




OSPF

Open Shortest Path First (OSPF) Protocol & Operation
Overview



In This Section

- + Recommended OSPF Resources
- + What is OSPF?
- + Why use OSPF?
- + How does OSPF Work?
 - + Forming Adjacencies
 - + Path Selection
 - + Topology Maintenance

Recommended OSPF Resources

+ Books

- + OSPF: Anatomy of an Internet Routing Protocol
- + Routing TCP/IP Volume 1
- + Cisco IP Routing: Packet Forwarding and Intra-domain Routing Protocols

+ Standards

- + [RFC 2328 "OSPF Version 2"](#)
- + [RFC 5340 "OSPF for IPv6"](#)

+ Documentation

- + [OSPF Technology Documentation](#)
- + [OSPF Configuration Guide](#)

What Is Open Shortest Path First (OSPF)?

- + Open Standards Based Interior Gateway Routing Protocol (IGP)
 - + OSPFv2 - RFC 2328 “OSPF Version 2”
 - + OSPFv3 - RFC 5340 “OSPF for IPv6”
- + Link-State Protocol
 - + Uses Dijkstra SPF Algorithm
- + “Classless” Protocol
 - + Supports VLSM And Summarization

Why Use OSPF?

- + Guarantees Loop-Free Topology
 - + All routers agree on overall topology
 - + Uses Dijkstra SPF Algorithm for calculation
- + Standards Based
 - + Inter-operability between vendors
- + Large Scalability
 - + Hierarchy through “areas”
 - + Topology summarization

Why Use OSPF? (cont.)

- + Fast Convergence
 - + Actively Tracks Neighbor Adjacencies
 - + Event Driven Incremental Updates
- + Efficient Updating
 - + Uses reliable multicast and unicast updates
 - + Non-OSPF devices do not need to process updates
- + Bandwidth Based Cost Metric
 - + More flexible than static hop count

Why Use OSPF? (cont.)

- + Control Plane Security
 - + Supports multiple forms of authentication
 - + E.g. Clear Text, MD5, SHA, IPsec, etc.
- + Extensible
 - + Future application support through “opaque” LSA
 - + E.g. MPLS Traffic Engineering

How OSPF Works

- + Step 1
 - + Discover OSPF Neighbors & Exchange Topology Information
- + Step 2
 - + Choose Best Path via SPF
- + Step 3
 - + Neighbor and Topology Table Maintenance

Step 1 – Neighbor & Topology Discovery

- + Like EIGRP, OSPF uses *Hello* packets to discover neighbors on OSPF enabled attached links
 - + Transport via IP protocol 89 (OSPF)
 - + Sent as multicast to 224.0.0.5 or 224.0.0.6, or unicast
 - + More on this later...
- + Hello packets contain attributes that neighbors must agree on to form “adjacency”
 - + Once adjacency is negotiated, LSDB is exchanged

Negotiating OSPF Adjacencies

- + OSPF adjacency occurs when connected neighbors use hello packets to agree on unique and common attributes
 - + Not all OSPF neighbors actually form adjacency
 - + Most OSPF configuration problems happen at this stage

Unique OSPF Adjacency Attributes

- + Router-ID
 - + Node ID in the Link State Graph
 - + Chosen based upon...
 - + Manual configuration
 - + Highest active Loopback IP
 - + Highest active Interface IP
- + Interface IP Address
 - + For OSPFv2 the interface's primary IP address
 - + For OSPFv3 the interface's link-local address

Common OSPF Adjacency Attributes

- + Interface Area-ID
- + Hello interval & dead interval
- + Interface network address
- + Interface MTU
- + Network Type
- + Authentication
- + Stub Flags
- + Other optional capabilities

OSPF Hello Packets

- + OSPF routers periodically send Hello packets out OSPF enabled links every *HelloInterval*
- + Hello packet contains
 - Local Router-ID
 - Local Area-ID
 - Local Interface Subnet Mask
 - Local Interface Priority
 - Hello Interval
 - Dead Interval
 - Authentication Type & Password
 - DR/BDR Addresses
 - Options (e.g. stub flags, etc.)
 - Router IDs of other neighbors on the link

OSPF Adjacency State Machine

- + OSPF adjacency process uses 8 states to determine progress of adjacency establishment
- + Down
 - + No hellos have been received from neighbor
- + Attempt
 - + Unicast hello packet has been sent to neighbor, but no hello has been received back
 - + Only used for manually configured NBMA neighbors (more on this later...)
- + Init
 - + I have received a hello packet from a neighbor, but they have not acknowledged a hello from me

OSPF Adjacency State Machine (cont.)

- + 2-Way
 - + I have received a hello packet from a neighbor and they have acknowledged a hello from me
 - + Indicated by my Router-ID in neighbor's hello packet
- + ExStart
 - + First step of actual adjacency
 - + Master & slave relationship is formed, where master has higher Router-ID
 - + Master chooses the starting sequence number for the Database Descriptor (DBD) packets that are used for actual LSA exchange

OSPF Adjacency State Machine (cont.)

- + Exchange
 - + Local link state database is sent through DBD packets
 - + DBD sequence number is used for reliable acknowledgement/retransmission
- + Loading
 - + Link State Request packets are sent to ask for more information about a particular LSA
- + Full
 - + Neighbors are fully adjacent and databases are synchronized

Step 2 – Choose Best Path via SPF

- + Once databases are synchronized, path selection begins
- + Each router's LSAs include a "cost" attribute for each described link
 - + Best path to that link is lowest end-to-end cost
 - + Multiple Equal Cost paths are allowed (ECMP)
- + Cisco's implementation uses bandwidth based cost, but per RFC it is arbitrary
 - + Default Cisco Cost = $100\text{Mbps} / \text{Link Bandwidth}$
 - + Reference bandwidth can be modified to accommodate higher speed links (e.g. TenGigabitEthernet)

Step 3 – Neighbor & Topology Maintenance

- + Once adjacencies established and SPT built, OSPF state machine tracks neighbor and topology changes
 - + Hello packets used to track neighbor changes
 - + LSA fields used to track topology changes

Tracking Neighbor Changes

- + Hello packets continue to be sent on each OSPF enabled link every HelloInterval
 - + 10 or 30 seconds by default depending on interface type
- + If a Hello packet is not received from a neighbor within *RouterDeadInterval*, the neighbor is declared down
 - + Defaults to 4 times HelloInterval
 - + Can be as low as 1 second for faster convergence

Tracking Topology Changes

- + When a new LSA is received it is checked against the database for changes such as...
 - + Sequence number
 - + Used to track new vs old LSAs
 - + Age
 - + Used to keep information new and withdraw old information
 - + Periodic flooding occurs after 30 minutes
 - + “paranoid” update
 - + LSAs that reach *MaxAge* (60 minutes) are withdrawn
 - + Checksum
 - + Used to avoid transmission & memory corruption

LSA Flooding

- + When change is detected new LSA is generated and “flooded” (sent) out all links
 - + OSPF does not use split horizon
 - + Self-originated LSAs are simply dropped
- + Not all LSA changes require SPF to recalculate
 - + e.g. link up/down event vs. seq number change
 - + See RFC 2328 “13. The Flooding Procedure” for details





OSPF

OSPF Single Area Configuration



In This Section

- + Configuring Basic OSPF
- + Verifying OSPF Adjacencies
- + Verifying the OSPF Database

OSPF Prerequisites

- + IP routing must be enabled
 - + I.e. **ip routing**
- + Must be an “up/up” interface running IP
 - + Used for OSPF Router-ID

Enabling OSPF

- + Enable global OSPF process
 - + **router ospf [process-id]**
 - + Process-id locally significant
 - + Exception is MPLS L3VPN
- + Enable interface OSPF process
 - + Process level
 - + **network [address] [wildcard] area [area-id]**
 - + Interface level
 - + **ip ospf [process-id] area [area-id]**

OSPF Network Statement

- + Useful for enabling OSPF on multiple interfaces
 - + Wildcard mask does not relate to subnet mask
- + Most specific match determines the area
 - + **network 0.0.0.0 255.255.255.255 area 0**
 - + **network 1.0.0.0 0.255.255.255 area 1**
 - + **network 1.2.0.0 0.0.255.255 area 2**
 - + **network 1.2.3.0 0.0.0.255 area 3**
 - + **network 1.2.3.4 0.0.0.0 area 4**

OSPF Interface Statement

- + Enables OSPF on the primary and secondary IP addresses
 - + Secondary advertisement can be disabled
- + OSPF stays enabled even if IP address changes

OSPF Verification

- + Verify OSPF is enabled
 - + **show ip ospf**
 - + **show ip ospf interface [brief]**
- + Verify OSPF adjacencies
 - + **show ip ospf neighbor**
 - + **debug ip ospf adj**
- + Verify OSPF database
 - + **show ip ospf database [router | network | summary | ...]**





OSPF

Troubleshooting OSPF Adjacencies



In This Section

- + Troubleshooting OSPF Adjacency
 - + Understanding the OSPF State Machine
 - + Interpreting show commands
 - + Interpreting debug commands

Troubleshooting OSPF Adjacencies

- + Where can problems arise?
 - + Transport problems
 - + Attribute negotiation problems
- + Useful troubleshooting commands
 - + show ip ospf neighbor
 - + show ip ospf database
 - + debug ip ospf adj
 - + debug ip packet
 - + **Use with caution**

OSPF Adjacency State Machine

- + Normal OSPF Adjacency State Machine Order
 - + Down/Attempt
 - + Init
 - + 2-Way
 - + Stop here for DROthers
 - + ExStart
 - + Exchange
 - + Loading
 - + Full

OSPF Adjacencies Attributes

- + Unique OSPF Adjacency Attributes
 - + Router-ID
 - + Interface IP Address
- + Common OSPF Adjacency Attributes
 - + Interface Area-ID
 - + Hello interval & dead interval
 - + Interface network address
 - + Interface MTU
 - + Network Type
 - + Authentication
 - + Stub Flags
 - + Other optional capabilities





OSPF

OSPF Areas and LSAs



In This Section

- + Scaling OSPF with Areas
- + OSPF LSA Types
- + OSPF Path Selection
 - + Intra-Area Path Selection
 - + Inter-Area Path Selection
 - + External Path Selection

OSPF Areas Overview

- + Areas add hierarchy and scalability to OSPF
- + An area defines a flooding domain
 - + All devices in the area agree on the topology
 - + Changes inside the area require LSA flooding and full SPF
- + Routing between areas hides topology details
 - + Inter-area routing is similar to distance vector
 - + Changes outside the area don't always require LSA flooding or SPF
 - + Limits impact on router resources

OSPF Two-Level Hierarchy

- + Backbone area
 - + Area 0 (0.0.0.0)
 - + Used to summarize topology information between other areas
 - + Traffic from one area to another must pass through area 0
 - + Must be contiguous
- + Non-backbone areas
 - + All other areas 1 – 2^{32} (0.0.0.1 – 255.255.255.255)
 - + Must use connections to area 0 to reach other areas

OSPF Router Types

- + Backbone routers
 - + At least one link in area 0
- + Internal routers
 - + All links in one non-backbone area

OSPF Router Types (cont.)

- + Area Border Router (ABR)
 - + Links in both area 0 and in non-backbone area(s)
 - + Used to summarize information between area 0 and non-backbone area
- + Autonomous System Boundary Router (ASBR)
 - + At least one link in the OSPF domain
 - + At least one link outside the OSPF domain
 - + EIGRP, IS-IS, BGP, etc.
 - + Used to redistribute information to/from other routing domains and OSPF

OSPF LSA Types

- + With different router types in the OSPF domain, different types of advertisements are required
 - + e.g DR, ABR, ASBR, etc.
- + Different LSA formats used to represent this information
 - + Format is defined by type code
 - + Type 1, type 2, etc.
- + Which LSA types are sent and received depends on
 - + Router's type
 - + OSPF network type
 - + Area type

OSPF LSA Types (cont.)

- + LSA types are...
 - + Type 1 – Router LSA
 - + Type 2 – Network LSA
 - + Type 3 – Network Summary LSA
 - + Type 4 – ASBR Summary LSA
 - + Type 5 – External LSA
 - + Type 7 – NSSA External LSA

OSPF LSA Types (cont.)

- + Other types exist outside our scope
- + Type 6 – Multicast LSA
 - + MOSPF not implemented by most vendors
- + Types 8, 9, 10 – Opaque LSA
 - + Used for extensibility
 - + E.g. MPLS Traffic Engineering

LSAs and Route Types

- + LSAs are grouped together by 3 route types...
- + Intra-Area Routes (O)
 - + LSA Types 1 & 2
- + Inter-Area Routes (O IA)
 - + LSA Types 3 & 4
- + External Routes
 - + E1/E2
 - + LSA Type 5
 - + N1/N2
 - + LSA Type 7

LSA Type 1 – Router LSA

- + Generated by every router in the OSPF domain
 - + Not flooded outside the area they originate in
- + Describes its directly connected links
 - + What are my link costs
 - + Who are my neighbors
- + Used to build the graph for intra-area SPF

LSA Type 2 – Network LSA

- + Generated by the Designated Router (DR) on broadcast and non-broadcast network types
 - + Not flooded outside the area they originate in
- + Describes who is adjacent with the DR
 - + What is my link cost to the DR
 - + Implies my link cost to all others adjacent to that DR
- + Used to reduce redundant information in the database
 - + $n*(n-1)/2$ and flooding scalability issue

LSA Type 3 – Network Summary LSA

- + Generated by ABR
 - + Flooded from area 0 into non-backbone areas and vice-versa
- + Describes ABR's reachability to links in other areas
 - + Includes cost, but hides ABR's actual path to destination
- + SPF not run for ABR advertised routes
 - + ABR can reach link A via SPT in cost X
 - + I can reach ABR via SPT in cost Y
 - + Implies I can reach link A via SPT in cost $X + Y$
- + This is why inter-area routing is like distance vector

LSA Type 4 – ASBR Summary LSA

- + Generated by ABR
 - + Flooded from area 0 into non-backbone area and vice-versa
- + Describes ABR's reachability to ASBRs in other areas
 - + Includes cost, but hides ABR's actual path to destination
- + SPF not run to reach inter-area ASBR
 - + ABR can reach ASBR via SPT in cost X
 - + I can reach ABR via SPT in cost Y
 - + Implies I can reach ASBR via SPT in cost $X + Y$
- + Inter-area external routing is also like distance vector

LSA Type 5 – External LSA

- + Generated by ASBR
 - + Flooded to all non-stub areas
- + Describes routes ASBR is redistributing
 - + Metric
 - + Metric Type
 - + Type 1 = E1
 - + Type 2 = E2 (default)
 - + Forward Address
 - + Who should I route towards to reach the link?
 - + Usually the ASBR itself, but could be someone else in some designs
 - + Route Tag

OSPF External Type 1 vs. Type 2

- + External route type controls how metric for external link is calculated
- + Type 1 (E1)
 - + Take the cost the ASBR reports in plus the cost to the ASBR
- + Type 2 (E2)
 - + Take just the cost the ASBR reports in
 - + If there is a tie, then take the cost to the ASBR as well
- + Type 1 is preferred over Type 2
 - + More on manipulating path selection later...

OSPF External Route Calculation

- + Performs like distance vector routing similar to inter-area calculation
- + Intra-area externals
 - + ASBR can reach link A in cost X
 - + I can reach ASBR via SPT in cost Y
 - + I can reach link A via SPT in cost $X + Y$
- + Inter-area externals
 - + ASBR can reach link A in cost X
 - + ABR can reach ASBR via SPT in cost Y
 - + I can reach ABR via SPT in cost Z
 - + I can reach link A via SPT in cost $X + Y + Z$

LSA Type 7 – NSSA External LSA


- + Generated by ASBR inside a Not-So-Stubby-Area
 - + More on this later...





OSPF

OSPF Media Dependencies & Network Types



In This Section

- + OSPF Network Types
- + DR/BDR Election Process
- + OSPF Next Hop Processing
- + OSPF Network Type Design Issues

OSPF Media Dependencies

- + OSPF's behavior changes depending on what type of media it is configured on
 - + e.g. Ethernet vs. Frame Relay vs. PPP
- + OSPF defines different “network types” to deal with different media characteristics
- + OSPF network types control...
 - + How are hellos & updates sent?
 - + Who forms adjacency?
 - + How is the next-hop calculated?

OSPF Network Types

- + Broadcast
- + Non-Broadcast
- + Point-to-Point
- + Point-to-Multipoint
- + Point-to-Multipoint Non-Broadcast
- + Loopback

OSPF Network Types & Forming Adjacencies

- + OSPF network type does not need to match to form adjacency...
 - + But they do need to be compatible
 - + Other attributes must still match
 - + E.g. timers
- + What makes the network types compatible?
 - + Usage of Type 2 LSA

OSPF Type 2 LSA Review

- + LSA Type 2 – Network LSA
- + Generated by the Designated Router
 - + Describes who is adjacent with DR
 - + Not flooded outside the area they originate in
- + Used to optimize OSPF operation on a shared segment
 - + Reduce number of OSPF adjacencies
 - + Reduce LSA flooding replication
 - + Simplify SPF calculation

OSPF Network Type Compatibilities

- + Network types that use Type 2 LSA
 - + Broadcast
 - + Non-Broadcast
- + Network types that *do not* use Type 2 LSA
 - + Point-to-Point
 - + Point-to-Multipoint
 - + Point-to-Multipoint Non-Broadcast

OSPF Network Broadcast

- + Default on multi-access broadcast medias
 - + Ethernet
 - + Token Ring
 - + FDDI
- + Sends hellos and updates as multicast
 - + 224.0.0.5 (AllSPFRouters)
 - + 224.0.0.6 (AllDRouters)
- + Uses DR & BDR
 - + I.e. uses Type 2 LSA

OSPF Network Non-Broadcast

- + Default on multipoint NBMA medias
 - + Frame Relay & ATM
- + Sends hellos as unicast
 - + Manually defined addresses with **neighbor** command
- + Uses DR & BDR
 - + I.e. uses Type 2 LSA

How The DR & BDR Work

- + Designated Router (DR)
 - + Forms adjacency with all routers on the link
 - + Listens for LSUs (224.0.0.6)
 - + Re-floods LSUs back to the segment (224.0.0.5)
 - + Does not modify next-hop value
- + Backup Designated Router (BDR)
 - + Used for redundancy of DR
 - + Does not re-flood LSUs

How DROthers Work

- + DROthers
 - + All other routers on link
 - + I.e. not the DR or BDR
 - + Form FULL adjacency with DR & BDR
 - + Stop at 2-Way adjacency with each other

DR / BDR Election

- + DR / BDR chosen through election process
- + Election based on interface priority and Router-ID
 - + Priority
 - + 0 – 255
 - + Higher better
 - + 0 = never
 - + Router-ID
 - + Highest loopback / interface IP
 - + Can be statically set
 - + Higher better
- + Uses WAIT timer to stop pre-emption of current DR/BDR
 - + Unlike IS-IS's Designated Intermediate System (DIS)

OSPF Network Point-to-Point

- + Default on point-to-point medias
 - + E.g. HDLC, PPP, GRE
- + Sends hellos as multicast
 - + 224.0.0.5
- + No DR/BDR Election
- + Supports only two neighbors on the link

OSPF Network Point-to-Multipoint

- + Treat network as a collection of P2P links
- + Sends hellos as multicast
 - + 224.0.0.5
- + No DR/BDR Election
- + Special next-hop processing
- + Usually the best design option for partial mesh NBMA networks

Point-to-Multipoint Non-Broadcast

- + Same as point-to-multipoint, but sends hellos as unicast
 - + Manually defined addresses with **neighbor** command
 - + Allows for per-VC OSPF cost over NBMA
- + No DR/BDR Election
- + Special next-hop processing

OSPF Network Loopback

- + Special case for Loopback and Looped-back interfaces
- + Advertises link as /32 stub host route
- + **ip ospf network point-to-point** used to disable this behavior





OSPF

Configuring OSPF Network Types



In This Section

- + Modifying OSPF DR/BDR Election
- + Modifying OSPF Network Types
- + Verifying the OSPF Database
- + Verifying Next Hop Processing
- + Verifying OSPF Network Type Design Issues





OSPF

OSPF Virtual Links



In This Section

- + OSPF Discontiguous Areas
- + OSPF Virtual Links
- + Path Selection with OSPF Virtual Links

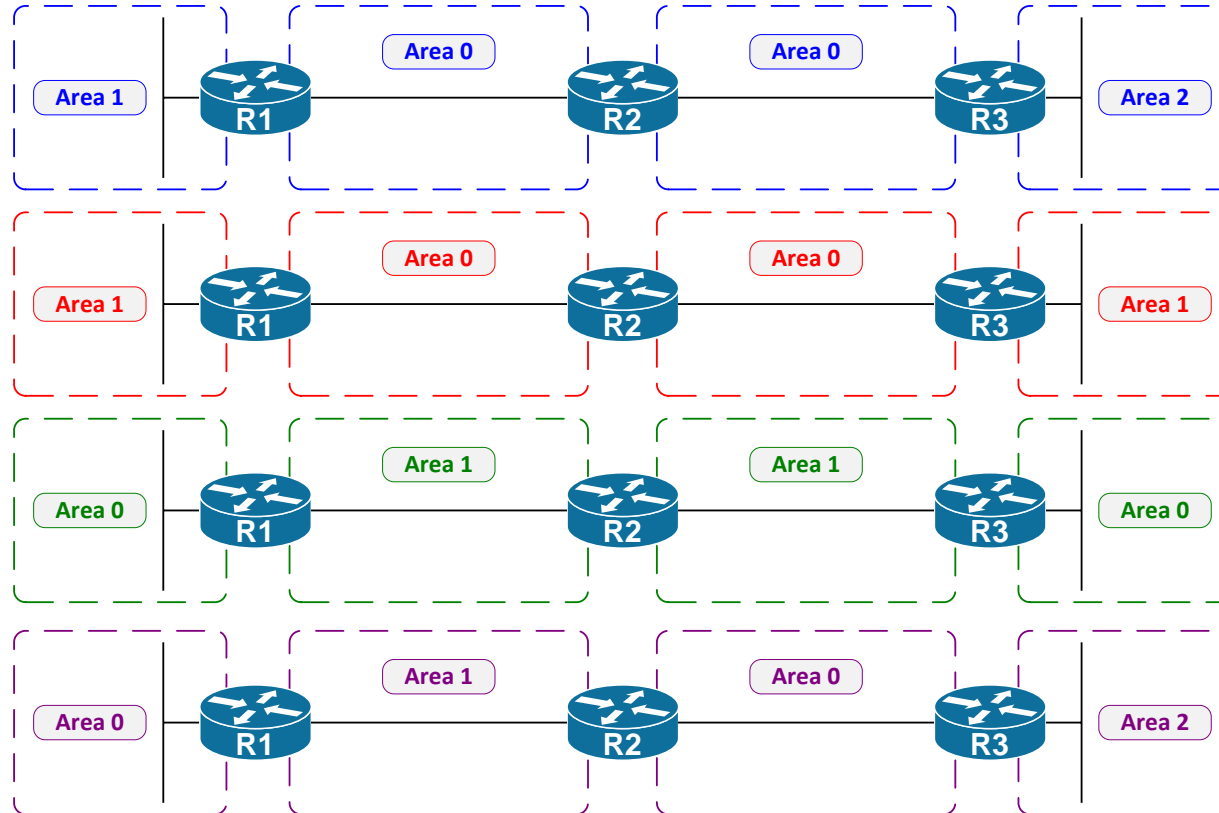
OSPF Inter-Area Routing Review

- + OSPF Inter-Area routing relies on the Network Summary LSA
- + LSA Type 3 – Network Summary LSA
 - + Generated by Area Border Routers (ABRs)
 - + Flooded from area 0 into non-transit area and vice-versa
 - + Describes ABR's reachability to links in other areas

OSPF Discontiguous Area Problems

- + SPF not run to reach an ABR's advertised Inter-Area routes
 - + Use the Intra-Area SPT to the ABR
 - + Add the ABR's cost to my SPT cost
- + What if the ABR isn't reachable via an SPT?
 - + Certain states cause ABRs to be unreachable
 - + These states are called "discontiguous" areas or discontiguous area 0

Which Following Designs Are (In)valid?



Repairing Discontiguous Areas

- + Broken area designs are fixed by adding new area 0 links & adjacencies
 - + Links could be either physical or virtual
 - + E.g. a GRE tunnel
- + OSPF Virtual Links are a form of virtual area 0 adjacencies

OSPF Virtual Links

- + How OSPF Virtual Links work
 - + Used to form multi-hop unicast area 0 adjacency
 - + Follows the already built SPT between ABRs to heal the database

OSPF Virtual Links Caveats

- + Endpoints must be reachable via a normal area
 - + I.e. not a stub area
- + Transit area must not have filtering applied
 - + I.e. LSA 3 filters, distribute lists, etc.
- + Inherits cost from SPT cost between endpoints
 - + Cost must be below 65535 (0xffff)
- + Runs as a demand circuit
 - + Errors in config could be hidden until flooding occurs





OSPF

OSPF Stub Areas



In This Section

- + Scaling OSPF with Stub Areas
- + OSPF Stub Area Variations
- + Type 7 LSA

OSPF Scalability

- + Scaling OSPF is a function of two variables...
- + How complex is the topology graph?
 - + i.e. how many routers are in the area
 - + large flooding domain means lots of SPF runs
- + How much reachability information is there?
 - + i.e. how many routes are being advertised
 - + large routing tables means it takes longer to flood
- + Scalability is achieved by minimizing these two

Topology vs. NLRI Summarization

- + Topology summarization achieved through OSPF Areas
 - + Hide the details of how the graph looks in other areas
 - + Only run SPF for intra-area destinations
 - + Areas don't hide reachability information though
- + NLRI summarization reduces the number of routes
 - + Take multiple longer match prefixes and combine them into smaller shorter matches

OSPF NLRI Summarization

- + OSPF summarization is implemented two ways
- + Per-prefix summarization
 - + e.g. two routes 100.0.0.0/16 and 100.1.0.0/16 become one route 100.0.0.0/15
- + Per-LSA summarization
 - + E.g. remove all Inter-Area routes and replace them with the shortest match possible, a default route
 - + This is what Stub Areas do

How OSPF Stub Areas Work

- + Filtering is enforced at common transit point of the OSPF topology
 - + i.e. the ABR
- + ABR controls which LSAs enter the area
 - + Type-3, Type-4, and/or Type-5 are filtered depending on stub type
- + Reachability information removed is then replaced with a default route
 - + Still allows reachability to removed routes (in most cases)
- + All routers in the area must agree on the stub flag
 - + Part of adjacency negotiation

OSPF Stub Area Types

- + Four stub area types control which routes (LSAs) can enter the area
- + Stub Area
 - + Stops external routes
- + Totally Stubby Area
 - + Stops inter-area and external routes
- + Not-So-Stubby Area (NSSA)
 - + Stops external routes, but allows local redistribution
- + Not-So-Totally-Stubby Area
 - + Stops inter-area and external routes, but allows local redistribution

OSPF Stub Areas

- + Stub Area logic
 - + I know how to get to my ABR
 - + My ABR knows how to get to the ASBRs
 - + The ASBRs knows how to get to the external routes
 - + If I default to the ABR, I don't need the specific external routes
- + Stub Area result
 - + ABR removes LSAs 4 (ASBR) & 5 (External)
 - + ABR originates default route

OSPF Totally Stubby Areas

- + Totally Stubby Area logic
 - + I know how to get to my ABR
 - + My ABR knows how to get to other areas and to the ASBRs
 - + The ASBRs knows how to get to the external routes
 - + If I default to the ABR, I don't need the specific inter-area or external routes
- + Totally Stubby Area result
 - + ABR removes LSAs 3 (Inter-Area), 4 (ASBR), & 5 (External)
 - + ABR originates default route

OSPF Not-So-Stubby Areas (NSSA)

- + NSSA logic
 - + Stub areas block external routes, but what if need to redistribute into the stub area?
 - + Filter like a stub area, but make an exception for local redistribution
- + NSSA Result
 - + Redistributing router generates NSSA External (LSA Type 7)
 - + ABR changes NSSA External into External into area 0
 - + ABR removes LSAs 4 (ASBR) & 5 (External)
 - + ABR does not automatically originate default route

LSA Type 7 – NSSA External LSA

- + Generated by ASBR inside NSSA
 - + Flooded only within NSSA
 - + Changed into Type 5 LSA by an ABR as it leaves the area
- + Describes routes ASBR is redistributing
 - + Metric
 - + Metric Type
 - + Type 1 = N1
 - + Type 2 = N2 (default)
 - + Forward Address
 - + Who should I route towards to reach the link?
 - + Usually the ASBR itself, but could be someone else in some designs
 - + Route Tag

OSPF Not-So-Totally Stubby Areas

- + Not-So-Totally Stubby Area logic
 - + Totally Stubby areas block inter-area and external routes, but what if I need to redistribute into the totally stubby area?
 - + Combine Totally Stubby and NSSA behaviors
- + Not-So-Totally Stubby Area result
 - + Redistributing router generates NSSA External (LSA Type 7)
 - + ABR changes NSSA External into External into area 0
 - + ABR removes LSAs 3 (Inter-Area), 4 (ASBR), & 5 (External)
 - + ABR originates default route





OSPF

Configuring OSPF Stub Areas



In This Section

- + Configuring OSPF Stub Areas
- + Stub Areas & Adjacency Problems
- + OSPF Database Verification with Stub Areas





OSPF

Traffic Engineering with OSPF Stub Areas



In This Section

- + OSPF Stub Areas with Multiple ABR Exit Points
- + Traffic Engineering with Stub Areas
 - + Longest Match Routing
 - + Area Default Cost





OSPF

Configuring OSPF Not so Stubby Areas (NSSA)



In This Section

- + Configuring OSPF Not-So-Stubby-Areas
- + Controlling NSSA Redistribution





OSPF

Default Routing with OSPF NSSA



In This Section


- + OSPF NSSA Default Routing
- + OSPF Route Type Preference





OSPF

OSPF NSSA Translator Election & Forwarding Address



In This Section

- + OSPF NSSA Type 7 to Type 5 Translator Election
- + OSPF Forwarding Address





OSPF

OSPF Path Selection



In This Section

- + OSPF Cost Calculation
- + OSPF Path Selection Order
- + Modifying OSPF Path Selection

OSPF Path Selection

- + Router LSAs include cost attribute for each link
 - + Value of 0 - 65535 (0xffff)
- + Best path to that link is lowest end-to-end cost
 - + I.e. the SPT
- + Cisco's implementation uses bandwidth based cost, but per RFC it is arbitrary
 - + Default Cisco Cost = $100\text{Mbps} / \text{Link Bandwidth}$
 - + Reference bandwidth can be modified to accommodate higher speed links (e.g. TenGigabitEthernet)

OSPF Path Selection Order

- + Per RFC, OSPF path selection state machine prefers...
 - + Intra Area Routes (O)
 - + Inter Area Routes (O IA)
 - + External Type 1 (E1)
 - + NSSA Type 1 (N1)
 - + External Type 2 (E2)
 - + NSSA Type 2 (N2)
- + Cannot be modified with metric or distance

Modifying OSPF Path Selection

- + OSPF uses bandwidth based cost
 - + $\text{COST} = \text{Reference_BW} / \text{Interface_BW}$
- + Cost can be modified by changing...
 - + Interface **bandwidth**
 - + Interface **ip ospf cost**
 - + Process **auto-cost**
 - + Process **neighbor cost**

OSPF Virtual Link Cost

- + Virtual Links inherit their cost from SPT cost between the Virtual Link endpoints
 - + SPT cost may exceed maximum link cost
- + Virtual Link must have cost below 65535 to initialize
 - + E.g. transiting 2 links of cost 50,000 isn't valid
- + Could occur if reference bandwidth is high and Virtual Link transits legacy links
 - + E.g reference BW is 40Gbps and VL transits a T1





OSPF

OSPF Summarization Overview



In This Section

- + Scaling OSPF with Summarization
- + Internal vs. External OSPF Summarization
- + Summary Discard Route

OSPF Scalability

- + Scaling OSPF is achieved two ways...
- + Summarize the topology graph
 - + Achieved through OSPF Areas
 - + Hides the details of how the graph looks in other areas
 - + Only run SPF for intra-area destinations
- + Summarize the NLRI
 - + Achieved through Stub Areas & per-prefix summaries
 - + Replace multiple longer match routes with a shorter match

OSPF NLRI Summarization

- + All devices in same area must have same LSDB
 - + Implies summarization can't be performed at arbitrary points in the topology
- + OSPF summarization can only occur...
 - + Between internal areas
 - + Between external domains
- + EIGRP & BGP win over OSPF & IS-IS here
 - + EIGRP & BGP hierarchy is arbitrary
 - + OSPF & IS-IS is strict leaf – spine – leaf

OSPF Internal vs. External Summaries

- + Internal Summarization
 - + Summarizes Type-1 into Type-3 LSAs
 - + Performed on ABRs
- + External Summarization
 - + Summarizes Type-5 into Type-5 LSAs
 - + Summarizes Type-7 into Type-7 LSAs
 - + Performed on ASBRs

Summary Discard Route

- + When summarizing, OSPF process automatically creates a local “discard” route
 - + I.e. route to Null0
- + Goal is to drop traffic if longest match is summary
 - + I.e. if you are summarizing it, you should always have a more specific route to it
 - + End result is that summary router cannot fallback to default
- + Can be disabled with **no discard-route**





OSPF

Configuring OSPF Summarization



In This Section

- + Configuring OSPF Summarization
- + Where OSPF Can Summarize
- + Other OSPF Summary Applications

Configuring OSPF Summarization

- + Internal Summaries
 - + **area x range**
 - + X is the source area
- + External Summaries
 - + **summary-address**

Where OSPF Can Summarize

- + Internal Summaries
 - + Only at the ABR who knows the LSA-1
 - + E.g. can summarize LSA-1 to LSA-3 but not LSA-3 to LSA-3
- + External Summaries
 - + Only at the ASBR who is the originator
 - + ASBR performing Type-5 or Type-7 redistribution
 - + ASBR/ABR performing Type-5 to Type-7 translation

Other Summary Applications

- + OSPF NLRI summary can also be used to...
- + Traffic Engineer
 - + Prefer longer match over shorter match
- + Filter routes
 - + **area range not-advertise**
 - + **summary-address not-advertise**
- + Enforce area-local scope for NSSA routes
 - + **summary-address nssa-only**





OSPF

OSPF Authentication



In This Section

- + OSPF Authentication
- + Supported OSPF Authentication Types
- + Configuring OSPF Authentication
- + Troubleshooting OSPF Authentication

OSPF Authentication

- + OSPF supports adjacency authentication to protect control plane
 - + E.g. prevent against routing injection attack
- + Every OSPF packet header includes authentication information
 - + I.e. Hello, LSU, LSR, etc.
- + Authentication does not mean encryption
 - + OSPFv2 payload is still clear text
 - + OSPFv3 supports IPsec encryption

OSPF Authentication Types

- + Three types of authentication
 - + Type 0 – Null
 - + Type 1 – Simple Password
 - + Type 2 – Cryptographic (MD5/SHA)

Implementing OSPF Authentication

- + OSPF authentication can be enabled on...
 - + OSPF process level
 - + **area [area-id] authentication...**
 - + Link level
 - + **ip ospf authentication...**
 - + Link level overrides process level
- + Password always configured on the link
 - + **ip ospf authentication-key...**
 - + **ip ospf message-digest-key...**
- + Key ID's must match for Cryptographic authentication

OSPF Virtual Link Authentication

- + Virtual Link is an Area 0 interface
 - + Implies same inheritance rules of authentication
- + Virtual Link is the interface
 - + Key goes at the interface
 - + Type can be configured globally or at the interface
- + Virtual Links runs as a demand circuit
 - + Always clear the VL after authentication





OSPF

OSPF Authentication Enhancements



In This Section

- + OSPF Authentication Enhancements
 - + OSPF SHA Cryptographic Authentication
 - + OSPF Key Chain Based Authentication

OSPF Authentication Enhancements

- + New enhancements defined in RFC 5709
 - + OSPFv2 HMAC-SHA Cryptographic Authentication
- + Does not define new authentication type
 - + Still 0 (Null), 1 (Simple Password), and 2 (Cryptographic)
- + Defines new algorithms for Type 2
 - + Keyed-MD5
 - + HMAC-SHA-1
 - + HMAC-SHA-256
 - + HMAC-SHA-384
 - + HMAC-SHA-512

OSPF Key Chain Authentication

- + Like EIGRP, OSPF now uses Key Chains
- + Allows for multiple enhancements...
 - + Multiple keys
 - + Automatic time-based key rotation
 - + Single key chain for multiple interfaces
- + Still backwards compatible with interface level MD5
 - + Key numbers must still match





OSPF

OSPF Route Filtering



In This Section

- + OSPF Filtering Overview
- + OSPF Filtering Methods
- + OSPF Filtering Examples

OSPF Filtering Overview

- + OSPF is a Link State IGP
 - + Routers know the entire topology graph within an area
 - + Input to SPF must be equal in the area to result in the same SPT
 - + Implies that filtering can only be applied at certain points

OSPF Route Filtering Methods

- + OSPF filtering is normally applied between areas
 - + Filtering with Summarization
 - + LSA Type-3 Filtering
 - + NSSA ABR External Prefix Filtering
 - + Transit prefix suppression
- + OSPF RIB can be filtered anywhere
 - + Distribute-list with ACL
 - + Distribute-list with Route-Map
 - + Administrative Distance
- + RIB filtering can be dangerous
 - + Does not stop the flooding of LSAs within the area

