service.	IP address in dotted quad" format: "192.168.1.100". Alternately, "none" can be used to reset the value to "0.0.0.0"
MAC Address	L2 MAC address associated with the IP address.	Six hexadecimal digits in the range 0-FF, optionally separated by colons
Gateway*	Address of the default gateway for subtending static IP address objects.	IP address in dotted quad" format: "192.168.1.100". Alternately, "none" can be used to reset the value to "0.0.0.0"
Net Mask*	IP Netmask for subscriber host.	IP address in dotted quad" format: "192.168.1.100". Alternately, "none" can be used to reset the value to "0.0.0.0"
Static Subnet		
Subnet Address*	Subscriber IP address subnet.	IP address in dotted quad" format: "192.168.1.100". Alternately, "none" can be used to reset the value to "0.0.0.0"
Gateway*	Address of the default gateway for subtending static IP address objects.	IP address in dotted quad" format: "192.168.1.100". Alternately, "none" can be used to reset the value to "0.0.0.0"
Subnet Mask*	Subnet mask for subtending static IP address objects.	IP address in dotted quad" format: "192.168.1.100". Alternately, "none" can be used to reset the value to "0.0.0.0"

*Required fields

To add a static IP host address

1. On the Navigation Tree, click the xDSL port or xDSL bonding group on which the service is provisioned where you want to add a static IP host address or subnet.
2. Click **Static IP/Subnet**.
3. If there are multiple Ethernet services provisioned on the selected port or bonding group, at the top of the Work Area, click the drop-down list 2-2-Data Gold-1-1 to select the service where you want to add a static IP host or subnet.
4. Click **Create** to open the Create Ethernet Service IP/Subnet dialog box.
5. Reference the table above to configure the parameters.
6. If you chose to add a subnet address, fill in the required addresses and subnet mask.
7. Click **Create**.

For CLI:

```
add static-ip-entry to-dsl-bond-interface <dsl-intfc> eth-svc <s-
name> type host ip <h-ip> netmask <n-ip> default-gw <gw-ip> [mac <m-
address>]

add static-ip-entry to-dsl-bond-interface <dsl-intfc> eth-svc <s-
name> type subnet ip <sub-ip> netmask <n-ip> default-gw <gw-ip>

add static-ip-entry to-interface <dsl-intfc> eth-svc <s-name> type
host ip <h-ip> netmask <n-ip> default-gw <gw-ip> [mac <m-address>]

add static-ip-entry to-interface <dsl-intfc> eth-svc <s-name> type
subnet ip <sub-ip> netmask <n-ip> default-gw <gw-ip>
```

Configuring Voice Services

This section describes how to create services on VDSL2 voice ports.

Topics Covered

This section covers the following **topics in bold** that are part of the overall VDSL2 services configuration process:

1. Configure network uplinks for VDSL2 services.
2. Create profiles that support VDSL2 applications.
3. **Configure voice subscriber services.**
    - **Overview of the VDSL2 voice services provisioning**
    - **Provisioning a Voice port**
    - **Configuring a SIP service**
    - **Configuring a TDM Gateway service**
    - **Configuring H.248 gateway voice service**

Overview: Configuring VDSL2 Voice Services

The E-Series supports three options for providing VDSL2 voice services:

- SIP gateway and H.248 gateway services options provide traditional VoIP offerings for the VDSL combo lines, where a gateway client on the VDSL2 unit interoperates with a SIP softswitch or H.248 media gateway controller (softswitch), and converts analog voice calls to packet format.

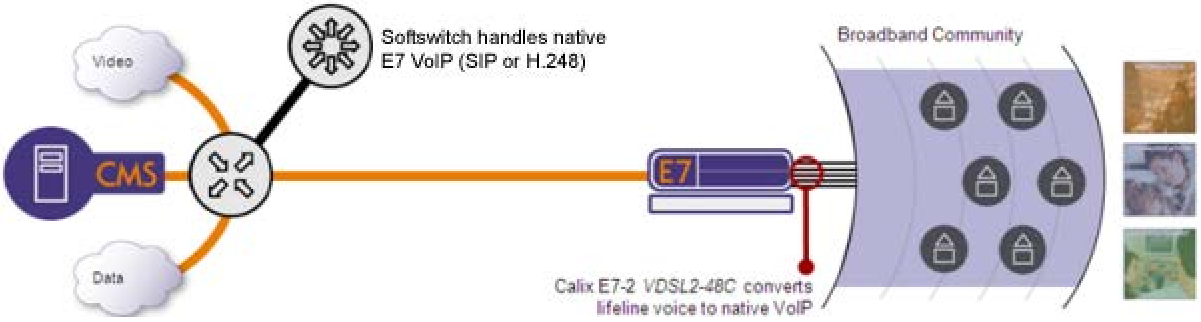
Video
CMS
Data
Softswitch handles native
E7 VoIP (SIP or H.248)
E7
Broadband Community
Calix E7-2 VDSL2-48C converts
lifeline voice to native VoIP

- TDM gateway option interoperates with a C7 voice gateway (EGW or VIPR), which converts VoIP traffic back to TDM format for exchange on the PSTN. 75

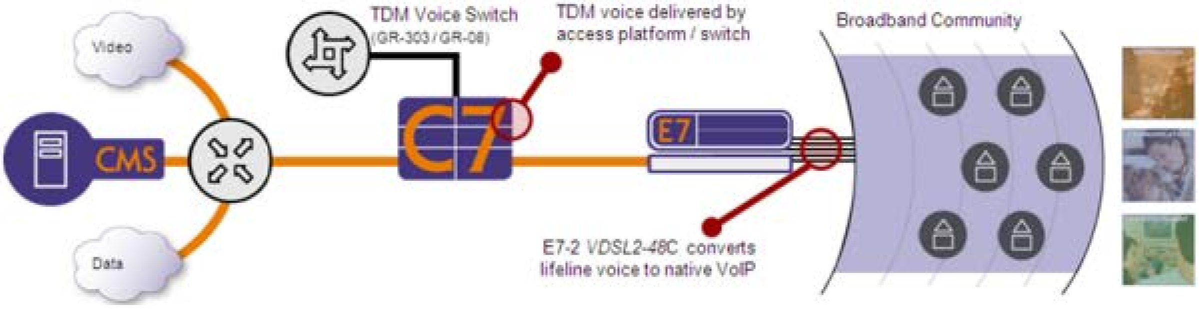
Video
CMS
Data
TDM Voice Switch (GR-303 / GR-08)
C7
TDM voice delivered by access platform / switch
E7
Broadband Community
E7-2 VDSL2-48C converts lifeline voice to native VoIP

This section describes how to turn-up an E-Series for VoIP service.

Note: Each VDSL2 card only supports one voice service type at a time, either SIP, H.248, or TDM gateway mode.

Turn-up process for VoIP Services

The VoIP service turn-up process for SIP, H.248, and TDM Gateway VoIP services includes the steps in the following sequence. You should have the following information on hand to configure a voice service:

- IP address (or domain name) of your VoIP provider's SIP server (or outbound proxy SIP server), H.248 media gateway controller, or C7 TDM gateway
- IP address to assign to the local VoIP services IP host on the VDSL2 unit
- VLAN ID for the VoIP service
- Service information per port, such as the SIP user name/password (for SIP mode, such as the VoIP telephone number), termination ID (for H.248 mode), or GR-303/GR-8 IG CRV (for C7 TDM Gateway mode).

Before starting

Before starting the voice services configuration process, check that the following conditions are met:

1. **The network uplinks for VDSL2 services are configured.**
    - Ethernet port interfaces
    - *Ethernet ports* (on page 21)
    - Service VLANs
    - *VLAN memberships* (on page 64)

    Uplink interfaces (or ERPS domain, if the uplink resides on a different shelf) must be added to the VLAN membership.

2. **The necessary system profiles that support VDSL2 voice applications are created.**
   - *IP host for voice services on a VDSL2 card* (on page 156)

     Each VDSL2 card voice service option requires an IP host definition object that specifies an IP host address and outer tag for VoIP. (The definition references a system default tag action that specifies the classifying and marking of packets from the subscriber port into the service VLAN specified in the IP host.)
   - *Dial plan for SIP voice service* (on page 158)
   - *H.248 Gateway* (on page 174)
   - *Profiles for voice services* (on page 155)

     For each type of voice service (*TDM Gateway* (on page 169), *SIP Gateway* (on page 161), *H.248 Gateway* (on page 172)), a profile must be defined.

Step 3. Configuring a Voice Service

After completing the steps outlined above, continue with the following topics in this section to configure a service on a Voice port:

- *Configure a Voice port* (on page 203)
- *Configure SIP service on a Voice port* (on page 205)

  or
- *Configure TDM gateway service on a Voice port* (on page 208)

  or
- *Configure H.248 gateway service on a Voice port* (on page 210)

Configuring a VDSL2 Voice Port

This topic describes how to configure a VDSL2 Voice (POTS) port for service.

Parameters

You can provision the following parameters for a VDSL2 Voice port:

Parameter	Description	Valid Options
Admin State	Administrative state of the ONT POTS port.	Enabled ‡ Enabled-no-alarms Disabled
Subscriber ID	Subscriber ID information, such as phone number, or account number.	String up to 27 characters (blank) ‡
Description	Optional description field for the port, subscriber address, name, or service.	String up to 27 characters (blank) ‡
Impedance	Impedance in ohms.	600-ohm , 900-ohm‡
Signal Type	Signal type to use for the POTS line. • **Loop start** (i.e. POTS) is the alternative to ground start.	loop-start ‡

Parameter	Description	Valid Options
System Tx Loss Plan	System transmit (tx) signal-level loss plan for POTS (voice) line provides attenuation settings, according to various standards, to reduce the perception of noise and echo on the line. • **GR-909** loss plan (less attenuation of -2 dB – higher signal level) is more compatible with GR-303 type phone systems. • **ANSI** loss plan (more attenuation of -3 dB – lower signal level) is more compatible with American VoIP type phone systems. • **ETSI-pstn** loss plan (more attenuation of -4 dB - lowest signal level) is more compatible with European VoIP type phone systems. • **Manual** loss plan allows you to set the transmit gain from the range -12.0 to 6.0 dB. **Note:** The TX and RX loss plan settings must match.	gr-909 ‡, ansi, etsi-pstn, manual
System Rx Loss Plan	System receive (rx) signal-level loss plan for POTS (voice) line provides attenuation settings, according to various standards, to reduce the perception of noise and echo on the line. • **GR-909** loss plan (Rx/Tx = -4/-2) is more compatible with GR-303 type phone systems. • **ANSI** loss plan (Rx/Tx = -9/-3) is more compatible with American VoIP type phone systems. • **ETSI-pstn** loss plan (ETSI EG 201 185 loss plan. Rx/Tx = -11/-4) is more compatible with European VoIP type phone systems. • **Manual** loss plan allows you to set the receive gain from the range -12.0 to 6.0 dB. **Note:** The TX and RX loss plan settings must match.	gr-909 ‡, ansi, etsi-pstn, manual
Transmit Gain	Transmit gain for a voice port. **Note:** This attribute is in effect only when the System Tx Loss Plan is set to manual.	-12.0 to 6.0 dB 0 ‡
Receive Gain	Receive gain for a voice port. **Note:** This attribute is in effect only when the System Rx Loss Plan is set to manual.	-12.0 to 6.0 dB in increments of 0.5 dB. 0 ‡

‡ Default

To configure a VDSL2 Voice (POTS) port for service

1. On the Navigation Tree, click the VDSL2 Voice port that you want to configure.
2. In the Work Area, click **Provisioning**.
3. Reference the table above to configure the parameters.
4. From the menu, click **Apply** to save changes.

Syntax:

```
set pots-port <p-port>
show pots-port
show pots-port detail
show pots-port <port> [detail|sip-svc|tdm-gw-svc]
```

Creating a SIP Service

This topic describes how to configure SIP voice service on a VDSL2 card Voice (POTS) port. Each VDSL2 card can only support one voice service type, either SIP, TDM gateway, or H.248 gateway.

Use the Voice Port Provisioning screen to activate Voice ports and apply the following previously-created profiles:

- SIP gateway profile
- Dial plan profile
- IP Host

Use the same screen to disable Voice ports.

Profiles referenced in service provisioning

When provisioning subscriber services on a VDSL2 card Voice port, you set up information that is common to multiple subscribers and maintained in profiles. Repeating the similar task of subscriber service activation then consists of referencing the existing profiles and sometimes providing subscriber-specific information to complement the profiles. Although an IP Host is configured for voice services on a VDSL2 card, it is not required to explicitly reference it when creating a SIP service because it applies to all voice services provisioned on the VDSL2 card. However, you are given the option to modify the existing IP Host or replace it with a new IP Host definition.

Note: When a profile is changed, all subscribers using that profile are affected. Profiles cannot be deleted while they are in use. Profiles can also be reloaded for a specific VDSL2 Voice port from the Service Table by clicking **Refresh**.

Information you need

You must have the following information on hand to configure SIP voice service:

- SIP server and call feature information with which to populate a SIP configuration file
- (Optional) IP addresses for SIP lines (only if using a static IP addressing scheme)
- SIP profile to use on the voice (POTS) port
- VoIP telephone number
- IP address of your VoIP provider's SIP server
- (For SIP VoIP mode only) The IP address or domain name of the SIP server or outbound proxy SIP server.

Parameters

You can provision the following parameters for a SIP service on a VDSL2 voice port:

Parameter	Description	Valid Options
Subscriber Port		
Subscriber ID	Subscriber ID information, such as phone number, or account number.	Any VDSL2 Voice port
Description	Optional description field for the port, subscriber address, name, or service.	String up to 27 characters (blank) ‡
Service Type	(Form view only) The type of voice service to create.	SIP ‡, TDM Gateway, H.248 Gateway
Create SIP Service		
SIP Gateway Profile*	Name of the SIP gateway profile to use, or removing the service from the port **Note:** When a profile is changed, all subscribers using that profile are affected. Profiles cannot be deleted while they are in use. Profiles can also be reloaded for a specific VDSL2 Voice port from the Service Table by clicking **Refresh.**	Any established SIP profile, Remove Service ‡
User Name*	User name for registration with the SIP server (it forms the UserId field, left of the @ in the SIP URL). For example, type the subscriber phone number. The “@” character is supported in the username for VDSL and T-Series ONT SIP services.	text string
Enable Caller ID	Enables or disables Caller ID feature.	select to enable, clear to disable ‡
Three-way Calling	Enables or disables Three-way Calling feature.	select to enable, clear to disable ‡
Dial Plan	Name of the dial plan to apply to the SIP VoIP service.	any available plans
Enable Direct Connect	Whether to enable support for Direct Connect (Hot Line) and Direct Connect Timer (Warm Line)	select to enable, clear to disable ‡
Direct Connect	If Direct Contact support is enabled, enter a hot line number to immediately dial after an off hook event (for example, taxi/hotel phones hanging on the walls of airports).	Number string between 0-15 numbers.
Direct Connect Timer	If Direct Contact support is enabled, optionally configure the delay in seconds before the Direct Contact number command is dialed. This feature is known as warm line support and may be used to redirect a call to a pre-recorded message, a front desk operator, or an emergency contact.	0‡-35
Universal Resource Indicator*	This is the Universal Resource Indicator (URI) of the port. If a local URI is "aaa@bbb", "aaa" is the telephone number configured, and "bbb" is the domain name of the SIP server configured. **Note:** The AoR domain can be defined in the SIP Service with the URI format of “<user>@<{host|domain}” When defining a domain in the URI field of the SIP Service, it will override the domain provisioned in the SIP Gateway Profile.	text string up to 32 characters
Password*	The SIP registration password that matches the user name, if required.	24-character text string, with the exception of ! ' "
Call Waiting	Enables or disables Call Waiting feature: Accept the new call and put the current call on hold.	select to enable, clear to disable ‡
T38 Fax Relay	Enables or disables the T.38 fax relay feature.	select to enable, clear to disable ‡
Msg Waiting Indicator	Whether to enable the audio and visual message indicator.	select to enable ‡, clear to disable

Parameter	Description	Valid Options
Update SIP IP Host		
ID	Name of IP Host to use or update, or indicate the creation of a new IP Host.	Any established IP host for the card
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 you set the Host Protocol to dhcp, any previously set 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.	static dhcp
S-VLAN (Outer Tag) *	Indicates the customer-specific tag.	2-4093 (Except for 1002-1005 which are reserved for E-Series operation.)
Name*	If you are creating a new IP host, enter a descriptive name for it. This field will auto-fill if you select a previously configured IP host.	Character string up to 31 characters.
Static IP	If the host protocol is static, IP address statically assigned to the VDSL2 card. This attribute is ignored, yet preserved, if the host protocol is later switched to DHCP. The static IP address must not be the same as the gateway address.	4-byte IP address
Static IP Mask	If the host protocol is static, IP network mask assigned to the VDSL2 card. This attribute is ignored, yet preserved, if the host protocol is later switched to DHCP.	4-byte IP address
Static IP Gateway	If the host protocol is static, the Static IP gateway 4-byte address to use in routing its traffic to the SIP server. This attribute is ignored, yet preserved, if the host protocol is later switched to DHCP. The gateway address and subtending IP addresses must belong to the same subnet, as indicated by the mask	4-byte IP address

*Required fields
‡ Default

To create a SIP voice service

1. On the Navigation Tree, click the Voice port on which you want to add a SIP service.
2. In the Work Area, click **Services** > **Table > Create > SIP Service** to open the Create dialog box.
3. Reference the table above to configure the parameters.
4. Click **Create** to activate a SIP voice service on the subscriber port.
5. Click **Provisioning** to verify the service operation.

For CLI:

```
add sip-svc to-pots-port <vdsl-port> ip-host <lc-ip> sip-gw-profile
<sipgw-profile> user <user name> password <pswd> uri <URI ID> [call-
waiting|caller-id|three-way-calling|t38-fax-relay|dial-plan|admin-
state|direct-connect|direct-connect-timer|msg-waiting-indicator]
```

Example CLI:

```
add sip-svc to-pots-port 2/1 ip-host 2/SIP_IPhost sip-gw-profile
SIP_gw user 2012042453 password password uri 2012042453
```

Creating a TDM Gateway Service

This topic describes how to configure TDM gateway voice service (to a Calix C7 voice gateway) on E-Series VDSL2 POTS ports. The TDM gateway solution employs VoIP between the E-Series POTS ports and the C7 voice gateway, which converts voice traffic back to TDM format for exchange on the PSTN. Each VDSL2 card or service unit can only support one voice service type (TDM gateway, SIP, or H.248 gateway).

Use the VoIP Port Provisioning screen to activate VoIP ports and apply the following previously-created profiles:

- TDM gateway profile
- TDM gateway service group (optional)

Use the same screen to disable Voice ports.

Note: For TDM gateway services, you can use the C7 gateway's internal DHCP server to provide IP addresses for the VoIP hosts (recommended), or use an external DHCP server option (requires Option 43 parameters included in the DHCP Offer).

Note: For instructions to configure the C7 TDM voice gateway, refer to the *Calix C7 VoIP Services Guide*.

Profiles referenced in service provisioning

When provisioning subscriber services on a VDSL2 Voice port, you set up information that is common to multiple subscribers and maintained in profiles. Repeating the similar task of subscriber service activation then consists of referencing the existing profiles and sometimes providing subscriber-specific information to complement the profiles.

Note: When a profile is changed, all subscribers using that profile are affected. Profiles cannot be deleted while they are in use. Profiles can also be reloaded for a specific VDSL2 Voice port from the Service Table by clicking **Refresh.**

Voice Service IP Host

Although an IP Host is configured for voice services on a VDSL2 card, it is not explicitly referenced when creating a voice service because it applies to all voice services provisioned on the VDSL2 card. However, you are given the option to modify the existing IP Host or replace it with a new IP Host definition, when provisioning the voice service.

TDM Gateway Service Report in CMS

CMS allows you to run a database search for provisioned TDM Gateway services on a multi-platform basis, and then generate a report on the results.

Do one of the following, depending on whether you have CMS Desktop or CMS Web open:

- (CMS Desktop) On the **Tools** menu, click **Search** > **Multi-Platform** > **TDM Service**, and then click **Submit**.
- (CMS Web) Open CMS Web. In the **Module** list on the left, click **Configuration** > **Multi-Platform** > **TDM Service**.

Information you need

You must have the following information on hand to configure TDM Gateway voice service:

- TDM Gateway profile to use on the Voice port
- Call Reference Value (CRV) identifiers to use for the subscriber voice line
- Service Group to define the C7 network and interface groups used for delivering the TDM Gateway service (available through CMS only)

Parameters

You can provision the following TDM Gateway voice service parameters accessed from the Services page:

Parameter	Description	Valid Options
	Subscriber Port	
Subscriber Port*	System address of the port where the subscriber services will be provisioned.	Any ONT Voice port
Subscriber ID	Subscriber ID information, such as phone number, or account number.	String up to 27 characters (blank) ‡
Description	Optional description field for the port, subscriber address, name, or service.	String up to 27 characters (blank) ‡
Service Type	(Form view only) Type of voice service to create.	SIP ‡, TDM Gateway, H.248 Gateway
	Create TDM Gateway Service	
TDM Gateway Profile*	Name of the TDM Gateway profile to use.	Any established TDM Gateway profile
Service Group (via CMS only)	(Optional) Defines the C7 network and interface groups used for delivering the TDM Gateway service, allowing CMS to automatically create and validate C7 gateway cross-connects. **Note:** If you do not reference a service group, you must configure the SIP VCG on the C7 network and add a VoIP connection to the E-Series. For details, see the *Calix C7 VoIP Services Guide*.	Any global TDM Service Group established from CMS
Call Reference Value*	Call Reference Value (CRV) identifies the subscriber line, as provisioned on the Calix C7 GR-303 interface group (upper case). For example, **N1-1-IG1-224**. This CRV ID must be provisioned on the CRVs used to build translation tables on the Class 5 switch that map remote connections to internal circuits at the switch. **Note:** The specified CRV must also be provisioned on the interface group at the Calix C7 voice gateway for service to work. The specified CRV must be entered in upper case as shown above.	Specific CRV for this line appearance. **Note:** The CRV MUST be in upper case.

Parameter	Description	Valid Options
Update TDM Gateway IP Host		
ID	Name of the IP Host to use for the voice services provisioned on the VDSL2 card.	An existing IP Host
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 ‡
Static IP	If the host protocol is static, the IP address statically assigned to the VDSL2 card. This attribute is ignored, yet preserved, if the host protocol is later switched to DHCP.	4-byte IP address
Static IP Mask	If the host protocol is static, the IP network mask assigned to the VDSL2 card. This attribute is ignored, yet preserved, if the host protocol is later switched to DHCP.	4-byte IP address
Static IP Gateway	If the host protocol is static, the Static IP gateway address for the VDSL2 card to use in routing its traffic to the SIP server. This attribute is ignored, yet preserved, if the host protocol is later switched to DHCP.	4-byte IP address
S-VLAN (Outer Tag) *	Outer tag VLAN ID (customer-specific) for the service tag action to reference.	2-4093 (Except for 1002-1005 which are reserved for E-Series operation.)

*Required fields
‡ Default

To create a TDM gateway service

1. On the Navigation Tree, click the Voice port on which you want to add a SIP service.
2. In the Work Area, click **Services** > **Table > Create > TDM Gateway Service** to open the Create dialog box.
3. Reference the table above to configure the parameters.
4. Click **Create** to activate a TDM Gateway voice service on the subscriber port.
5. Click **Provisioning** to verify the service operation.

For CLI:

```
add tdm-gw-svc to-pots-port <pots port> ip-host <card ip> tdm-gw-
profile <profile name> crv <crv> [admin-state]
```

Example CLI:

```
add tdm-gw-svc to-pots-port 1/1 ip-host 1/tdmgw tdm-gw-profile tdmgw
crv N1-1-IG1-1
```

Creating an H.248 Service

This topic describes how to configure H.248 gateway voice service on a VDSL2 card Voice (POTS) port. Each VDSL2 card can only support one voice service type, either SIP, TDM gateway, or H.248 gateway.

Use the Voice Port Provisioning screen to activate Voice ports and apply the following previously-created profiles:

- H.248 gateway profile
- H.248 gateway
- IP Host

Use the same screen to disable Voice ports.

Profiles referenced in service provisioning

When provisioning subscriber services on a VDSL2 card Voice port, you set up information that is common to multiple subscribers and maintained in profiles. Repeating the similar task of subscriber service activation then consists of referencing the existing profiles and sometimes providing subscriber-specific information to complement the profiles. Although an IP Host is configured for voice services on a VDSL2 card, it is not required to explicitly reference it when creating a voice service because it applies to all voice services provisioned on the VDSL2 card. However, you are given the option to modify the existing IP Host or replace it with a new IP Host definition.

Note: When a profile is changed, all subscribers using that profile are affected. Profiles cannot be deleted while they are in use. Profiles can also be reloaded for a specific VDSL2 Voice port from the Service Table by clicking **Refresh**.

Media gateway considerations

There is a default inactivity timeout of 25 seconds in the Media Gateway (MG) to determine when the MG to Media Gateway Controller (MGC) association is down. Under normal operation, there is a keep alive heartbeat/polling interval that is initiated by the MGC to the MG. If the interval value is longer than 25 seconds, the E-Series reports an alarm condition. Therefore, the MGC (H.248 switch) must have a polling interval less than 25 seconds to avoid the E-Series alarm condition described.

Parameters

You can provision the following parameters for an H.248 Gateway service on a VDSL2 port:

Parameter	Description	Valid Options
	Subscriber Port	
Subscriber Port*	System address of the port where the subscriber services will be provisioned.	Any ONT Voice port
Subscriber ID	Subscriber ID information, such as phone number, or account number.	String up to 27 characters (blank) ‡
Description	Optional description field for the port, subscriber address, name, or service.	String up to 27 characters (blank) ‡
Service Type	(Form view only) Type of voice service to create. Select **H.248 Gateway Service**.	SIP ‡, TDM Gateway, H.248 Gateway

Parameter	Description	Valid Options
Create H.248 Gateway Service		
Gateway Profile*	Name of the H.248 Gateway profile to use for the service.	Any established H.248 Gateway profile
Termination ID*	The ID to use for ephemeral terminations.	Specific CRV for this line appearance, up to a maximum of 48 characters. **Note:** The CRV MUST be in upper case.
Update H.248 Gateway IP Host		
ID	Name of the IP Host to use for the voice services provisioned on the VDSL2 card.	An existing IP Host
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 ‡
Static IP	If the host protocol is static, the IP address statically assigned to the VDSL2 card. This attribute is ignored, yet preserved, if the host protocol is later switched to DHCP.	4-byte IP address
Static IP Mask	If the host protocol is static, the IP network mask assigned to the VDSL2 card. This attribute is ignored, yet preserved, if the host protocol is later switched to DHCP.	4-byte IP address
Static IP Gateway	If the host protocol is static, the Static IP gateway 4-byte address for the VDSL2 card to use in routing its traffic to the SIP server. This attribute is ignored, yet preserved, if the host protocol is later switched to DHCP.	4-byte IP address
S-VLAN (Outer Tag) *	Outer tag VLAN ID for the service tag action to reference.	2-4093 (Except for 1002-1005 which are reserved for E-Series operation.)

*Required fields
‡ Default

To create an H.248 gateway voice service

1. On the Navigation Tree, click the VDSL2 **Voice** port on which to create the service.
2. In the Work Area, click **Services** > **Form**.
3. Reference the table above to configure the parameters.
4. Click **Create** to activate an H.248 Gateway voice service on the subscriber port.
5. Click **Provisioning** to verify the service operation.

For CLI:

- **`add h248-gw-svc to-pots-port <vdsl-port> h248-gw <gw-name> termination-id <t-id> h248-gw-profile <p-name> [admin-state]`**

Configuring T1/E1 PWE3 Services

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.*

xDSL Port Power Save

The xDSL ports have a power save feature disabled by default so that when the power source changes from AC to battery backup during an AC power failure, the system switches-off lower priority xDSL services, leaving the remaining battery power available for higher priority voice connections to emergency services (for example, 911).

To facilitate a smooth transition of switching off xDSL services, a power save delay timer is available on a per-port basis so services are maintained for a number of minutes after an AC power failure, and then the port transitions to a power save state.

This power save feature is also known as power-shedding or CAT3 power shutdown.

To modify the power save parameters for an xDSL port

1. On the Navigation Tree, click the xDSL port on which you want to modify the power save parameters.
2. In the Work area, do the following:
   a. In the Power Save checkbox, choose whether to power save feature.
      - Select the checkbox to enable the feature.
      - Clear the checkbox to disable the feature.

   b. In the Power Save Timeout box, enter the number of minutes to wait after an AC power failure, before the port transitions to a power save state
3. Click **Apply**.

For CLI:

```
set dsl-port <name> basic [power-save|power-save-timeout]
show dsl-port <name> basic
```

Appendix A

Reference Information

This appendix provides general reference information about the Calix E-Series Ethernet service access platform.

Topics Covered

This appendix covers the following topics:

- System capacities
- LED behavior
- Line card additional status descriptions

System Support Capacities

The Calix E7 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 (E7 line cards share a common table)	32,000
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 rate limiting	• 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.
Concurrent Multicast Streams	• E7-20: 2000 • 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

Description	Capacity
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 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. 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
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)	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
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
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	264
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
Security profiles	16 Security profiles used by VDSL interfaces (including bonded-Links) must have DHCP lease limt = 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 VDSL2-based MCCs in an ERPS ring, the maximum VDSL2 combo subscribers/node and nodes/ring are lower than in other configurations. These capacities are shown in **bold** in the table below, and apply to all nodes in a ring when at least one VDSL2-based MCC is present in the ring.

	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**	**ERPS Ring, VDSL2 MCC**
Max Shelves/Node	1	10	10	10	**5 (combo)** 10 (overlay)
Max Nodes/Ring	16	N/A	N/A	16	**8**
Max Subs/Node	combo — 96 overlay — 48	combo — 864 overlay — 432	combo — 960 overlay — 480	combo — (Note 2) overlay — 432	**combo — 480** overlay — 480

Note 1) For modular chassis nodes with GPON/AE-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) A GPON/AE-based MCC node supports the maximum configuration of 864 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.

General xDSL Support Capacities

Description	Capacity
MTU Packet Size (DSL services)	1536 Bytes total frame size
VDSL2 Bandwidth	Up to 100 Mbps downstream, 50 Mbps upstream (Profile 17a). For VDSL2 bonded pairs, up to 100 Mbps downstream also applies.
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	24 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.
Adjacent xDSL ports bonded train rate limits	120 Mbps Downstream 60 Mbps Upstream
Bandwidth Shaping (DSL ports)	64 kbps increments
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

E-Series 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.
- 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

CTRL LED	SRVC LED	FAIL LED	State Description
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 Activiation
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
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

E-Series Line Card Additional Status Descriptions

The table below shows the possible states for an E-Series line card Additional Status.

Line Card Additional Status	Description
default-prov	Indicates that the object's parameters have never been changed from default values.
child-prov	Indicates that the object has subtending records provisioned. these may be ports or interfaces that have been updated by the user.
present	Indicates that the object is present.
system-disabled	Indicates an object has a service affecting alarm reported against it, or any of its parents. For example: • An ONT parent is the system object. • A Card parent is the shelf, and then the system.
user-disabled	Indicates an object is disabled by the user (Admin status = disabled).
degraded	Indicates an object (ONT or a Card) has a non-service affecting alarm reported against it.
active	Indicates that the card is 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.
standby	Indicates the card is in standby status.

Note: If the card has the default-prov state and does not have the child-prov state, this indicates that the card hierarchy is completely default and will be deleted from the database upon card departure.