Alcatel·Lucent

Alcatel-Lucent 7705

SERVICE AGGREGATION ROUTER OS | RELEASE 3.0

ROUTING PROTOCOLS GUIDE

Alcatel-Lucent assumes no responsibility for the accuracy of the information presented, which is subject to change without notice.

Disclaimers

Alcatel-Lucent products are intended for commercial uses. Without the appropriate network design engineering, they must not be sold, licensed or otherwise distributed for use in any hazardous environments requiring fail-safe performance, such as in the operation of nuclear facilities, aircraft navigation or communication systems, air traffic control, direct life-support machines, or weapons systems, in which the failure of products could lead directly to death, personal injury, or severe physical or environmental damage. The customer hereby agrees that the use, sale, license or other distribution of the products for any such application without the prior written consent of Alcatel-Lucent, shall be at the customer's sole risk. The customer hereby agrees to defend and hold Alcatel-Lucent harmless from any claims for loss, cost, damage, expense or liability that may arise out of or in connection with the use, sale, license or other distribution of the products in such applications.

This document may contain information regarding the use and installation of non-Alcatel-Lucent products. Please note that this information is provided as a courtesy to assist you. While Alcatel-Lucent tries to ensure that this information accurately reflects information provided by the supplier, please refer to the materials provided with any non-Alcatel-Lucent product and contact the supplier for confirmation. Alcatel-Lucent assumes no responsibility or liability for incorrect or incomplete information provided about non-Alcatel-Lucent products.

However, this does not constitute a representation or warranty. The warranties provided for Alcatel-Lucent products, if any, are set forth in contractual documentation entered into by Alcatel-Lucent and its customers.

This document was originally written in English. If there is any conflict or inconsistency between the English version and any other version of a document, the English version shall prevail.

When printed by Alcatel-Lucent, this document is printed on recycled paper.

Table of Contents

List of Tables

Getting Started 29

OSPF 33

IS-IS 141

List of Figures

OSPF

List of Acronyms

Acronym	Expansion
2G	second generation wireless telephone technology
3DES	triple DES (data encryption standard)
3G	third generation mobile telephone technology
5620 SAM	5620 Service Aware Manager
7705 SAR	7705 Service Aggregation Router
7710 SR	7710 Service Router
7750 SR	7750 Service Router
9500 MPR	9500 Microwave Packet Radio
A/S	active/standby
ABR	available bit rate area border router
AC	alternating current attachment circuit
ACL	access control list
ACR	adaptive clock recovery
ADP	automatic discovery protocol
AFI	authority and format identifier
AIS	alarm indication signal
ANSI	American National Standards Institute
Apipe	ATM VLL
APS	automatic protection switching
ARP	address resolution protocol
AS	autonomous system
ASAP	any service, any port

Acronym	Expansion
ASBR	autonomous system boundary router
ASN	autonomous system number
ATM	asynchronous transfer mode
ATM PVC	ATM permanent virtual circuit
B3ZS	bipolar with three-zero substitution
Batt A	battery A
B-bit	beginning bit (first packet of a fragment)
Bellcore	Bell Communications Research
BFD	bidirectional forwarding detection
BGP	border gateway protocol
BITS	building integrated timing supply
BOF	boot options file
BRAS	Broadband Remote Access Server
BSC	Base Station Controller
BSTA	Broadband Service Termination Architecture
BTS	base transceiver station
CAS	channel associated signaling
CBN	common bonding networks
CBS	committed buffer space
CC	control channel continuity check
CCM	continuity check message
CE	customer edge circuit emulation
CEM	circuit emulation
CES	circuit emulation services

Acronym	Expansion
CESoPSN	circuit emulation services over packet switched network
CFM	connectivity fault management
CIDR	classless inter-domain routing
CIR	committed information rate
CLI	command line interface
CLP	cell loss priority
CoS	class of service
CPE	customer premises equipment
Cpipe	circuit emulation (or TDM) VLL
CPM	Control and Processing Module (CPM is used instead of CSM when referring to CSM filtering – to align with CLI syntax used with other SR products)
CPU	central processing unit
CRC	cyclic redundancy check
CRON	a time-based scheduling service (from chronos = time)
CSM	Control and Switching Module
CSNP	complete sequence number PDU
CSPF	constrained shortest path first
CV	connection verification customer VLAN (tag)
CW	control word
DC	direct current
DC-C	DC return - common
DCE	data communications equipment
DC-I	DC return - isolated
DCO	digitally controlled oscillator
DDoS	distributed DoS

Acronym	Expansion
DES	data encryption standard
DHCP	dynamic host configuration protocol
DIS	designated intermediate system
DM	delay measurement
DNS	domain name server
DoS	denial of service
dot1p	IEEE 802.1p bits, found in Ethernet or VLAN ingress packet headers and used to map traffic to up to eight forwarding classes
dot1q	IEEE 802.1q encapsulation for Ethernet interfaces
DPLL	digital phase locked loop
DSCP	differentiated services code point
DSL	digital subscriber line
DSLAM	digital subscriber line access multiplexer
DTE	data termination equipment
DU	downstream unsolicited
DV	delay variation
e911	enhanced 911 service
E-bit	ending bit (last packet of a fragment)
ECMP	equal cost multi-path
EFM	Ethernet in the first mile
EGP	exterior gateway protocol
EIA/TIA-232	Electronic Industries Alliance/Telecommunications Industry Association Standard 232 (also known as RS-232)
ELER	egress label edge router

Acronym	Expansion
E&M	ear and mouth earth and magneto exchange and multiplexer
Epipe	Ethernet VLL
ERO	explicit route object
ESD	electrostatic discharge
ESMC	Ethernet synchronization message channel
ETE	end-to-end
ETH-CFM	Ethernet connectivity fault management (IEEE 802.1ag)
EVDO	evolution - data optimized
EXP bits	experimental bits (currently known as TC)
FC	forwarding class
FCS	frame check sequence
FDB	forwarding database
FDL	facilities data link
FEAC	far-end alarm and control
FEC	forwarding equivalence class
FF	fixed filter
FIB	forwarding information base
FIFO	first in, first out
FNG	fault notification generator
FOM	figure of merit
FRR	fast reroute
FTN	FEC-to-NHLFE
FTP	file transfer protocol
GFP	generic framing procedure

Acronym	Expansion
GigE	Gigabit Ethernet
GRE	generic routing encapsulation
GSM	Global System for Mobile Communications (2G)
HCM	high capacity multiplexing
HDB3	high density bipolar of order 3
HEC	header error control
HMAC	hash message authentication code
HSDPA	high-speed downlink packet access
HSPA	high-speed packet access
IANA	internet assigned numbers authority
IBN	isolated bonding networks
ICMP	Internet control message protocol
ICP	IMA control protocol cells
IEEE	Institute of Electrical and Electronics Engineers
IEEE 1588v2	Institute of Electrical and Electronics Engineers standard 1588-2008
IES	Internet Enhanced Service
IETF	Internet Engineering Task Force
IGP	interior gateway protocol
ILER	ingress label edge router
ILM	incoming label map
IMA	inverse multiplexing over ATM
IOM	input/output module
IP	Internet Protocol
IPCP	Internet Protocol Control Protocol
IPIP	IP in IP

Acronym	Expansion
Ipipe	IP interworking VLL
IPoATM	IP over ATM
IS-IS	Intermediate System-to-Intermediate System
IS-IS-TE	IS-IS-traffic engineering (extensions)
ISO	International Organization for Standardization
LB	loopback
lbf-in	pound force inch
LBM	loopback message
LBR	loopback reply
LCP	link control protocol
LDP	label distribution protocol
LER	label edge router
LIB	label information base
LLF	link loss forwarding
LLID	loopback location ID
LM	loss measurement
LSA	link-state advertisement
LSDB	link-state database
LSP	label switched path link-state PDU (for IS-IS)
LSR	label switch router link-state request
LSU	link-state update
LT	linktrace
LTE	line termination equipment
LTM	linktrace message

Acronym	Expansion
LTN	LSP ID to NHLFE
LTR	linktrace reply
MA	maintenance association
MA-ID	maintenance association identifier
MAC	media access control
MBB	make-before-break
MBS	maximum buffer space maximum burst size media buffer space
MBSP	mobile backhaul service provider
MC-MLPPP	multi-class multilink point-to-point protocol
MD	maintenance domain
MD5	message digest version 5 (algorithm)
MDA	media dependent adapter
MDDB	multidrop data bridge
MDL	maintenance data link
ME	maintenance entity
MED	multi-exit discriminator
MEF	Metro Ethernet Forum
MEG	maintenance entity group
MEG-ID	maintenance entity group identifier
MEN	Metro Ethernet network
MEP	maintenance association end point
MFC	multi-field classification
MHF	MIP half function
MIB	management information base

Acronym	Expansion
MIP	maintenance association intermediate point
MIR	minimum information rate
MLPPP	multilink point-to-point protocol
MP	merge point multilink protocol
MP-BGP	multiprotocol border gateway protocol
MPLS	multiprotocol label switching
MPR	see 9500 MPR
MRRU	maximum received reconstructed unit
MRU	maximum receive unit
MSDU	MAC Service Data Unit
MS-PW	multi-segment pseudowire
MTSO	mobile trunk switching office
MTU	maximum transmission unit multi-tenant unit
MW	microwave
N·m	newton meter
NBMA	non-broadcast multiple access (network)
NE	network element
NET	network entity title
NHLFE	next hop label forwarding entry
NHOP	next-hop
NLRI	network layer reachability information
NNHOP	next next-hop
NNI	network-to-network interface
Node B	similar to BTS but used in 3G networks — term is used in UMTS (3G systems) while BTS is used in GSM (2G systems)

Acronym	Expansion
NSAP	network service access point
NSSA	not-so-stubby area
NTP	network time protocol
OAM	operations, administration, and maintenance
OAMPDU	OAM protocol data units
OC3	optical carrier, level 3
ORF	outbound route filtering
OS	operating system
OSI	Open Systems Interconnection (reference model)
OSINLCP	OSI Network Layer Control Protocol
OSPF	Open Shortest Path First
OSPF-TE	OSPF-traffic engineering (extensions)
OSS	operations support system
OSSP	Organization Specific Slow Protocol
PDU	protocol data units
PDV	packet delay variation
PDVT	packet delay variation tolerance
PE	provider edge router
PHB	per-hop behavior
PHY	physical layer
PID	protocol ID
PIR	peak information rate
PLCP	Physical Layer Convergence Protocol
PLR	point of local repair
POP	point of presence
POS	packet over SONET

Acronym	Expansion
PPP	point-to-point protocol
PSN	packet switched network
PSNP	partial sequence number PDU
PTP	precision time protocol performance transparency protocol
PVC	permanent virtual circuit
PVCC	permanent virtual channel connection
PW	pseudowire
PWE	pseudowire emulation
PWE3	pseudowire emulation edge-to-edge
QL	quality level
QoS	quality of service
RADIUS	Remote Authentication Dial In User Service
RAN	Radio Access Network
RBS	robbed bit signaling
RD	route distinguisher
RDI	remote defect indication
RED	random early discard
RESV	reservation
RIB	routing information base
RJ-45	registered jack 45
RNC	Radio Network Controller
RRO	record route object
RS-232	Recommended Standard 232 (also known as EIA/TIA-232)
RSVP-TE	resource reservation protocol - traffic engineering
R&TTE	Radio and Telecommunications Terminal Equipment

Acronym	Expansion
RT	receive/transmit
RTM	routing table manager
RTN	battery return
RTP	real-time protocol
RTU	remote terminal unit
SAA	service assurance agent
SAP	service access point
SAR-8	7705 Service Aggregation Router - 8-slot chassis
SAR-F	7705 Service Aggregation Router - fixed form-factor chassis
SAToP	structure-agnostic TDM over packet
SCADA	surveillance, control and data acquisition
SCP	secure copy
SD	signal degrade
SDH	synchronous digital hierarchy
SDI	serial data interface
SDP	service destination point
SE	shared explicit
SF	signal fail
SFP	small form-factor pluggable (transceiver)
SGT	self-generated traffic
SHA-1	secure hash algorithm
SIR	sustained information rate
SLA	Service Level Agreement
SNMP	Simple Network Management Protocol
SNPA	subnetwork point of attachment
SNTP	simple network time protocol

Acronym	Expansion
SONET	synchronous optical networking
S-PE	switching provider edge router
SPF	shortest path first
SPT	shortest path tree
SR	service router (includes 7710 SR, 7750 SR)
SRLG	shared risk link group
SSH	secure shell
SSM	synchronization status messaging
SSU	system synchronization unit
STM1	synchronous transport module, level 1
SVC	switched virtual circuit
TACACS+	Terminal Access Controller Access-Control System Plus
TC	traffic class (formerly known as EXP bits)
TCP	transmission control protocol
TDM	time division multiplexing
TE	traffic engineering
TFTP	trivial file transfer protocol
TLDP	targeted LDP
TLV	type length value
ToS	type of service
T-PE	terminating provider edge router
TPID	tag protocol identifier
TTL	time to live
TTM	tunnel table manager
U-APS	unidirectional automatic protection switching
UBR	unspecified bit rate

Acronym	Expansion
UDP	user datagram protocol
UMTS	Universal Mobile Telecommunications System (3G)
UNI	user-to-network interface
V.35	V-series Recommendation 35
VC	virtual circuit
VCC	virtual channel connection
VCCV	virtual circuit connectivity verification
VCI	virtual circuit identifier
VID	VLAN ID
VLAN	virtual LAN
VLL	virtual leased line
VoIP	voice over IP
VP	virtual path
VPC	virtual path connection
VPI	virtual path identifier
VPLS	virtual private LAN service
VPN	virtual private network
VPRN	virtual private routed network
VRF	virtual routing and forwarding table
VSE	vendor-specific extension
VSO	vendor-specific option
WCDMA	wideband code division multiple access (transmission protocol used in UMTS networks)
WRED	weighted random early discard
WTR	wait to restore

Preface

About This Guide

This guide describes routing protocols supported by the 7705 SAR and provides configuration and implementation examples.

This guide is organized into functional chapters and provides concepts and descriptions of the implementation flow, as well as Command Line Interface (CLI) syntax and command usage.

Audience

This guide is intended for network administrators who are responsible for configuring the 7705 SAR. It is assumed that the network administrators have an understanding of networking principles and configurations, routing processes, and protocols and standards, including:

- CLI concepts
- Interior Gateway Protocols (IGP)
- Open Shortest Path First (OSPF) routing protocol
- Intermediate system to Intermediate system (IS-IS) routing protocol
- traffic engineering extensions to IGP
- Constrained Shortest Path First (CSPF)
- Border Gateway Protocols (BGP)

List of Technical Publications

The 7705 SAR OS documentation set is composed of the following guides:

- 7705 SAR OS Basic System Configuration Guide

  This guide describes basic system configurations and operations.
- 7705 SAR OS System Management Guide

  This guide describes system security and access configurations as well as event logging and accounting logs.
- 7705 SAR OS Interface Configuration Guide

  This guide describes card and port provisioning.
- 7705 SAR OS Router Configuration Guide

  This guide describes logical IP routing interfaces, IP-based filtering, and routing policies.
- 7705 SAR OS MPLS Guide

  This guide describes how to configure Multiprotocol Label Switching (MPLS), Resource Reservation Protocol for Traffic Engineering (RSVP-TE), and Label Distribution Protocol (LDP).
- 7705 SAR OS Services Guide

  This guide describes how to configure service parameters such as service access points (SAPs), service destination points (SDPs), customer information, and user services.
- 7705 SAR OS Quality of Service Guide

  This guide describes how to configure Quality of Service (QoS) policy management.
- 7705 SAR OS Routing Protocols Guide

  This guide provides an overview of dynamic routing concepts and describes how to configure them.
- 7705 SAR OS OAM and Diagnostics Guide

  This guide provides information on Operations, Administration and Maintenance (OAM) tools.

Technical Support

If you purchased a service agreement for your 7705 SAR router and related products from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance. If you purchased an Alcatel-Lucent service agreement, contact your welcome center at:

Web: http://www.alcatel-lucent.com/support

Getting Started

In This Chapter

This chapter provides process flow information to configure dynamic IP routing protocols.

The 7705 SAR router can function as an LSR (label switch router), allowing it to be used in more complex networks; for example:

- tier-2 aggregation (aggregator of aggregator sites) – traffic is aggregated from other tier-3 7705 SAR nodes (aggregated small cell sites), and this traffic, along with local traffic, is switched to tier-1 SR nodes
- ring-based configurations – multiple tier-3/tier-2 7705 SAR nodes are linked via a ring and an SR tier-2/tier-1 node, which acts as a gateway from the ring to a higher level or directly to the MTSO

➡ **Note:** For information on the 7705 SAR as an LSR, refer to the 7705 SAR OS MPLS Guide.

To switch traffic from one router to another in the network, the 7705 SAR must support IP forwarding. To support these larger and more complex topologies, dynamic routing protocols are provided. In Release 3.0, the 7705 SAR supports OSPF, IS-IS, and BGP as dynamic routing protocols.

Alcatel-Lucent 7705 SAR Routing Configuration Process

Table 1 lists the tasks necessary to configure OSPF, IS-IS, and BGP.

This guide is presented in an overall logical configuration flow. Each section describes a software area and provides CLI syntax and command usage to configure parameters for a functional area.

Table 1: Configuration Process

Area	Task	Chapter
Protocol configuration	Configure OSPF	OSPF on page 33
	Configure IS-IS	IS-IS on page 141
	Configure BGP	BGP on page 227
Reference	List of IEEE, IETF, and other proprietary entities	Standards and Protocol Support on page 347

Notes on 7705 SAR-8 and 7705 SAR-F

The 7705 SAR-8 and the 7705 SAR-F run the same operating system software. The main difference between the products is their hardware platforms.

The 7705 SAR-8 has an 8-slot chassis that supports two CSMs, a Fan module and six adapter cards.

The 7705 SAR-F chassis has a fixed hardware configuration. The 7705 SAR-F replaces the CSM, Fan module, and the 16-port T1/E1 ASAP Adapter card and 8-port Ethernet Adapter card with an all-in-one unit that provides comparable functional blocks, as detailed in Table 2.

The fixed configuration of the 7705 SAR-F means that provisioning the router at the “card slot” and “type” levels is preset and is not user-configurable. Operators begin configurations at the port level.

Note: Unless stated otherwise, references to the terms “Adapter card” and “CSM” throughout the 7705 SAR OS documentation set include the equivalent functional blocks on the 7705 SAR-F.

Table 2: 7705 SAR-8 and 7705 SAR-F Comparison

7705 SAR-8	7705 SAR-F	Notes
CSM	Control and switching functions	The control and switching functions include the console and management interfaces, the alarm and fan functions, the synchronization interfaces, system LEDs, and so on.
Fan module	Integrated with the control and switching functions	
16-port T1/E1 ASAP Adapter card	16 individual T1/E1 ports on the faceplate	The T1/E1 ports on the 7705 SAR-F are equivalent to the T1/E1 ports on the 16-port T1/E1 ASAP Adapter card, except that the 16 T1/E1 ports on the 7705 SAR-F support multiple synchronization sources to support two timing references. On the 7705 SAR-8, the CLI indicates the MDA type for the 16-port T1/E1 ASAP Adapter card as `a16-chds1`. On the 7705 SAR-F, the CLI indicates the MDA type for the 7705 SAR-F ports as `a16-chds1v2`.

Table 2: 7705 SAR-8 and 7705 SAR-F Comparison (Continued)

7705 SAR-8	7705 SAR-F	Notes
8-port Ethernet Adapter card	8 individual Ethernet ports on the faceplate	The –48 VDC versions of the 7705 SAR-8 support two versions of the 8-port Ethernet Adapter card, with version 2 having additional support for Synchronous Ethernet. The Ethernet ports on the 7705 SAR-F are equivalent to the Ethernet ports on version 2 of the 8-port Ethernet Adapter card and support multiple synchronization sources to support two timing references. The +24 VDC version of the 7705 SAR-8 only supports version 2 of the 8-port Ethernet Adapter card. On the 7705 SAR-8, the CLI indicates the MDA type for the 8-port Ethernet Adapter card as `a8-eth` or `a8-ethv2`. On the 7705 SAR-F, the CLI indicates the MDA type for the 7705 SAR-F Ethernet ports as `a8-ethv3`, to distinguish it from the actual version 2 of the 8-port Ethernet Adapter card.
Requires user configuration at card (IOM) and MDA (adapter card) levels	Configuration at card (IOM) and MDA (adapter card) levels is preset and users cannot change these types	

OSPF

In This Chapter

This chapter provides information about configuring the Open Shortest Path First (OSPF) protocol.

Topics in this chapter include:

Overview of OSPF

OSPF (Open Shortest Path First) is an interior gateway protocol (IGP) that is used within large autonomous systems (ASs). An autonomous system is a group of networks and network equipment under a common administration. OSPF is a link-state protocol; each router maintains an identical database (called the link-state database, topological database, or routing information database [RIB]) of the AS, including information about the local state of each router (for example, its usable interfaces and reachable neighbors).

OSPF-TE (OSPF with traffic engineering extensions) is used to advertise reachability information and traffic engineering information such as bandwidth.

OSPF routers exchange status, cost, and other relevant interface information with neighboring routers. The information exchange enables all participating routers to establish their link-state database.

OSPF uses a cost metric that represents the status of the link and the bandwidth of the interface in an algorithm to determine the best route to a destination. The algorithm used is called the SPF (shortest path first) or Dijkstra algorithm. Path selection is based on lowest cost, which might not necessarily be the shortest route but is the best route in regards to bandwidth. Each router applies the algorithm to calculate the shortest path to each destination in the network.

When the best route to a particular destination is determined, the route information is sent to the routing table manager (RTM). The RTM may contain more than one best route to a destination from multiple protocols. Because metrics from different protocols are not comparable, the RTM uses preference to select the best route. The route with the lowest preference value is selected. For more information, see Configuring Route Preferences on page 60.

The best routes from the RTM are then added to the forwarding table (also known as the forwarding database [or FIB]). All forwarding decisions are based on the information in the forwarding database.

The forwarding (or dropping) of packets is controlled by filters applied to the interface and route policies applied to the OSPF protocol. Refer to the 7705 SAR OS Router Configuration Guide for information on filters and route policies.

Alcatel-Lucent’s implementation of OSPF conforms to OSPF Version 2 specifications presented in RFC 2328, *OSPF Version 2*. Routers running OSPF can be enabled with minimal configuration. All default and command parameters can be modified.

The following major OSPF features are supported:

- areas – backbone, stub, and not-so-stubby areas (NSSAs)
- virtual links
- neighbors and adjacencies
- link-state advertisements (LSAs)
- metrics
- authentication
- route redistribution and summarization
- OSPF traffic engineering (TE) extensions (to track and advertise available bandwidth – used by MPLS traffic engineering; that is, RSVP-TE)

OSPF Areas

An autonomous system can be divided into areas, with each area containing a group of networks. An area's topology is concealed from the rest of the AS, which significantly reduces OSPF protocol traffic (LSA updates), simplifies the network topology, and simplifies the routing table by populating it with summarized routes rather than exact routes on each router. This decrease in LSA updates, link-state database size, and CPU time, all required for OSPF route calculations, results in a decrease in route calculation time.

All routers in an area have identical link-state databases for that area.

Areas within the same AS are linked to each other via area border routers (ABRs). An ABR is a router that belongs to more than one area. An ABR maintains a separate topological database for each area it is connected to.

Routing in the AS takes place on two levels, depending on whether the source and destination of a packet reside in the same area (intra-area routing) or different areas (inter-area routing). In intra-area routing, the packet is routed solely on information obtained within the area; that is, routing updates are only passed within the area. In inter-area routing, routing updates are passed between areas.

External routes refer to routing updates passed from another routing protocol into the OSPF domain.

Routers that pass information between an OSPF routing domain and a non-OSPF network are called autonomous system boundary routers (ASBRs).

Backbone Area

Every OSPF system requires a backbone area. The OSPF backbone area is uniquely identified as area 0 and uses the area ID 0.0.0.0. All other areas must be connected to the backbone area, either physically or logically. The backbone distributes routing information between areas. If it is not practical or possible to connect an area to the backbone (see area 0.0.0.5 in Figure 1), the ABRs (routers Y and Z in the figure) must be connected via a virtual link. The two ABRs form a point-to-point-like adjacency across the transit area (area 0.0.0.4).

Figure 1: Backbone Area

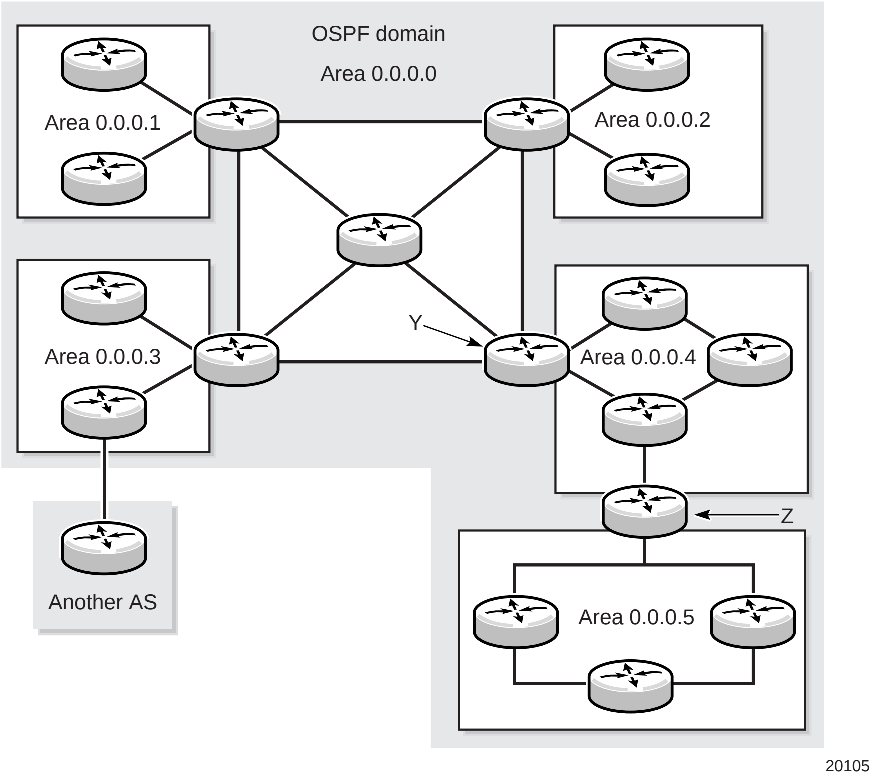

Stub Area

A stub area is a designated area that does not allow external route advertisements and cannot contain ASBRs. Virtual links cannot pass through stub areas.

To route to external destinations, the ABR of the stub area advertises a single default route into the stub area (0.0.0.0). A default route is the network route used by a router when no other known route exists for a given IP packet's destination address. All packets for destinations not known by the router's routing table are sent to the default route and thus out to the network.

This feature reduces the size of the router's database and reduces OSPF protocol traffic, memory usage, and CPU route calculation time.

In Figure 1, areas 0.0.0.1, 0.0.0.2 and 0.0.0.5 could be configured as stub areas.

Not-So-Stubby Area

Another OSPF area type is called a not-so-stubby area (NSSA). NSSAs are similar to stub areas except that limited importing of external routes is allowed. Only routes within the AS are advertised. External routes learned by OSPF routers in the NSSA area are advertised as type 7 LSAs (external route advertisements only within the NSSA area) and are translated by ABRs into type 5 external route advertisements for distribution into other areas of the OSPF domain.

For information on LSA types, see Link-State Advertisements.

An NSSA area cannot be designated as the transit area of a virtual link.

In Figure 1, area 0.0.0.3 could be configured as an NSSA area.

Virtual Links

The backbone area in an OSPF AS must be contiguous and all other areas must be directly connected to the backbone area via an ABR. If it is not practical or possible to physically connect an area to the backbone, virtual links can be used to connect to the backbone through a non-backbone area.

A virtual link functions as a point-to-point link that passes through a transit area. Figure 1 depicts routers Y and Z as the start and end points of the virtual link while area 0.0.0.4 is the transit area. In order to configure virtual links, the router must be an ABR. Virtual links are identified by the router ID of the other endpoint, which is another ABR.

These two endpoint routers must be attached to a common area, called the transit area. The area through which the virtual link passes must have full routing information.

Transit areas pass traffic from an area adjacent to the backbone or to another area. The traffic does not originate or terminate in the transit area. The transit area cannot be a stub area or an NSSA area.

Virtual links are part of the backbone and behave as if they were unnumbered point-to-point networks between the two routers. A virtual link uses the intra-area routing of its transit area to forward packets. Virtual links are brought up and down through the building of the shortest-path trees for the transit area.

Neighbors and Adjacencies

A router uses the OSPF Hello protocol to discover neighbors. Neighbors are routers that interface to a common network. In a broadcast-supported topology, one router sends Hello packets to a multicast address and receives Hello packets in return. Unicast Hello packets are used in non-broadcast topologies.

The neighbors then attempt to form adjacencies by exchanging link-state information with the goal of having identical link-state databases. When the link-state databases of two neighbors are synchronized, they are considered to be adjacent.

Designated Routers and Backup Designated Routers

In multi-access broadcast networks, such as Ethernet networks, with at least two attached routers, a designated router and a backup designated router can be elected. The concept of a designated router was developed in order to avoid the formation of adjacencies between all attached routers. Without a designated router, the area would be flooded with LSAs – a router would send LSAs to all its adjacent neighbors, and each in turn would send LSAs to all their neighbors, and so on. This would create multiple copies of the same LSA on the same link.

The designated router reduces the number of adjacencies required because each router forms an adjacency only with the designated router and backup designated router. Only the designated router sends LSAs in multicast format to the rest of the network, reducing the amount of routing protocol traffic and the size of the link-state database. If the designated router fails, the backup designated router becomes active.

The designated router is automatically elected based on priority – the router with the highest priority becomes the designated router and the router with the second-highest priority becomes the backup. If two routers have the same priority, the one with the highest router ID wins.

A router with a priority set to 0 can never become a designated router.

After a designated router is elected, it begins sending Hello packets to all other attached routers in order to form adjacencies.

Notes:

- In point-to-point networks, where a single pair of routers are connected, no designated or backup designated router is elected. An adjacency must be formed with the neighbor router.
- To significantly improve adjacency forming and network convergence, a network should be configured as point-to-point if only two routers are connected, even if the network is a broadcast media such as Ethernet.

Link-State Advertisements

Link-state advertisements (LSAs) describe the state of a router or network, including router interfaces and adjacency states. Each LSA is flooded throughout an area. The collection of LSAs from all routers and networks form the protocol's link-state (or topological) database.

The distribution of topology database updates takes place along adjacencies. A router sends LSAs to advertise its state according to the configured interval and when the router's state changes. These packets include information about the router's adjacencies, which allows detection of non-operational routes.

When a router discovers a routing table change or detects a change in the network, link-state information is advertised to other routers to maintain identical routing tables. Router adjacencies are reflected in the contents of its link-state advertisements. The relationship between adjacencies and the link states allow the protocol to detect non-operating routers. Link-state advertisements flood the area. The flooding mechanism ensures that all routers in an area have the same topological database. The database consists of the collection of LSAs received from each router belonging to the area.

OSPF sends only the changed information, not the whole topology information or whole link-state database, when a change takes place. From the topological database, each router constructs a tree of shortest paths with itself as root (that is, runs the Dijkstra algorithm). OSPF distributes routing information between routers belonging to a single AS.

Table 3 lists the types of LSAs generated by routers.

Table 3: LSA types

LSA	Definition
Type 1 - Router	Router link advertisements generated by each internal router for each area it belongs to LSAs are flooded only in the area in which they were originated Router LSAs list all the router's links and the state and cost of the links
Type 2 - Network	Network link advertisements generated by designated routers describing the set of routers attached to a particular network LSAs are flooded only in the area of the router that originated them Network LSAs list all attached routers, including the designated router
Type 3 - Network Summary	Summary link advertisements generated by ABRs describing inter-area routes (areas within the AS but outside the area they are sent into) LSAs let internal routers know which destinations can be reached by the ABR LSAs are sent in both directions – into a non-zero area and into the backbone area
Type 4 - ASBR Summary	Summary link advertisements generated by ABRs indicating the location of ASBRs An ABR generates a type 4 LSA after receiving a type 5 LSA from an ASBR
Type 5 - AS External	Generated by an ASBR and describes destinations external to the AS or a default route external to the AS LSAs are flooded to all areas except stub areas
Type 6 - Group membership	Group membership link entry generated by multicast OSPF routers Not applicable in this release
Type 7 - NSSA External	NSSA external routes generated by an ASBR and used by the NSSA to import external routes into a stub area LSAs are flooded only to the NSSA The ABR converts type 7 LSAs into type 5 LSAs before flooding them into the backbone, where they are then flooded to all areas except stub areas

Metrics

In OSPF, all interfaces have a cost value or routing metric used in the OSPF link-state calculation. A metric value is configured based on hop count, bandwidth, or other parameters, to compare different paths through an AS. OSPF uses cost values to determine the best path to a particular destination – the lower the cost value, the more likely the interface will be used to forward data traffic.

Costs are also associated with externally derived routing data, such as those routes learned from an Exterior Gateway Protocol (EGP), for example, BGP, and are passed transparently throughout the AS. This data is kept separate from the OSPF protocol's link-state data. Each external route can be tagged by the advertising router, enabling the passing of additional information between routers on the boundaries of the AS.

Authentication

All OSPF protocol exchanges can be authenticated. This guarantees that only trusted routers can participate in autonomous system routing. Alcatel-Lucent's implementation of OSPF supports plain text (simple password) and Message Digest 5 (MD5) authentication.

When authentication is enabled on a link, a text string password must be configured. Neighbor OSPF routers must supply the password in all OSPF packets they send to an interface.

Plain text authentication includes the password in each OSPF packet sent on a link.

MD5 authentication is more secure than plain text authentication. MD5 authentication uses the password as an encryption key. Routers in the same routing domain must be configured with the same key. When the MD5 hashing algorithm is used for authentication, MD5 is used to verify data integrity by creating a 128-bit message digest from the data input that is included in each packet. The packet is transmitted to the router neighbor and can only be decrypted if the neighbor has the correct password.

By default, authentication is not enabled on an interface.

Route Redistribution and Summarization

Route redistribution is the taking of routes from one protocol and sending them to another protocol. The 7705 SAR supports the redistribution of static routes into OSPF. These routes are advertised as type 5 or type 7 LSAs (external routes) and are included in each router's link-state database.

Route redistribution involves the use of routing policies. For information on routing policies, refer to the 7705 SAR OS Router Configuration Guide, "Route Policies".

Route summarization allows an ABR or ASBR to summarize routes with the same prefix into a single route and distribute it to other areas. Routes redistributed into OSPF from static routes can also be summarized.

Route summarization reduces the amount of routing information across areas and the size of routing tables on the routers, thus improving the calculation speed of the routers.

OSPF-TE Extensions

OSPF traffic engineering (TE) extensions enable the 7705 SAR to include traffic engineering information in the algorithm in order to calculate the best route to a destination. The traffic information includes:

- maximum reservable bandwidth
- unreserved bandwidth
- available bandwidth

IP Subnets

OSPF enables the flexible configuration of IP subnets. Each distributed OSPF route has a destination and mask. A network mask is a 32-bit number that indicates the range of IP addresses residing on a single IP network/subnet. This specification displays network masks as hexadecimal numbers; for example, the network mask for a class C IP network is displayed as 0xffffff00. This mask is often displayed as 255.255.255.0.

Two different subnets with the same IP network number might have different masks, called variable-length subnets. A packet is routed to the longest or most specific match. Host routes are considered to be subnets whose masks are all ones (0xffffffff).

For example, for a packet destined for IP address 10.1.1.1, 10.1.1.0/24 is a longer (better) match than 10.1.1.0/16. If both entries are in the routing table, the route designated by 10.1.1.0/24 will be used.

OSPF Instances

A routing instance is a routing entity for a router. In Release 3.0, the 7705 SAR supports the default routing instance only; it does not support multiple instances. The default routing instance is associated with the global routing table.

Bidirectional Forwarding Detection (BFD) for OSPF

BFD is a simple protocol for detecting failures in a network. BFD uses a “hello” mechanism that sends control messages periodically to the far end and receives periodic control messages from the far end. BFD can detect device, link, and protocol failures.

When BFD is enabled on an OSPF interface, the state of the interface is tied to the state of the BFD session between the local node and remote (far-end) node. BFD is implemented in asynchronous mode only, meaning that neither end responds to control messages; rather, the messages are sent in the time period configured at each end.

If the configured number of consecutive BFD missed messages is reached, the link is declared down and OSPF takes the appropriate action (for example, generates an LSA update against the failed link or reroutes around the failed link).

Due to the lightweight nature of BFD, it can detect failures faster than other detection protocols, making it ideal for use in applications such as mobile transport.

Preconfiguration Requirements

The router ID must be available before OSPF can be configured. The router ID is a 32-bit IP address assigned to each router running OSPF. This number uniquely identifies the router within an AS. OSPF routers use the router IDs of the neighbor routers to establish adjacencies. Neighbor IDs are learned when Hello packets are received from the neighbor.

Before configuring OSPF parameters, ensure that the router ID is derived by one of the following methods:

- define the value using the `config>router` *`router-id`* command
- define the system interface using the `config>router>interface` *`ip-int-name`* command (used if the router ID is not specified with the `config>router router-id` command)

  A system interface must have an IP address with a 32-bit subnet mask. The system interface is assigned during the primary router configuration process when the interface is created in the logical IP interface context.
- if you do not specify a router ID, the last 4 bytes of the MAC address are used

OSPF Configuration Process Overview

Figure 2 displays the process to provision basic OSPF parameters.

Figure 2: OSPF Configuration Process

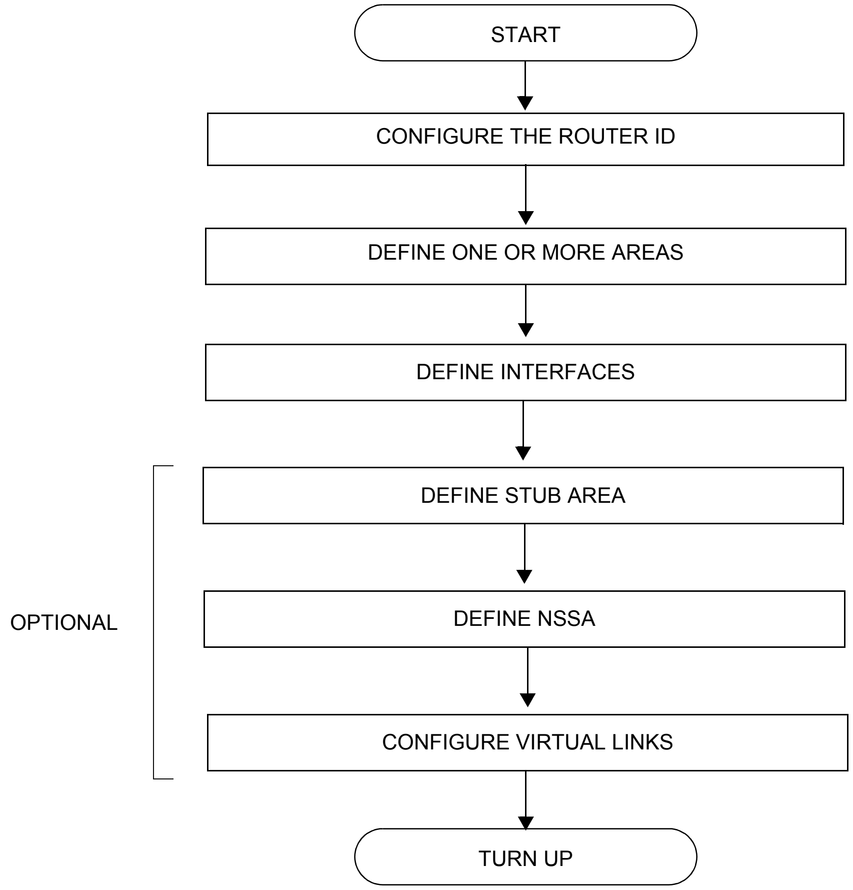

Configuration Notes

General

- Before OSPF can be configured, the router ID must be configured.
- The basic OSPF configuration includes at least one area and an associated interface.
- All default and command parameters can be modified.
- By default, a router has no configured areas.
- The base OSPF instance is created in the administratively enabled state.

Reference Sources

For information on supported IETF drafts and standards, as well as standard and proprietary MIBs, refer to Standards and Protocol Support on page 347.

Configuring OSPF with CLI

This section provides information to configure the Open Shortest Path First (OSPF) protocol using the command line interface.

Topics in this section include:

OSPF Configuration Guidelines

Configuration planning is essential to organize routers, backbone, non-backbone, stub, NSSA areas, and transit links. OSPF provides essential defaults for basic protocol operability. You can configure or modify most commands and parameters.

The minimal OSPF parameters that are necessary to deploy OSPF are:

- router ID

  Each router running OSPF must be configured with a unique router ID. The router ID is used by the OSPF routing protocol to establish adjacencies.

  If a new router ID is defined, the OSPF protocol is not automatically restarted with the new ID. The router must be shut down and restarted in order to initialize the new router ID.
- area

  At least one OSPF area must be created. An interface must be assigned to each OSPF area.
- interfaces

  An interface is the connection between a router and one of its attached networks. An interface has state information associated with it, which is obtained from the underlying lower-level protocols and the routing protocol. An interface to a network has associated with it a single IP address and mask (unless the network is an unnumbered point-to-point network). An interface is sometimes also referred to as a link.

Basic OSPF Configuration

This section provides information to configure OSPF as well as configuration examples of common configuration tasks.

The minimal OSPF parameters that need to be configured are:

- a router ID
- one or more areas
- interfaces `(interface "system")`

The following is an example of a basic OSPF configuration:

```
ALU-A>config>router>ospf# info
----------------------------------------------
            area 0.0.0.0
                interface "system"
                exit
            exit
            area 0.0.0.20
                nssa
                exit
                interface "to-104"
                    priority 10
                exit
            exit
            area 0.0.1.1
            exit
----------------------------------------------
ALU-A>config>router>ospf#
```

Configuring the Router ID

The router ID uniquely identifies the router within an AS. In OSPF, routing information is exchanged between autonomous systems, which are groups of networks that share routing information. The router ID can be set to be the same as the system interface address (loopback address). This is the default setting.

The router ID is derived by one of the following methods:

- defining the value using the `config>router` *`router-id`* command
- defining the system interface using the `config>router>interface` *`ip-int-name`* command (used if the router ID is not specified with the `config>router router-id` command)
- inheriting the last 4 bytes of the MAC address

When configuring a new router ID, protocols are not automatically restarted with the new router ID. The next time a protocol is initialized, the new router ID is used. To force the new router ID, issue the `shutdown` and `no shutdown` commands for OSPF or restart the entire router.

Use the following CLI syntax to configure a router ID (in the `config>router` context):

CLI Syntax: `router-id` *`ip-address`*

The following displays a router ID configuration example:

```
A:ALU-B>config>router# info
#------------------------------------------
# IP Configuration
#------------------------------------------
        interface "system"
            address 10.10.10.104/32
        exit
        interface "to-103"
            address 10.0.0.104/24
            port 1/1/1
        exit
        router-id 10.10.10.104
...
#------------------------------------------
A:ALU-B>config>router#
```

Configuring an OSPF Area

An OSPF area consists of routers configured with the same area ID. To include a router in a specific area, the common area ID must be assigned and an interface identified.

If your network consists of multiple areas, you must also configure a backbone area (0.0.0.0) on at least one router. The backbone contains the area border routers and other routers not included in other areas. The backbone distributes routing information between areas. To maintain backbone connectivity, there must be at least one interface in the backbone area or a virtual link must be configured to another router in the backbone area.

The minimal configuration must include an area ID and an interface. Modifying other command parameters is optional.

Use the following CLI syntax to configure an OSPF area (in the `config>router` context):

CLI Syntax:

```
ospf
     area area-id
        area-range ip-prefix/mask [advertise|not-
          advertise]
        blackhole-aggregate
```

The following displays an OSPF area configuration example:

```
A:ALU-A>config>router>ospf# info
----------------------------------------------
            area 0.0.0.0
            exit
            area 0.0.0.20
            exit
----------------------------------------------
ALU-A>config>router>ospf#
```

Configuring an Interface

In OSPF, an interface can be configured to act as a connection between a router and one of its attached networks. An interface includes state information that was obtained from underlying lower-level protocols and from the routing protocol itself. An interface to a network is associated with a single IP address and mask (unless the network is an unnumbered point-to-point network). Note that if the address is removed from an interface, all OSPF data for the interface is also removed. If the address is merely changed, the OSPF configuration is preserved.

The `passive` command enables the passive property to and from the OSPF interface where passive interfaces are advertised as OSPF interfaces but do not run the OSPF protocol. By default, only interface addresses that are configured for OSPF are advertised as OSPF interfaces. The passive parameter allows an interface to be advertised as an OSPF interface without running the OSPF protocol. When enabled, the interface will ignore ingress OSPF protocol packets and not transmit any OSPF protocol packets.

Use the following CLI syntax to configure an OSPF interface (in the `config>router` context):

CLI Syntax:

```
ospf
     area area-id
         interface ip-int-name
             advertise-subnet
             authentication-key [authentication-key|hash-
             key] [hash|hash2]
             authentication-type [password|message-digest]
             bfd-enable [remain-down-on-failure]
             dead-interval seconds
             hello-interval seconds
             interface-type {broadcast|point-to-point}
             message-digest-key key-id md5 [key|hash-key]
             [hash|hash2]
             metric metric
             mtu bytes
             passive
```

```
priority number
retransmit-interval seconds
no shutdown
transit-delay seconds
```

The following displays an interface configuration example:

```
A:ALU-49>config>router>ospf# info
----------------------------------------------
            asbr
            overload
            overload-on-boot timeout 60
            traffic-engineering
            export "OSPF-Export"
            exit
            area 0.0.0.0
                virtual-link 1.2.3.4 transit-area 1.2.3.4
                    hello-interval 9
                    dead-interval 40
                exit
                interface "system"
                exit
            exit
            area 0.0.0.20
                stub
                exit
                interface "to-103"
                exit
            exit
            area 0.0.0.25
                nssa
                exit
            exit
            area 1.2.3.4
            exit
----------------------------------------------
A:ALU-49>config>router>ospf#
```

Configuring Other OSPF Components

The following sections show the CLI syntax for:

- Configuring a Stub Area
- Configuring a Not-So-Stubby Area
- Configuring a Virtual Link
- Configuring Authentication
- Assigning a Designated Router
- Configuring Route Summaries
- Configuring Route Preferences

Configuring a Stub Area

Configure stub areas to control external advertisement flooding and to minimize the size of the topological databases on an area's routers. A stub area cannot also be configured as an NSSA. The area ID cannot be 0.0.0.0 – this address is reserved for the backbone area.

By default, summary route advertisements (type 3 LSAs) are sent into stub areas. The **no** form of the summary command disables sending summary route advertisements, and only the default route is advertised by the ABR.

Stub areas cannot be used as transit areas. If the area was originally configured as a transit area for a virtual link, existing virtual links are removed when its designation is changed to NSSA or stub.

Use the following CLI syntax to configure a stub area:

CLI Syntax:

```
ospf
    area area-id
        stub
            default-metric metric
            summaries
```

The following displays a stub configuration example:

```
ALU-A>config>router>ospf>area># info
----------------------------------------------
...
            area 0.0.0.0
            exit
            area 0.0.0.20
                stub
                exit
            exit
```

Configuring a Not-So-Stubby Area

NSSAs are similar to stub areas in that no external routes are imported into the area from other OSPF areas. The major difference between a stub area and an NSSA is that an NSSA can flood external routes that it learns throughout its area and from an area border router to the entire OSPF domain. An area cannot be both a stub area and an NSSA. The area ID cannot be 0.0.0.0 – this address is reserved for the backbone area.

NSSAs cannot be used as transit areas. If the area was originally configured as a transit area for a virtual link, existing virtual links are removed when its designation is changed to NSSA or stub.

Use the following CLI syntax to configure NSSAs:

CLI Syntax:
```
ospf
     area area-id
        nssa
           area-range ip-prefix/mask [advertise|not-
           advertise]
           originate-default-route [type-7]
           redistribute-external
           summaries
```

The following displays an NSSA configuration example:

```
A:ALU-49>config>router>ospf# info
----------------------------------------------
...

            area 0.0.0.25
                nssa
                exit
            exit
----------------------------------------------
A:ALU-49>config>router>ospf#
```

Configuring a Virtual Link

The backbone area (area 0.0.0.0) must be contiguous and all other areas must be connected to the backbone area. If it is not possible or practical to connect an area to the backbone, the area border routers must be connected via a virtual link. Two area border routers will form a point-to-point-like adjacency across the transit area. A virtual link can only be configured while in the context of area 0.0.0.0. The transit area cannot be a stub area or an NSSA.

The `router-id` parameter specified in the `virtual-link` command must be associated with the virtual neighbor; that is, the router ID of the far-end router must be specified, not the local router ID.

Use the following CLI syntax to configure a virtual link:

```
CLI Syntax:  ospf
                 area area-id
                    virtual-link router-id transit-area area-id
                       authentication-key [authentication-key|hash-
                       key] [hash|hash2]
                       authentication-type [password|message-digest]
                       dead-interval seconds
                       hello-interval seconds
                       message-digest-key key-id md5 [key|hash-key]
                       [hash|hash2]
                       retransmit-interval seconds
                       transit-delay
                       no shutdown
```

The following displays a virtual link configuration example:

```
A:ALU-49>config>router>ospf# info
----------------------------------------------
...

            area 0.0.0.0
                virtual-link 1.2.3.4 transit-area 1.2.3.4
                    hello-interval 9
                    dead-interval 40
                exit
            exit
            area 0.0.0.20
                stub
                exit
            exit
            area 0.0.0.25
                nssa
                exit
            exit
            area 1.2.3.4
            exit
----------------------------------------------
A:ALU-49>config>router>ospf#
```

Configuring Authentication

Authentication must be explicitly configured. The following authentication commands can be configured on the interface level or the virtual link level:

- `authentication-key` — configures the password used by the OSPF interface or virtual link to send and receive OSPF protocol packets on the interface when simple password authentication is configured
- `authentication-type` — enables authentication and specifies the type of authentication to be used on the OSPF interface, either password or message digest
- `message-digest-key` — command used when the `message-digest` keyword is selected in the `authentication-type` command. The Message Digest 5 (MD5) hashing algorithm is used for authentication. MD5 is used to verify data integrity by creating a 128-bit message digest from the data input. It is unique to that specific data.

A special checksum is included in transmitted packets and is used by the far-end router to verify the packet by using an authentication key (a password). Routers on both ends must use the same authentication key.

MD5 can be configured on each interface and each virtual link. If MD5 is enabled on an interface, that interface accepts routing updates only if the MD5 authentication is accepted. Updates that are not authenticated are rejected. A router accepts only OSPF packets sent with the same `key-id` value defined for the interface.

If the `hash` parameter is not used in the authentication commands, unencrypted characters can be entered. If the `hash` parameter is used, all keys specified in the command are stored in encrypted format in the configuration file. When the `hash` keyword is specified, the password must be entered in encrypted form. Hashing cannot be reversed. To configure an unhashed key, issue the `no message-digest-key` *`key-id`* command and then re-enter the command without the `hash` parameter.

Use the following CLI syntax to configure authentication:

CLI Syntax:
```
ospf
    area area-id
        interface ip-int-name
            authentication-key [authentication-key|hash-
            key] [hash|hash2]
            authentication-type [password|message-digest]
            message-digest-key key-id md5 [key|hash-key]
            [hash|hash2]
        virtual-link router-id transit-area area-id
            authentication-key [authentication-key|hash-
            key] [hash|hash2]
            authentication-type [password|message-digest]
            message-digest-key key-id md5 [key|hash-key]
            [hash|hash2]
```

The following displays authentication configuration examples:

```
A:ALU-49>config>router>ospf# info
----------------------------------------------
...

            area 0.0.0.40
                interface "test1"
                    authentication-type password
                    authentication-key "3WErEDozxyQ" hash
                exit
            exit
            area 1.2.3.4
            exit
----------------------------------------------
A:ALU-49>config>router>ospf#


A:ALU-49>config>router>ospf# info
----------------------------------------------
...

            area 0.0.0.0
                virtual-link 10.0.0.1 transit-area 0.0.0.1
                    authentication-type message-digest
                    message-digest-key 2 md5 "Mi6BQAFi3MI" hash
                exit
                virtual-link 1.2.3.4 transit-area 1.2.3.4
                    hello-interval 9
                    dead-interval 40
                exit
                interface "system"
                exit
            exit
----------------------------------------------
A:ALU-49>config>router>ospf#
```

Assigning a Designated Router

The designated router is responsible for flooding network link advertisements on a broadcast network to describe the routers attached to the network. A router uses Hello packets to advertise its priority. The router with the highest-priority interface becomes the designated router. If routers have the same priority, the designated router is elected based on the highest router ID. A router with priority 0 is not eligible to be a designated router or a backup designated router. At least one router on each logical IP network or subnet must be eligible to be the designated router. By default, routers have a priority value of 1.

When a router starts up, it checks for a current designated router. If a designated router is present, the router accepts that designated router, regardless of its own priority designation. If the designated and backup designated routers fail, new designated and backup routers are elected according to their priority numbers or router IDs (in case of a priority tie).

Designated routers are used only in multi-access (broadcast) networks.

Use the following CLI syntax to configure the designated router:

CLI Syntax:

```
ospf
area area-id
     interface ip-int-name
        priority number
```

The following displays a priority designation example:

```
A:ALU-49>config>router>ospf# info
----------------------------------------------
...

            area 0.0.0.25
                nssa
                exit
                interface "if2"
                    priority 100
                exit
            exit
----------------------------------------------
A:ALU-49>config>router>ospf#
```

Configuring Route Summaries

ABRs send summary advertisements (type 3 LSAs) into a stub area or NSSA to describe the routes to other areas. This command is particularly useful in order to reduce the size of the link-state database within the stub or NSSA.

By default, summary route advertisements are sent into the stub area or NSSA. The `no` form of the `summaries` command disables sending summary route advertisements and, in stub areas, the default route is advertised by the area border router.

Use the following CLI syntax to configure a route summary:

CLI Syntax:

```
ospf
    area area-id
       stub
          summaries
       nssa
          summaries
```

The following displays a stub route summary configuration example:

```
A:ALU-49>config>router>ospf# info
----------------------------------------------
...
            area 0.0.0.20
                stub
                    summaries
                exit
                interface "to-103"
                exit
            exit
----------------------------------------------
A:ALU-49>config>router>ospf#
```

Configuring Route Preferences

A router can learn routes from different protocols and distribute them into OSPF, in which case, the costs are not comparable. When this occurs, the preference value is used to decide which route is installed in the forwarding table and used as the path to the destination. The route with the lowest preference value is selected.

The 7705 SAR supports the redistribution of static routes and routes from directly attached and aggregated networks into OSPF.

Different protocols should not be configured with the same preference. If this occurs, the tiebreaker is based on the default preferences as defined in Table 4.

If multiple routes are learned with an identical preference using the same protocol, the lowest-cost route is used. If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, the decision of what route to use is determined by the configuration of ECMP in the `config>router` context. Refer to the 7705 SAR OS Router Configuration Guide for information on ECMP.

Table 4: Route Preference Defaults by Route Type

Route Type	Preference	Configurable
Direct attached	0	No
Static routes	5	Yes
OSPF internal	10	Yes
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes

Note: To configure a preference for static routes, use the config>router>static-route command. Refer to the 7705 SAR OS Router Configuration Guide, “IP Router Command Reference”, for information.

Use the following CLI syntax to configure a route preference for OSPF internal and external routes:

CLI Syntax:
```
ospf
    preference preference
    external-preference preference
```

The following displays a route preference configuration example:

```
A:ALU-49>config>router>ospf# info
----------------------------------------------
            asbr
            overload
            overload-on-boot timeout 60
            traffic-engineering
            preference 9
            external-preference 140
            exit
----------------------------------------------
A:ALU-49>config>router>ospf#
```

OSPF Configuration Management Tasks

This section discusses the following OSPF configuration management tasks:

- Modifying a Router ID
- Deleting a Router ID
- Modifying OSPF Parameters

Modifying a Router ID

Because the router ID is defined in the `config>router` context, not in the OSPF configuration context, the protocol instance is not aware of changes to the ID value. Changing the router ID on a device could cause configuration inconsistencies if associated values are not also modified.

After you have changed a router ID, manually shut down and restart the protocol using the `shutdown` and `no shutdown` commands in order for the changes to be incorporated.

Use the following CLI syntax to change a router ID number:

CLI Syntax: `config>router# router-id` *`router-id`*

The following displays an NSSA router ID modification example:

```
A:ALU-49>config>router# info
------------------------------------------
IP Configuration
------------------------------------------
        interface "system"
            address 10.10.10.104/32
        exit
        interface "to-103"
            address 10.0.0.103/24
            port 1/1/1
        exit
        router-id 10.10.10.104
------------------------------------------
A:ALU-49>config>router#
```

Deleting a Router ID

You can modify a router ID, but you cannot delete the parameter. If the `no router` *`router-id`* command is issued, the router ID reverts to the default value, the system interface address (which is also the loopback address). If a system interface address is not configured, the last 4 bytes of the chassis MAC address are used as the router ID.

Modifying OSPF Parameters

You can change or remove existing OSPF parameters in the CLI. The changes are applied immediately.

The following example displays an OSPF modification in which an interface is removed and another interface added.

Example:

```
config>router# ospf
config>router>ospf# area 0.0.0.20
config>router>ospf>area# no interface "to-103"
config>router>ospf>area# interface "to-HQ"
config>router>ospf>area>if$ priority 50
config>router>ospf>area>if# exit
config>router>ospf>area# exit
```

The following example displays the OSPF configuration with the modifications entered in the previous example:

```
A:ALU-49>config>router>ospf# info
----------------------------------------------
            asbr
            overload
            overload-on-boot timeout 60
            traffic-engineering
            preference 9
            external-preference 140
            export "OSPF-Export"
            exit
            area 0.0.0.0
                virtual-link 10.0.0.1 transit-area 0.0.0.1
                    authentication-type message-digest
                    message-digest-key 2 md5 "Mi6BQAFi3MI" hash
                exit
                virtual-link 1.2.3.4 transit-area 1.2.3.4
                    hello-interval 9
                    dead-interval 40
                exit
                interface "system"
                exit
            exit
            area 0.0.0.1
            exit
            area 0.0.0.20
                stub
                exit
                interface "to-HQ"
                    priority 50
                exit
            exit
----------------------------------------------
A:ALU-49>config>router>ospf#
```

OSPF Command Reference

Command Hierarchies

- Configuration Commands
- Show Commands
- Clear Commands
- Debug Commands
- Tools Commands (refer to the Tools chapter in the 7705 SAR OS Services Guide)

Configuration Commands

```
config
    — router
        — [no] ospf
            — [no] area area-id
                — area-range ip-prefix/mask [advertise | not-advertise]
                — no area-range ip-prefix/mask
                — [no] blackhole-aggregate
                — [no] interface ip-int-name
                    — [no] advertise-subnet
                    — authentication-key [authentication-key | hash-key] [hash |
                      hash2]
                    — no authentication-key
                    — authentication-type {password | message-digest}
                    — no authentication-type
                    — [no] bfd-enable [remain-down-on-failure]
                    — dead-interval seconds
                    — no dead-interval
                    — hello-interval seconds
                    — no hello-interval
                    — interface-type {broadcast | point-to-point}
                    — no interface-type
                    — message-digest-key key-id md5 {key | hash-key} [hash |
                      hash2]
                    — no message-digest-key key-id
                    — metric metric
                    — no metric
                    — mtu bytes
                    — no mtu
                    — [no] passive
                    — priority number
                    — no priority
                    — retransmit-interval seconds
                    — no retransmit-interval
                    — [no] shutdown
                    — transit-delay seconds
                    — no transit-delay
                — [no] nssa
                    — area-range ip-prefix/mask [advertise | not-advertise]
                    — no area-range ip-prefix/mask
                    — originate-default-route [type-7]
                    — no originate-default-route
                    — [no] redistribute-external
                    — [no] summaries
                — [no] stub
                    — default-metric metric
                    — no default-metric
                    — [no] summaries
```

```
          — [no] virtual-link router-id  transit-area area-id
                    — authentication-key [authentication-key | hash-key] [hash |
                      hash2]
                    — no authentication-key
                    — authentication-type {password | message-digest}
                    — no authentication-type
                    — dead-interval seconds
                    — no dead-interval
                    — hello-interval seconds
                    — no hello-interval
                    — message-digest-key key-id md5 {key | hash-key} [hash |
                      hash2]
                    — no message-digest-key key-id
                    — retransmit-interval seconds
                    — no retransmit-interval
                    — [no] shutdown
                    — transit-delay seconds
                    — no transit-delay
— [no] asbr [trace-path domain-id]
— [no] disable-ldp-sync
— export policy-name [ policy-name...(up to 5 max)]
— no export
— external-db-overflow limit seconds
— no external-db-overflow
— external-preference preference
— no external-preference
— overload [timeout seconds]
— no overload
— [no] overload-include-stub
— overload-on-boot [timeout seconds]
— no overload-on-boot
— preference preference
— no preference
— reference-bandwidth bandwidth-in-kbps
— no reference-bandwidth
— router-id ip-address
— no router-id
— [no] shutdown
— timers
          — lsa-arrival lsa-arrival-time
          — no lsa-arrival
          — lsa-generate max-lsa-wait [lsa-initial-wait [lsa-second-wait]]
          — no lsa-generate
          — spf-wait max-spf-wait [spf-initial-wait [spf-second-wait]]
          — no spf-wait
— [no] traffic-engineering
```

Show Commands

```
show
     — router
              — ospf
                       — area [area-id] [detail]
                       — database [type {router | network | summary | asbr-summary | external | nssa |
                         all}] [area area-id] [adv-router router-id] [link-state-id] [detail]
                       — interface [area area-id] [detail]
                       — interface [ip-int-name | ip-address] [detail]
                       — neighbor [ip-int-name] [router-id]  [detail]
                       — neighbor [remote ip-address] [detail]
                       — range [area-id]
                       — spf
                       — statistics
                       — status
                       — virtual-link [detail]
                       — virtual-neighbor [remote ip-address] [detail]
```

Clear Commands

```
clear
     — router
              — ospf
                       — database [purge]
                       — export
                       — neighbor [ip-int-name | ip-address]
                       — statistics
```

Debug Commands

```
debug
     — router
              — ospf
                       — area [area-id]
                       — no area
                       — area-range [ip-address]
                       — no area-range
                       — cspf [ip-addr]
                       — no cspf
                       — interface [ip-int-name | ip-address]
                       — no interface
                       — leak [ip-address]
                       — no leak
                       — lsdb [type] [ls-id] [adv-rtr-id] [area area-id]
                       — no lsdb
                       — [no] misc
                       — neighbor [ip-int-name | router-id]
```

— **no neighbor**
— **nssa-range** [*ip-address*]
— **no nssa-range**
— **packet** [*packet-type*] [*ip-address*]
— **no packet**
— **rtm** [*ip-address*]
— **no rtm**
— **spf** [*type*] [*dest-addr*]
— **no spf**
— **virtual-neighbor** [*ip-address*]
— **no virtual-neighbor**

Command Descriptions

Configuration Commands

- Generic Commands on page 72
- Global Commands on page 73
- Area Commands on page 84
- Interface/Virtual Link Commands on page 89

Generic Commands

shutdown

Syntax [**no**] **shutdown**

Context config>router>ospf
config>router>ospf>area>interface
config>router>ospf>area>virtual-link

Description This command administratively disables the entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many entities must be explicitly enabled using the **no shutdown** command.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they can be deleted.

Unlike other commands and parameters where the default state is not indicated in the configuration file, **shutdown** and **no shutdown** are always indicated in system-generated configuration files.

The **no** form of the command puts an entity into the administratively enabled state.

Default **OSPF Protocol —** The Open Shortest Path First (OSPF) protocol is created in the **no shutdown** state.

OSPF Interface — When an IP interface is configured as an OSPF interface, OSPF on the interface is in the **no shutdown** state by default.

Global Commands

ospf

Syntax [**no**] **ospf**

Context config>router

Description This command activates OSPF on the router and enables access to the context to define OSPF parameters.

Before OSPF can be activated on the router, the router ID must be configured.

The router ID uniquely identifies the router within an AS. In OSPF, routing information is exchanged between autonomous systems, which are groups of networks that share routing information. The router ID can be set to be the same as the system interface address (loopback address).

The router ID is derived by one of the following methods:

- defining the value using the `config>router` *`router-id`* command
- defining the system interface using the `config>router>interface` *`ip-int-name`* command (used if the router ID is not specified with the `config>router router-id` command)
- inheriting the last 4 bytes of the MAC address

When configuring a new router ID, protocols are not automatically restarted with the new router ID. The next time a protocol is initialized, the new router ID is used. To force the new router ID, issue the `shutdown` and `no shutdown` commands for OSPF or restart the entire router.

The **no** form of the command reverts to the default value.

Default **no ospf**

asbr

Syntax [**no**] **asbr** [**trace-path** *domain-id*]

Context config>router>ospf

Description This command configures the router as an Autonomous System Boundary Router (ASBR) if the router is to be used to distribute external routes into the OSPF domain. When a router is configured as an ASBR, the export policies into the OSPF domain take effect. If no policies are configured, no external routes are redistributed into the OSPF domain.

The **no** form of the command removes the ASBR status and withdraws the routes redistributed from the routing table into OSPF from the link-state database.

In Release 3.0, only the base OSPF instance is supported; therefore, the domain ID may not need to be configured. However, in order to prevent routing loops (where routes learned from one domain are redistributed back into the domain), the domain ID can be used to tag external LSAs – indicating which domain or network they have learned the route from.

Default **no asbr —** the router is not an ASBR

Parameters *domain-id* — specifies the domain ID

Values 1 to 31

Default 0x0

disable-ldp-sync

Syntax [**no**] **disable-ldp-sync**

Context config>router>ospf

Description This command disables the IGP-LDP synchronization feature on all interfaces participating in the OSPF or IS-IS routing protocol. When this command is executed, IGP immediately advertises the actual value of the link cost for all interfaces that have the IGP-LDP synchronization enabled if the currently advertised cost is different. IGP-LDP synchronization will then be disabled for all interfaces. This command does not delete the interface configuration.

The **no** form of this command restores the default settings and re-enables IGP-LDP synchronization on all interfaces participating in the OSPF or IS-IS routing protocol and for which the **ldp-sync-timer** is configured (refer to the 7705 SAR OS Router Configuration Guide for information on configuring the **ldp-sync-timer**).

Default **no disable-ldp-sync**

export

Syntax **export** *policy-name* [*policy-name…(*up to 5 max*)*]
no export

Context config>router>ospf

Description This command associates export route policies to determine which routes are exported from the route table to OSPF. Export polices are only in effect if OSPF is configured as an ASBR.

If no export policy is specified, non-OSPF routes are not exported from the routing table manager to OSPF.

If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered will override the previous command. A maximum of five policy names can be specified.

The **no** form of the command removes all policies from the configuration.

Refer to the 7705 SAR OS Router Configuration Guide for information on defining route policies.

Default **no export —** no export route policies specified

Parameters *policy-name* — the export route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

The specified name(s) must already be defined.

external-db-overflow

Syntax **external-db-overflow** *limit seconds*
no external-db-overflow

Context config>router>ospf

Description This command enables limits on the number of non-default, AS-external-LSA entries that can be stored in the link-state database (LSDB) and specifies a wait timer before processing these after the limit is exceeded.

The *limit* value specifies the maximum number of entries that can be stored in the LSDB. Placing a limit on these LSAs in the LSDB protects the router from receiving an excessive number of external routes that consume excessive memory or CPU resources. If the number of routes reaches or exceeds the *limit*, the table is in an overflow state. When in an overflow state, the router will not originate any new AS-external-LSAs and will withdraw all the self-originated non-default external LSAs.

The *seconds* value specifies the amount of time to wait after an overflow state before regenerating and processing non-default, AS-external-LSAs. The waiting period acts like a dampening period, preventing the router from continuously running shortest path first (SPF) calculations caused by the excessive number of non-default, AS-external LSAs.

The **external-db-overflow** must be set identically on all routers attached to any regular OSPF area. OSPF stub areas and not-so-stubby areas (NSSAs) are excluded.

The **no** form of the command disables limiting the number of non-default, AS-external-LSA entries.

Default **no external-db-overflow**

Parameters *limit* — the maximum number of non-default, AS-external-LSA entries that can be stored in the LSDB before going into an overflow state, expressed as a decimal integer

Values 0 to 2147483674

seconds — the number of seconds after entering an overflow state before attempting to process non-default, AS-external-LSAs, expressed as a decimal integer

Values 0 to 2147483674

external-preference

Syntax **external-preference** *preference*
no external-preference

Context config>router>ospf

Description This command configures the preference for OSPF external routes. The preference for internal routes is set with the preference command.

A route can be learned by the router from different protocols, in which case, the costs are not comparable. When this occurs, the preference is used to decide which route will be used.

Different protocols should not be configured with the same preference. If this occurs, the tiebreaker is based on the default preferences as defined in Table 5.

Table 5: Route Preference Defaults by Route Type

Route Type	Preference	Configurable
Direct attached	0	No
Static routes	5	Yes
OSPF internal	10	Yes
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes

If multiple routes are learned with an identical preference using the same protocol, the lowest-cost route is used. If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, the decision of what route to use is determined by the configuration of ECMP in the **config>router** context. Refer to the 7705 SAR OS Router Configuration Guide for information on ECMP.

Note: To configure a preference for static routes, use the config>router>static-route command. Refer to the 7705 SAR OS Router Configuration Guide for information.

The **no** form of the command reverts to the default value.

Default **external-preference 150 —** OSPF external routes have a default preference of 150

Parameters *preference* — the preference for external routes expressed as a decimal integer

Values 1 to 255

overload

Syntax **overload** [**timeout** *seconds*]
no overload

Context config>router>ospf

Description This command changes the overload state of the local router so that it appears to be overloaded. When overload is enabled, the router can participate in OSPF routing, but is not used for transit traffic. Traffic destined for directly attached interfaces continues to reach the router.

To put the IGP in an overload state, enter a timeout value. The IGP will enter the overload state until the timeout timer expires or a **no overload** command is executed.

If no timeout is specified, the overload state is maintained indefinitely.

If the **overload** command is encountered during the execution of an overload-on-boot command, the **overload** command takes precedence. This could occur as a result of a saved configuration file where both parameters are saved. When the file is saved by the system, the **overload-on-boot** command is saved after the **overload** command.

Use the **no** form of this command to return to the default. When the **no overload** command is executed, the overload state is terminated regardless of the reason the protocol entered the overload state.

Default **no overload**

Parameters *seconds* — the number of seconds to reset overloading

Values 1 to 1800

overload-include-stub

Syntax [**no**] **overload-include-stub**

Context config>router>ospf

Description This command is used to determine if the OSPF stub networks should be advertised with a maximum metric value when the system goes into an overload state for any reason. When enabled, the system uses the maximum metric value. When this command is enabled and the router is in overload, all stub interfaces, including loopback and system interfaces, will be advertised at the maximum metric.

Default **no overload-include-stub**

overload-on-boot

Syntax **overload-on-boot** [**timeout** *seconds*]
no overload-on-boot

Context config>router>ospf

Description When the router is in an overload state, the router is used only if there is no other router to reach the destination. This command configures OSPF upon bootup in the overload state until one of the following events occurs:

- the timeout timer expires (if a timeout has been specified)
- a manual override of the current overload state is entered with the **no overload** command

If no timeout is specified, the overload state is maintained indefinitely.

The **no overload** command does not affect the **overload-on-boot** function.

The **no** form of the command removes the overload-on-boot functionality from the configuration.

Default **no overload-on-boot**

Parameters *seconds* — the number of seconds to reset overloading

Values 1 to 1800

preference

Syntax **preference** *preference*
no preference

Context config>router>ospf

This command configures the preference for OSPF internal routes.

A route can be learned by the router from different protocols, in which case, the costs are not comparable. When this occurs, the preference is used to decide which route will be used.

Different protocols should not be configured with the same preference. If this occurs, the tiebreaker is based on the default preferences as defined in Table 5. If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, the decision of what route to use is determined by the configuration of ECMP in the **config>router** context. Refer to the 7705 SAR OS Router Configuration Guide for information on ECMP.

The **no** form of the command reverts to the default value.

Default **preference 10 —** OSPF internal routes have a preference of 10

Parameters *preference* — the preference for internal routes expressed as a decimal integer

Values 1 to 255

reference-bandwidth

Syntax **reference-bandwidth** *bandwidth-in-kbps*
no reference-bandwidth

Context config>router>ospf

Description This command configures the reference bandwidth that provides the reference for the default costing of interfaces based on their underlying link speed.

The default interface cost is calculated as follows:

cost = reference bandwidth/bandwidth

The default reference bandwidth is 100 000 000 kb/s or 100 Gb/s; therefore, the default auto-cost metrics for various link speeds are as follows:

- 10 Mb/s link default cost of 10000
- 100 Mb/s link default cost of 1000
- 1 Gb/s link default cost of 100

The **reference-bandwidth** command assigns a default cost to the interface based on the interface speed. To override this default cost on a particular interface, use the **metric** *metric* command in the **config>router>ospf>area>interface** *ip-int-name* context.

The **no** form of the command reverts the reference bandwidth to the default value.

Default **reference-bandwidth 100000000**

Parameters *bandwidth-in-kbps* — the reference bandwidth in kilobits per second expressed as a decimal integer

Values 1 to 1000000000

router-id

Syntax **router-id** *ip-address*
no router-id

Context config>router>ospf

Description This command configures the router ID for the OSPF instance. This is a required command when configuring multiple instances and the instance being configured is not the base instance. The 7705 SAR does not support multiple instances of OSPF; therefore, it is recommended that this command not be used.

When configuring the router ID in the base instance of OSPF, the value overrides the router ID configured in the **config>router** context.

The default value for the base instance is inherited from the configuration in the **config>router** context. If the router ID in the **config>router** context is not configured, the following applies:

- the system uses the system interface address (which is also the loopback address)
- if a system interface address is not configured, the last 4 bytes of the chassis MAC address are used

When configuring a new router ID, the instance is not automatically restarted with the new router ID. The next time the instance is initialized, the new router ID is used.

To force the new router ID to be used, issue the **shutdown** and **no shutdown** commands for the instance, or reboot the entire router.

The **no** form of the command to reverts to the default value.

Default **0.0.0.0 (base OSPF)**

Parameters *ip-address* — a 32-bit, unsigned integer uniquely identifying the router in the Autonomous System

timers

Syntax **timers**

Context config>router>ospf

Description This command enables the context that allows for the configuration of OSPF timers. Timers control the delay between receipt of a link-state advertisement (LSA) requiring an SPF calculation and the minimum time between successive SPF calculations.

Changing the timers affects CPU utilization and network reconvergence times. Lower values reduce reconvergence time but increase CPU utilization. Higher values reduce CPU utilization but increase reconvergence time.

Default **none**

lsa-arrival

Syntax **lsa-arrival** *lsa-arrival-time*
no lsa-arrival

Context config>router>ospf>timers

Description This command defines the minimum delay that must pass between receipt of the same link-state advertisements (LSAs) arriving from neighbors.

It is recommended that the neighbors' configured **lsa-generate** *lsa-second-wait* interval be equal to or greater than the *lsa-arrival-time*.

Use the **no** form of this command to return to the default.

Default **no lsa-arrival**

Parameters *lsa-arrival-time* — the timer in milliseconds. Values entered that do not match this requirement will be rejected.

Values 0 to 600000

lsa-generate

Syntax **lsa-generate** *max-lsa-wait* [*lsa-initial-wait* [*lsa-second-wait*]]
no lsa-generate

Context config>router>ospf>timers

Description This command customizes the throttling of OSPF LSA generation. Timers that determine when to generate the first, second, and subsequent LSAs can be controlled with this command. Subsequent LSAs are generated at increasing intervals of the *lsa-second-wait* timer until a maximum value is reached.

It is recommended that the *lsa-arrival-time* be equal to or less than the *lsa-second-wait* interval.

Use the **no** form of this command to return to the default.

Default **no lsa-generate**

Parameters *max-lsa-wait* — the maximum interval, in milliseconds, between two consecutive ocurrences of an LSA being generated

Values	10 to 600000
Default	5000

lsa-initial-wait — the first waiting period between LSAs generated, in milliseconds. When the LSA exceeds the *lsa-initial-wait* timer value and the topology changes, there is no wait period and the LSA is immediately generated.

When an LSA is generated, the initial wait period commences. If, within the specified *lsa-initial-wait* period, another topology change occurs, the *lsa-initial-wait* timer applies.

Values	10 to 600000
Default	5000

lsa-second-wait — the hold time in milliseconds between the first and second LSA generation. The next topology change is subject to this second wait period. With each subsequent topology change, the wait time doubles (this is 2x the previous wait time.). This assumes that each failure occurs within the relevant wait period.

Values	10 to 600000
Default	5000

spf-wait

Syntax **spf-wait** *max-spf-wait* [*spf-initial-wait* [*spf-second-wait*]]
no spf-wait

Context config>router>ospf>timers

Description This command defines the maximum interval between two consecutive SPF calculations in milliseconds. Timers that determine when to initiate the first, second, and subsequent SPF calculations after a topology change occurs can be controlled with this command. Subsequent SPF runs (if required) will occur at exponentially increasing intervals of the *spf-second-wait* interval. For example, if the *spf-second-wait* interval is 1000, the next SPF will run after 2000 milliseconds, and the next SPF will run after 4000 milliseconds, etc., until it reaches the **spf-wait** value. The SPF interval will stay at the **spf-wait** value until there are no more SPF runs scheduled in that interval. After a full interval without any SPF runs, the SPF interval will drop back to *spf-initial-wait*.

The timer must be entered in increments of 100 milliseconds. Values entered that do not match this requirement will be rejected.

Use the **no** form of this command to return to the default.

Default **no spf-wait**

Parameters *max-spf-wait* — the maximum interval in milliseconds between two consecutive SPF calculations

Values 10 to 120000

Default 1000

spf-initial-wait — the initial SPF calculation delay in milliseconds after a topology change

Values 10 to 100000

Default 1000

spf-second-wait — the hold time in milliseconds between the first and second SPF calculation

Values 10 to 100000

Default 1000

traffic-engineering

Syntax [**no**] **traffic-engineering**

Context config>router>ospf

Description This command enables traffic engineering route calculations constrained by nodes or links.

Traffic engineering enables the router to perform route calculations constrained by nodes or links. The traffic engineering capabilities of this router are limited to calculations based on link and nodal constraints.

The **no** form of the command disables traffic engineered route calculations.

Default **no traffic-engineering**

Area Commands

area

Syntax [**no**] **area** *area-id*

Context config>router>ospf

Description This command creates the context to configure an OSPF area. An area is a collection of network segments within an AS that have been administratively grouped together. The area ID can be specified in dotted-decimal notation or as a 32-bit decimal integer.

The **no** form of the command deletes the specified area from the configuration. Deleting the area also removes the OSPF configuration of all the interfaces, virtual-links, address ranges, and so on, that are currently assigned to this area.

In Release 3.0, the 7705 SAR supports a maximum of four areas.

Default **no area —** no OSPF areas are defined

Parameters *area-id* — the OSPF area ID expressed in dotted-decimal notation or as a 32-bit decimal integer

Values 0.0.0.0 to 255.255.255.255 (dotted-decimal), 0 to 4294967295 (decimal integer)

area-range

Syntax **area-range** *ip-prefix*/*mask* [**advertise** | **not-advertise**]
no area-range *ip-prefix*/*mask*

Context config>router>ospf>area
config>router>ospf>area>nssa

Description This command creates ranges of addresses on an Area Border Router (ABR) for the purpose of route summarization or suppression. When a range is created, the range is configured to be advertised or not advertised into other areas. Multiple range commands can be used to summarize or hide different ranges. In the case of overlapping ranges, the most specific range command applies.

ABRs send summary link advertisements to describe routes to other areas. To minimize the number of advertisements that are flooded, you can summarize a range of IP addresses and send reachability information about these addresses in an LSA.

The **no** form of the command deletes the range advertisement or non-advertisement.

Default **no area-range —** no range of addresses is defined

Special Cases

NSSA Context — In the NSSA context, the option specifies that the range applies to external routes (via type 7 LSAs) learned within the NSSA when the routes are advertised to other areas as type 5 LSAs.

Area Context — If this command is not entered under the NSSA context, the range applies to summary LSAs even if the area is an NSSA.

Parameters *ip-prefix* — the IP prefix in dotted-decimal notation for the range

Values a.b.c.d (host bits must be 0)

mask — the subnet mask for the range expressed as a decimal integer

Values 0 to 32

advertise | **not-advertise —** specifies whether or not to advertise the summarized range of addresses into other areas

Default advertise

blackhole-aggregate

Syntax [**no**] **blackhole-aggregate**

Context config>router>ospf>area

Description This command installs a low-priority blackhole route for the entire aggregate. Existing routes that make up the aggregate will have a higher priority and only the components of the range for which no route exists are blackholed.

It is possible that when performing area aggregation, addresses may be included in the range for which no actual route exists. This can cause routing loops. To avoid this problem, configure the blackhole aggregate option.

The **no** form of this command removes this option.

Default **blackhole-aggregate**

nssa

Syntax [**no**] **nssa**

Context config>router>ospf>area

Description This command creates the context to configure an OSPF Not So Stubby Area (NSSA) and adds or removes the NSSA designation from the area.

NSSAs are similar to stub areas in that no external routes are imported into the area from other OSPF areas. The major difference between a stub area and an NSSA is that an NSSA has the capability to flood external routes that it learns throughout its area and via an ABR to the entire OSPF domain.

Existing virtual links of a non-stub area or NSSA are removed when the designation is changed to NSSA or stub.

An area can be designated as stub or NSSA but never both at the same time.

By default, an area is not configured as an NSSA area.

The **no** form of the command removes the NSSA designation and configuration context from the area.

Default **no nssa**

originate-default-route

Syntax **originate-default-route** [**type-7**]
no originate-default-route

Context config>router>ospf>area>nssa

Description This command enables the generation of a default route and its LSA type (3 or 7) into a Not So Stubby Area (NSSA) by an NSSA Area Border Router (ABR) or Autonomous System Border Router (ASBR).

Include the **type-7** parameter to inject a type 7 LSA default route instead the type 3 LSA into the NSSA configured with no summaries.

To revert to a type 3 LSA, enter **originate-default-route** without the **type-7** parameter.

When configuring an NSSA with no summaries, the ABR will inject a type 3 LSA default route into the NSSA area. Some older implementations expect a type 7 LSA default route.

The **no** form of the command disables origination of a default route.

Default **no originate-default-route**

Parameters **type-7** — specifies that a type 7 LSA should be used for the default route

Default type 3 LSA for the default route

redistribute-external

Syntax [no] **redistribute-external**

Context config>router>ospf>area>nssa

Description This command enables the redistribution of external routes into the Not So Stubby Area (NSSA) on an NSSA area border router (ABR) that is exporting the routes into non-NSSA areas.

NSSAs are similar to stub areas in that no external routes are imported into the area from other OSPF areas. The major difference between a stub area and an NSSA is that the NSSA has the capability to flood external routes that it learns (providing it is an ASBR) throughout its area and via an ABR to the entire OSPF domain.

The **no** form of the command disables the default behavior to automatically redistribute external routes into the NSSA area from the NSSA ABR.

Default **redistribute-external**

summaries

Syntax [no] **summaries**

Context config>router>ospf>area>nssa
config>router>ospf>area>stub

Description This command enables sending summary (type 3) advertisements into a stub area or NSSA on an ABR.

This parameter is particularly useful to reduce the size of the routing and link-state database (LSDB) tables within the stub or NSSA area.

By default, summary route advertisements are sent into the stub area or NSSA.

The **no** form of the command disables sending summary route advertisements and, for stub areas, only the default route is advertised by the ABR.

Default **summaries**

stub

Syntax **[no] stub**

Context config>router>ospf>area

Description This command enables access to the context to configure an OSPF stub area and adds or removes the stub designation from the area.

External routing information is not flooded into stub areas. All routers in the stub area must be configured with the **stub** command.

Existing virtual links of a non-stub area or NSSA are removed when its designation is changed to NSSA or stub.

An OSPF area cannot be both an NSSA and a stub area at the same time.

By default, an area is not a stub area.

The **no** form of the command removes the stub designation and configuration context from the area.

Default **no stub**

default-metric

Syntax **default-metric** *metric*
no default-metric

Context config>router>ospf>area>stub

Description This command configures the metric used by the ABR for the default route into a stub area.

The default metric should only be configured on an ABR of a stub area.

An ABR generates a default route if the area is a **stub** area.

The **no** form of the command reverts to the default value.

Default **default-metric 1**

Parameters *metric* — the metric expressed as a decimal integer for the default route cost to be advertised into the stub area

Values 1 to 65535

Interface/Virtual Link Commands

interface

Syntax [**no**] **interface** *ip-int-name*

Context config>router>ospf>area

Description This command creates a context to configure an OSPF interface.

By default, interfaces are not activated in any interior gateway protocol, such as OSPF, unless explicitly configured.

The **no** form of the command deletes the OSPF interface configuration for this interface. The **shutdown** command in the **config>router>ospf>interface** context can be used to disable an interface without removing the configuration for the interface.

Default **no interface**

Parameters *ip-int-name* — the IP interface name. Interface names must be unique within the group of defined IP interfaces for the **config router interface** command. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

If the IP interface name does not exist or does not have an IP address configured, an error message will be returned.

If the IP interface exists in a different area, it will be moved to this area.

advertise-subnet

Syntax [**no**] **advertise-subnet**

Context config>router>ospf>area>interface

Description This command enables advertising point-to-point interfaces as subnet routes (network number and mask). When disabled, point-to-point interfaces are advertised as host routes.

The **no** form of the command disables advertising point-to-point interfaces as subnet routes, meaning they are advertised as host routes.

Default **advertise-subnet**

authentication-key

Syntax **authentication-key** [*authentication-key* **|** *hash-key*] [**hash | hash2**]
no authentication-key

Context config>router>ospf>area>interface
config>router>ospf>area>virtual-link

Description This command configures the password used by the OSPF interface or virtual link to send and receive OSPF protocol packets on the interface when simple password authentication is configured.

All neighboring routers must use the same type of authentication and password for proper protocol communication. If the **authentication-type** is configured as password, this key must be configured.

By default, no authentication key is configured.

The **no** form of the command removes the authentication key.

Default **no authentication-key**

Parameters *authentication-key* — the authentication key can be any combination of ASCII characters up to 8 characters in length (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

hash-key — the hash key can be any combination of ASCII characters up to 22 characters in length (encrypted) or 121 characters in length (if the **hash2** parameter is used). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash — specifies that the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2 — specifies that the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted hash form is assumed.

authentication-type

Syntax **authentication-type** {**password | message-digest**}
no authentication-type

Context config>router>ospf>area>interface
config>router>ospf>area>virtual-link

Description This command enables authentication and specifies the type of authentication to be used on the OSPF interface.

Both simple **password** and **message-digest** authentication are supported.

By default, authentication is not enabled on an interface.

The **no** form of the command disables authentication on the interface.

Default **no authentication-type**

Parameters **password —** enables simple password (plain text) authentication. If authentication is enabled and no authentication type is specified in the command, simple **password** authentication is enabled.

message-digest — enables message digest MD5 authentication in accordance with RFC 1321. If this option is configured, at least one message-digest-key must be configured.

bfd-enable

Syntax [**no**] **bfd-enable** [**remain-down-on-failure**]

Context config>router>ospf>area>interface

Description This command enables the use of bidirectional forwarding (BFD) to control the state of the associated OSPF interface. By enabling BFD on a given OSPF interface, the state of the interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for BFD are set via the BFD command under the IP interface.

The optional **remain-down-on-failure** parameter can be specified on OSPF interfaces that are enabled for BFD to keep OSPF from reaching the Full state if the BFD session to that neighbor cannot be established. This option is disabled by default and should be used only if there is a chance that unicast packets might be discarded while multicast packets are forwarded.

The **no** form of this command removes BFD from the associated OSPF adjacency.

Default **no bfd-enable**

dead-interval

Syntax **dead-interval** *seconds*
no dead-interval

Context config>router>ospf>area>interface
config>router>ospf>area>virtual-link

Description This command configures the time, in seconds, that OSPF waits before declaring a neighbor router down. If no Hello packets are received from a neighbor for the duration of the dead interval, the router is assumed to be down. The minimum interval must be two times the hello interval.

The **no** form of the command reverts to the default value.

Default **40**

Special Cases

OSPF Interface — If the **dead-interval** configured applies to an interface, all nodes on the subnet must have the same dead interval.

Virtual Link — If the **dead-interval** configured applies to a virtual link, the interval on both termination points of the virtual link must have the same dead interval.

Parameters *seconds* — the dead interval expressed in seconds

Values 1 to 65535

hello-interval

Syntax **hello-interval** *seconds*
no hello-interval

Context config>router>ospf>area>interface
config>router>ospf>area>virtual-link

Description This command configures the interval between OSPF hellos issued on the interface or virtual link.

The hello interval, in combination with the dead interval, is used to establish and maintain the adjacency. Use this parameter to edit the frequency that Hello packets are sent.

Reducing the interval, in combination with an appropriate reduction in the associated **dead-interval,** allows for faster detection of link and/or router failures but results in higher processing costs.

The **no** form of this command reverts to the default value.

Default **10**

Special Cases

OSPF Interface — If the **hello-interval** configured applies to an interface, all nodes on the subnet must have the same hello interval.

Virtual Link — If the **hello-interval** configured applies to a virtual link, the interval on both termination points of the virtual link must have the same hello interval.

Parameters *seconds* — the hello interval in seconds expressed as a decimal integer

Values 1 to 65535

interface-type

Syntax **interface-type** {**broadcast** | **point-to-point**}
no interface-type

Context config>router>ospf>area>interface

Description This command configures the interface type to be either broadcast or point-to-point.

Use this command to set the interface type of an Ethernet link to point-to-point to avoid having to carry the broadcast adjacency maintenance overhead of the Ethernet link, provided the link is used as a point-to-point link.

If the interface type is not known at the time the interface is added to OSPF, and the IP interface is subsequently bound (or moved) to a different interface type, this command must be entered manually.

The **no** form of the command reverts to the default value.

Default **broadcast** – if the physical interface is Ethernet or unknown

point-to-point – if the physical interface is T1, E1, or SONET/SDH

Special Cases

Virtual Link — A virtual link is always regarded as a point-to-point interface and is not configurable.

Parameters **broadcast —** configures the interface to maintain this link as a broadcast network. To significantly improve adjacency forming and network convergence, a network should be configured as point-to-point if only two routers are connected, even if the network is a broadcast media such as Ethernet.

point-to-point — configures the interface to maintain this link as a point-to-point link

message-digest-key

Syntax **message-digest-key** *keyid* **md5** {*key* | *hash-key*} [**hash** | **hash2**]
no message-digest-key *keyid*

Context config>router>ospf>area>interface
config>router>ospf>area>virtual-link

Description This command configures a message digest key when MD5 authentication is enabled on the interface. Multiple message digest keys can be configured.

The **no** form of the command removes the message digest key identified by the *key-id*.

Default **no message-digest-key**

Parameters *keyid* — the *keyid* is expressed as a decimal integer

Values 1 to 255

key — the MD5 key, any alphanumeric string up to 16 characters in length

hash-key — the MD5 hash key, any combination of ASCII characters up to 33 characters in length (encrypted) or 132 characters in length (if the **hash2** parameter is used). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash — specifies that the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2 — specifies that the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted hash form is assumed.

metric

Syntax **metric** *metric*
no metric

Context config>router>ospf>area>interface

Description This command configures an explicit route cost metric for the OSPF interface that overrides the metrics calculated based on the speed of the underlying link.

The **no** form of the command deletes the manually configured interface metric, so the interface uses the computed metric based on the **reference-bandwidth** command setting and the speed of the underlying link.

Default **no metric**

Parameters *metric* — the metric to be applied to the interface expressed as a decimal integer

Values 1 to 65535

mtu

Syntax **mtu** *bytes*
no mtu

Context config>router>ospf>area>interface

Description This command configures the OSPF packet size used on this interface. If this parameter is not configured, OSPF derives the MTU value from the MTU configured (default or explicitly) in the following contexts:

- **config>port>ethernet**
- **config>port>tdm>t1-e1>channel-group**

If this parameter is configured, the smaller value between the value configured with this command and the MTU configured (default or explicitly) in the preceding listed contexts is used.

To determine the actual packet size, add 14 bytes for an Ethernet packet and 18 bytes for a tagged Ethernet packet to the size of the OSPF (IP) packet MTU configured with this command.

Use the **no** form of this command to revert to the default.

Default **no mtu —** uses the value derived from the MTU configured in the **config>port** context

Parameters *bytes* — the MTU to be used by OSPF for this logical interface in bytes

Values 512 to 2084 (2098 – 14) (depends on the physical media)

passive

Syntax [**no**] **passive**

Context config>router>ospf>area>interface

Description This command adds the passive property to the OSPF interface where passive interfaces are advertised as OSPF interfaces but do not run the OSPF protocol.

By default, only interface addresses that are configured for OSPF will be advertised as OSPF interfaces. The **passive** parameter allows an interface to be advertised as an OSPF interface without running the OSPF protocol.

While in passive mode, the interface will ignore ingress OSPF protocol packets and not transmit any OSPF protocol packets.

The **no** form of the command removes the passive property from the OSPF interface.

Default Service interfaces defined with the **config**>**router**>**service-prefix** command are passive. All other interfaces are not passive.

priority

Syntax **priority** *number*
no priority

Context config>router>ospf>area>interface

Description This command configures the priority of the OSPF interface that is used in an election of the designated router on the subnet.

This parameter is only used if the interface is of type broadcast. The router with the highest priority interface becomes the designated router. A router with priority 0 is not eligible to be a designated router or backup designated router.

The **no** form of the command reverts the interface priority to the default value.

Default **1**

Parameters *number* — the interface priority expressed as a decimal integer

Values 0 to 255

retransmit-interval

Syntax **retransmit-interval** *seconds*
no retransmit-interval

Context config>router>ospf>area>interface
config>router>ospf>area>virtual-link

Description This command specifies the length of time, in seconds, that OSPF will wait before retransmitting an unacknowledged link-state advertisement (LSA) to an OSPF neighbor.

The value should be longer than the expected round-trip delay between any two routers on the attached network. If the retransmit interval expires and no acknowledgement has been received, the LSA will be retransmitted.

The **no** form of this command reverts to the default interval.

Default **5**

Parameters *seconds* — the retransmit interval in seconds expressed as a decimal integer

Values 1 to 1800

transit-delay

Syntax **transit-delay** *seconds*
no transit-delay

Context config>router>ospf>area>interface
config>router>ospf>area>virtual-link

Description This command configures the estimated time, in seconds, that it takes to transmit a link-state advertisement (LSA) on the interface or virtual link.

The **no** form of this command reverts to the default delay time.

Default **1**

Parameters *seconds* — the transit delay in seconds expressed as a decimal integer

Values 1 to 1800

virtual-link

Syntax [**no**] **virtual-link** *router-id* **transit-area** *area-id*

Context config>router>ospf>area

Description This command configures a virtual link to connect ABRs to the backbone.

The backbone area (area 0.0.0.0) must be contiguous and all other areas must be connected to the backbone area. If it is not practical or possible to connect an area to the backbone, the ABRs must be connected via a virtual link. The two ABRs form a point-to-point-like adjacency across the transit area. A virtual link can only be configured while in the area 0.0.0.0 context.

The *router-id* specified in this command must be associated with the virtual neighbor. The transit area cannot be a stub area or an NSSA.

The **no** form of the command deletes the virtual link.

Default **no virtual-link**

Parameters *router-id* — the router ID of the virtual neighbor in IP address dotted-decimal notation

area-id — the area ID specified identifies the transit area that links the backbone area to the area that has no physical connection with the backbone, expressed in dotted-decimal notation or as a 32-bit decimal integer

Values 0.0.0.0 to 255.255.255.255 (dotted-decimal), 0 to 4294967295 (decimal integer)

Show Commands

ospf

Syntax **ospf**

Context show>router

Description This command enables the context to display OSPF information.

area

Syntax **area** [*area-id*] [**detail**]

Context show>router>ospf

Description This command displays configuration information about all areas or the specified area. When **detail** is specified, operational and statistical information will be displayed.

Parameters *area-id* — the OSPF area ID expressed in dotted-decimal notation or as a 32-bit decimal integer

detail — displays detailed information on the area

Output The following output is an example of OSPF area information, and Table 6 describes the fields.

Sample Output

```
A:ALU-A# show router ospf area detail
===============================================================================
OSPF Areas (Detailed)
===============================================================================
-------------------------------------------------------------------------------
Area Id: 0.0.0.0
-------------------------------------------------------------------------------
Area Id          : 0.0.0.0              Type             : Standard
Virtual Links    : 0                    Total Nbrs       : 2
Active IFs       : 3                    Total IFs        : 3
Area Bdr Rtrs    : 0                    AS Bdr Rtrs      : 0
SPF Runs         : 7                    Last SPF Run     : 10/26/2008 10:09:18
Router LSAs      : 3                    Network LSAs     : 3
Summary LSAs     : 0                    Asbr-summ LSAs   : 0
Nssa ext LSAs    : 0                    Area opaque LSAs : 3
Total LSAs       : 9                    LSA Cksum Sum    : 0x28b62
Blackhole Range  : True                 Unknown LSAs     : 0
===============================================================================
```

Table 6: Show Area Output Fields

Label	Description
`Area Id`	A 32-bit integer uniquely identifying an area
`Type`	`NSSA` — this area is configured as an NSSA area
	`Standard` — this area is configured as a standard area (not NSSA or stub)
	`Stub` — this area is configured as a stub area
`SPF Runs`	The number of times that the intra-area route table has been calculated using this area's link-state database
`LSA Count`	The total number of link-state advertisements in this area's link-state database, excluding AS External LSAs
`LSA Cksum Sum`	The 32-bit unsigned sum of the link-state database advertisements LS checksums contained in this area's link-state database. This checksum excludes AS External LSAs (type 5).
`No. of OSPF Areas`	The number of areas configured on the router
`Virtual Links`	The number of virtual links configured through this transit area
`Active IFs`	The active number of interfaces configured in this area
`Area Bdr Rtrs`	The total number of ABRs reachable within this area
`AS Bdr Rtrs`	The total number of ASBRs reachable within this area
`Last SPF Run`	The time that the last intra-area SPF was run on this area
`Router LSAs`	The total number of router LSAs in this area
`Network LSAs`	The total number of network LSAs in this area
`Summary LSAs`	The summary of LSAs in this area
`Asbr-summ LSAs`	The summary of ASBR LSAs in this area
`Nssa-ext LSAs`	The total number of NSSA-EXT LSAs in this area
`Area opaque LSAs`	The total number of opaque LSAs in this area
`Total Nbrs`	The total number of neighbors in this area
`Total IFs`	The total number of interfaces configured in this area
`Total LSAs`	The sum of LSAs in this area excluding autonomous system external LSAs

Table 6: Show Area Output Fields (Continued)

Label	Description
`Blackhole Range`	`False` — no blackhole route is installed for aggregates configured in this area
	`True` — a lowest-priority blackhole route is installed for aggregates configured in this area
`Unknown LSAs`	The total number of unknown LSAs in this area

database

Syntax **database** [**type** {**router** | **network** | **summary** | **asbr-summary** | **external** | **nssa** | **all**}] [**area** *area-id*] [**adv-router** *router-id*] [*link-state-id*] [**detail**]

Context show>router>ospf

Description This command displays information about the OSPF link-state database.

When no command line options are specified, the command displays a summary output for all database entries.

Parameters **type** *keyword* — filters the OSPF link-state database information based on the type specified by *keyword*

type router — displays only router (type 1) LSAs in the link-state database

type network — displays only network (type 2) LSAs in the link-state database

type summary — displays only summary (type 3) LSAs in the link-state database

type asbr-summary — displays only ASBR summary (type 4) LSAs in the link-state database

type external — displays only AS external (type 5) LSAs in the link-state database. External LSAs are maintained globally, not per area. If the display of external links is requested, the area parameter, if present, is ignored.

type nssa — displays only NSSA area-specific AS external (type 7) LSAs in the link-state database

type all — displays all LSAs in the link-state database. The **all** keyword is intended to be used with either the **area** *area-id* or the **adv-router** *router-id* [*link-state-id*] parameters.

area *area-id* — displays link-state database information associated with the specified OSPF *area-id*

adv-router *router-id* [*link-state-id*] — displays link-state database information associated with the specified advertising router. To further narrow the number of items displayed, the *link-state-id* can optionally be specified.

detail — displays detailed information on the link-state database entries

Output The following output is an example of OSPF database information, and Table 7 describes the fields.

Sample Output

```
A:ALU-A# show router ospf database
===============================================================================
OSPF Link State Database (Type : All)
===============================================================================
Type         Area ID    Link State Id   Adv Rtr Id      Age  Sequence   Cksum
-------------------------------------------------------------------------------
Router       0.0.0.0    180.0.0.2       180.0.0.2       1800 0x800000b6 0xf54
Router       0.0.0.0    180.0.0.5       180.0.0.5       1902 0x8000009d 0xcb7c
Router       0.0.0.0    180.0.0.8       180.0.0.8       1815 0x8000009a 0x529b
Router       0.0.0.0    180.0.0.9       180.0.0.9       1156 0x80000085 0xd00f
Router       0.0.0.0    180.0.0.10      180.0.0.10      533  0x8000009d 0x3f1f
Router       0.0.0.0    180.0.0.11      180.0.0.11      137  0x80000086 0xc58f
Router       0.0.0.0    180.0.0.12      180.0.0.12      918  0x8000009d 0x4cf3
Router       0.0.0.0    180.0.0.13      180.0.0.13      1401 0x800000a2 0x879c
Network      0.0.0.0    180.0.53.28     180.0.0.28      149  0x80000083 0xe5cd
Network      0.0.0.0    180.0.54.28     180.0.0.28      1259 0x80000083 0xdad7
Summary      0.0.0.0    180.0.0.15      180.0.0.10      378  0x80000084 0xeba1
Summary      0.0.0.0    180.0.0.15      180.0.0.12      73   0x80000084 0xdfab
Summary      0.0.0.0    180.0.0.18      180.0.0.10      1177 0x80000083 0xcfbb
Summary      0.0.0.1    180.100.25.4    180.0.0.12      208  0x80000091 0x3049
AS Summ      0.0.0.1    180.0.0.8       180.0.0.10      824  0x80000084 0x3d07
AS Summ      0.0.0.1    180.0.0.8       180.0.0.12      1183 0x80000095 0x4bdf
AS Summ      0.0.0.1    180.0.0.9       180.0.0.10      244  0x80000082 0x73cb
AS Ext       n/a        7.1.0.0         180.0.0.23      1312 0x80000083 0x45e7
AS Ext       n/a        7.2.0.0         180.0.0.23      997  0x80000082 0x45e6
AS Ext       n/a        10.20.0.0       180.0.0.23      238  0x80000081 0x2d81
...
-------------------------------------------------------------------------------
No. of LSAs: 339
===============================================================================
A:ALU-A#  show router ospf database detail
===============================================================================
OSPF Link State Database (Type : All) (Detailed)
-------------------------------------------------------------------------------
Router LSA for Area 0.0.0.0
-------------------------------------------------------------------------------
Area Id          : 0.0.0.0              Adv Router Id    : 180.0.0.2
Link State Id    : 180.0.0.2            LSA Type         : Router
Sequence No      : 0x800000b7           Checksum         : 0xd55
Age              : 155                  Length           : 192
Options          : E
Flags            : None                 Link Count       : 14
Link Type (1)    : Point To Point
Nbr Rtr Id (1)   : 180.0.0.13           I/F Address (1)  : 180.0.22.2
No of TOS (1)    : 0                    Metric-0 (1)     : 25
Link Type (2)    : Stub Network
Network (2)      : 180.0.22.0           Mask (2)         : 255.255.255.0
No of TOS (2)    : 0                    Metric-0 (2)     : 25
Link Type (3)    : Point To Point
Nbr Rtr Id (3)   : 180.0.0.12           I/F Address (3)  : 180.0.5.2
No of TOS (3)    : 0                    Metric-0 (3)     : 25
Link Type (4)    : Stub Network
Network (4)      : 180.0.5.0            Mask (4)         : 255.255.255.0
No of TOS (4)    : 0                    Metric-0 (4)     : 25
Link Type (5)    : Point To Point
Nbr Rtr Id (5)   : 180.0.0.8            I/F Address (5)  : 180.0.13.2
No of TOS (5)    : 0                    Metric-0 (5)     : 6
```

```
Link Type (6)    : Stub Network
Network (6)      : 180.0.13.0           Mask (6)         : 255.255.255.0
No of TOS (6)    : 0                    Metric-0 (6)     : 6
Link Type (7)    : Point To Point
Nbr Rtr Id (7)   : 180.0.0.5            I/F Address (7)  : 180.0.14.2
No of TOS (7)    : 0                    Metric-0 (7)     : 6
Link Type (8)    : Stub Network
Network (8)      : 180.0.14.0           Mask (8)         : 255.255.255.0
No of TOS (8)    : 0                    Metric-0 (8)     : 6
Link Type (9)    : Point To Point
Nbr Rtr Id (9)   : 180.0.0.11           I/F Address (9)  : 180.0.17.2
No of TOS (9)    : 0                    Metric-0 (9)     : 25
Link Type (10)   : Stub Network
Network (10)     : 180.0.17.0           Mask (10)        : 255.255.255.0
No of TOS (10)   : 0                    Metric-0 (10)    : 25
Link Type (11)   : Stub Network
Network (11)     : 180.0.0.2            Mask (11)        : 255.255.255.255
No of TOS (11)   : 0                    Metric-0 (11)    : 1
Link Type (12)   : Stub Network
Network (12)     : 180.0.18.0           Mask (12)        : 255.255.255.0
No of TOS (12)   : 0                    Metric-0 (12)    : 24
Link Type (13)   : Point To Point
Nbr Rtr Id (13)  : 180.0.0.10           I/F Address (13) : 180.0.3.2
No of TOS (13)   : 0                    Metric-0 (13)    : 25
Link Type (14)   : Stub Network
Network (14)     : 180.0.3.0            Mask (14)        : 255.255.255.0
No of TOS (14)   : 0                    Metric-0 (14)    : 25
-------------------------------------------------------------------------------



AS Ext LSA for Network 180.0.0.14
-------------------------------------------------------------------------------
Area Id          : N/A                  Adv Router Id    : 180.0.0.10
Link State Id    : 180.0.0.14           LSA Type         : AS Ext
Sequence No      : 0x80000083           Checksum         : 0xa659
Age              : 2033                 Length           : 36
Options          : E
Network Mask     : 255.255.255.255      Fwding Address   : 180.1.6.15
Metric Type      : Type 2               Metric-0         : 4
Ext Route Tag    : 0
-------------------------------------------------------------------------------
...
A:ALU-A#
```

Table 7: Show Database Output Fields

Label	Description
Type LSA Type	The LSA type
Area Id	The OSPF area identifier
Link State Id	The link-state ID is an LSA type-specific field containing either a number to distinguish several LSAs from the same router, an interface ID, or a router ID; it identifies the piece of the routing domain being described by the advertisement
Adv Rtr Id/ Adv Router Id	The router identifier of the router advertising the LSA
Age	The age of the link-state advertisement in seconds
Sequence/ Sequence No	The signed 32-bit integer sequence number
Cksum/ Checksum	The 32-bit unsigned sum of the link-state advertisements' LS checksums
No. of LSAs	The number of LSAs displayed
Options	EA — external attribute LSA support
	DC — demand circuit support
	R — if clear, a node can participate in OSPF topology distribution without being used to forward transit traffic
	N — type 7 LSA support
	MC — multicast support (not applicable)
	E — external routes support
	V6 — not applicable
Prefix Options	P — propagate NSSA LSA
	MC — multicast support (not applicable)
	LA — local address capability; if set, the prefix is an IPv6 interface address of the advertising router (not applicable)
	NU — no unicast capability; if set, the prefix is excluded from IPv6 unicast calculations (not applicable)
Flags	None — no flags set
	V — the router is an endpoint for one or more fully adjacent virtual links having the described area as the transit area

Table 7: Show Database Output Fields (Continued)

Label	Description
`Flags (cont.)`	`E` — the router is an AS Boundary Router
	`B` — the router is an Area Border Router
`Link Count`	The number of links advertised in the LSA
`Link Type (`*`n`*`)`	The link type of the *n*th link in the LSA
`Network (`*`n`*`)`	The network address of the *n*th link in the LSA
`Metric-0 (`*`n`*`)`	The cost metric of the *n*th link in the LSA

interface

Syntax **interface** [**area** *area-id*] [**detail**]
interface [*ip-int-name* | *ip-address*] [**detail**]

Context show>router>ospf

Description This command displays the details of the OSPF interface, which can be identified by IP address or IP interface name. If neither is specified, all in-service interfaces are displayed.

The **area** option displays all interfaces configured in the specified area.

The **detail** option produces a great amount of data. It is recommended that this option be used only when requesting a specific interface.

Parameters *area-id* — displays all interfaces configured in this area

ip-int-name — displays only the interface identified by this interface name

ip-address — displays only the interface identified by this IP address

detail — displays detailed information on the interface

Output The following outputs are examples of OSPF interface information:

- OSPF standard interface information (Sample Output, Table 8)
- OSPF detailed interface information (Sample Output, Table 9)

Sample Output

```
A:ALU-A# show router ospf interface

===============================================================================
OSPF Interfaces
===============================================================================
If Name               Area Id         Designated Rtr  Bkup Desig Rtr  Adm  Oper
-------------------------------------------------------------------------------
system                0.0.0.1         10.10.10.104    0.0.0.0         Up   DR
to-103                0.0.0.20        0.0.0.0         0.0.0.0         Up   Down
-------------------------------------------------------------------------------
No. of OSPF Interfaces: 2
===============================================================================
```

Table 8: Show Interface Output Fields

Label	Description
`If Name`	The interface name
`Area Id`	A 32-bit integer uniquely identifying the area to which this interface is connected. Area ID 0.0.0.0 is used for the OSPF backbone.
`Designated RTR`	The IP interface address of the router identified as the designated router for the network in which this interface is configured Set to 0.0.0.0 if there is no designated router
`Bkup Desig Rtr`	The IP interface address of the router identified as the backup designated router for the network in which this interface is configured Set to 0.0.0.0 if there is no backup designated router
`Adm`	`Dn` — OSPF on this interface is administratively shut down
	`Up` — OSPF on this interface is administratively enabled
`Opr`	`Down` — the initial interface state. In this state, the lower-level protocols have indicated that the interface is unusable.
	`Wait` — the router is trying to determine the identity of the (backup) designated router for the network
	`PToP` — the interface is operational, and connects either to a physical point-to-point network or to a virtual link
	`DR` — this router is the designated router for this network
	`BDR` — this router is the backup designated router for this network

Table 8: Show Interface Output Fields (Continued)

Label	Description
Opr (cont.)	ODR — the interface is operational and part of a broadcast or NBMA network on which another router has been selected to be the designated router
No. of OSPF Interfaces	The number of interfaces listed

Sample Output

```
A:ALU-A# show router ospf interface detail

===============================================================================
OSPF Interface (Detailed) :  system
===============================================================================
-------------------------------------------------------------------------------
Configuration
-------------------------------------------------------------------------------
IP Address       : 10.10.10.104
Area Id          : 0.0.0.1              Priority         : 1
Hello Intrvl     : 10 sec               Rtr Dead Intrvl  : 40 sec
Retrans Intrvl   : 5 sec                Poll Intrvl      : 120 sec
Cfg Metric       : 0                    Advert Subnet    : True
Transit Delay    : 1                    Auth Type        : None
Passive          : True                 Cfg MTU          : 0
-------------------------------------------------------------------------------
State
-------------------------------------------------------------------------------
Admin Status     : Enabled              Oper State       : Designated Rtr
Designated Rtr   : 10.10.10.104         Backup Desig Rtr : 0.0.0.0
IF Type          : Broadcast            Network Type     : Transit
Oper MTU         : 1500                 Last Enabled     : 01/26/2009 23:28:25
Oper Metric      : 0                    Bfd Enabled      : No
Te Metric        : 0                    Te State         : Down
Admin Groups     : None
Ldp Sync         : outOfService         Ldp Sync Wait    : Disabled
Ldp Timer State  : Disabled             Ldp Tm Left      : 0
-------------------------------------------------------------------------------
Statistics
-------------------------------------------------------------------------------
Nbr Count        : 0                    If Events        : 2
Tot Rx Packets   : 0                    Tot Tx Packets   : 0
Rx Hellos        : 0                    Tx Hellos        : 0
Rx DBDs          : 0                    Tx DBDs          : 0
Rx LSRs          : 0                    Tx LSRs          : 0
Rx LSUs          : 0                    Tx LSUs          : 0
Rx LS Acks       : 0                    Tx LS Acks       : 0
Retransmits      : 0                    Discards         : 0
Bad Networks     : 0                    Bad Virt Links   : 0
Bad Areas        : 0                    Bad Dest Addrs   : 0
Bad Auth Types   : 0                    Auth Failures    : 0
Bad Neighbors    : 0                    Bad Pkt Types    : 0
Bad Lengths      : 0                    Bad Hello Int.   : 0
Bad Dead Int.    : 0                    Bad Options      : 0
Bad Versions     : 0                    Bad Checksums    : 0
LSA Count        : 0                    LSA Checksum     : 0x0
-------------------------------------------------------------------------------
```

Table 9: Show Detailed Interface Output Fields

Label	Description
`Interface`	The IP address of this OSPF interface
`IP Address`	The IP address and mask of this OSPF interface
`Interface Name`	The interface name
`Area Id`	A 32-bit integer uniquely identifying the area to which this interface is connected. Area ID 0.0.0.0 is used for the OSPF backbone.
`Priority`	The priority of this interface. Used in multi-access networks, this field is used in the designated router election algorithm.
`Hello Intrvl`	The length of time, in seconds, between the Hello packets that the router sends on the interface. This value must be the same for all routers attached to a common network.
`Rtr Dead Intrvl`	The number of seconds that a router's Hello packets have not been seen before its neighbors declare the router down. This should be some multiple of the Hello interval. This value must be the same for all routers attached to a common network.
`Retrans Intrvl`	The number of seconds between link-state advertisement retransmissions, for adjacencies belonging to this interface. This value is also used when retransmitting database description and link-state request packets.
`Poll Intrvl`	The larger time interval, in seconds, between the Hello packets sent to an inactive non-broadcast multi-access neighbor
`Cfg Metric`	The metric to be advertised for this interface
`Advert Subnet`	`False` — when a point-to-point interface is configured as false, then the subnet is not advertised and the endpoints are advertised as host routes
	`True` — when a point-to-point interface is configured as true, then the subnet is advertised
`Transit Delay`	The estimated number of seconds it takes to transmit a link-state update packet over this interface

Table 9: Show Detailed Interface Output Fields (Continued)

Label	Description
Auth Type	Identifies the authentication procedure to be used for the packet
	None — routing exchanges over the network/subnet are not authenticated
	Simple — a 64-bit field is configured on a per-network basis. All packets sent on a particular network must have this configured value in their OSPF header 64-bit authentication field. This essentially serves as a “clear” 64-bit password.
	MD5 — a shared secret key is configured on all routers attached to a common network or subnet. For each OSPF protocol packet, the key is used to generate and verify a “message digest” that is appended to the end of the OSPF packet.
Passive	False — this interfaces operates as a normal OSPF interface with regard to adjacency forming and network and link behavior
	True — no OSPF Hellos will be sent out on this interface and the router advertises this interface as a stub network or link in its router LSAs
Cfg MTU	The desired size of the largest packet that can be sent or received on this OSPF interface, specified in octets. This size does include the underlying IP header length, but not the underlying layer headers and trailers.
Admin Status	Disabled — OSPF on this interface is administratively shut down
	Enabled — OSPF on this interface is administratively enabled
Oper State	Down — the initial interface state. In this state, the lower-level protocols have indicated that the interface is unusable.
	Waiting — the router is trying to determine the identity of the (backup) designated router for the network
	Point To Point — the interface is operational and connects either to a physical point-to-point network or to a virtual link
	Designated Rtr — this router is the designated router for this network
	Other Desig Rtr — the interface is operational and part of a broadcast or NBMA network on which another router has been selected to be the designated router
	Backup Desig Rtr — this router is the backup designated router for this network

Table 9: Show Detailed Interface Output Fields (Continued)

Label	Description
`Designated Rtr`	The IP interface address of the router identified as the designated router for the network in which this interface is configured Set to 0.0.0.0 if there is no designated router
`Backup Desig Rtr`	The IP interface address of the router identified as the backup designated router for the network in which this interface is configured Set to 0.0.0.0 if there is no backup designated router
`IF Type`	`Broadcast` — LANs, such as Ethernet
	`NBMA` — X.25, Frame Relay and similar technologies
	`Point-To-Point` — links that are definitively point-to-point
`Network Type`	`Stub` — OPSF has not established a neighbor relationship with any other OSPF router on this network; therefore, only traffic sourced or destined for this network will be routed to this network
	`Transit` — OPSF has established at least one neighbor relationship with another OSPF router on this network; therefore, traffic en route to other networks may be routed via this network
`Oper MTU`	The operational size of the largest packet that can be sent or received on this OSPF interface, specified in octets. This size includes the underlying IP header length, but not the underlying layer headers and trailers.
`Last Enabled`	The time that this interface was last enabled to run OSPF on this interface
`Te Metric`	The TE metric configured for this interface. This metric is flooded out in the TE metric sub-TLV in the OSPF TE LSAs. Depending on the configuration, either the TE metric value or the native OSPF metric value is used in CSPF computations.
`Te State`	The MPLS interface TE status from OSPF standpoint
`Admin Groups`	The bit-map inherited from the MPLS interface that identifies the admin groups to which this interface belongs
`Ldp Sync`	Specifies whether the IGP-LDP synchronization feature is enabled or disabled on all interfaces participating in the OSPF routing protocol
`Ldp Sync Wait`	The time to wait for the LDP adjacency to come up
`Ldp Timer State`	The state of the LDP sync time left on the OSPF interface

Table 9: Show Detailed Interface Output Fields (Continued)

Label	Description
`Ldp Tm Left`	The time left before OSPF reverts back to advertising normal metrics for this interface
`Nbr Count`	The number of OSPF neighbors on the network for this interface
`If Events`	The number of times this OSPF interface has changed its state, or an error has occurred since this interface was last enabled
`Tot Rx Packets`	The total number of OSPF packets received on this interface since this interface was last enabled
`Tot Tx Packets`	The total number of OSPF packets transmitted on this interface since this interface was last enabled
`Rx Hellos`	The total number of OSPF Hello packets received on this interface since this interface was last enabled
`Tx Hellos`	The total number of OSPF Hello packets transmitted on this interface since this interface was last enabled
`Rx DBDs`	The total number of OSPF database description packets received on this interface since this interface was last enabled
`Tx DBDs`	The total number of OSPF database description packets transmitted on this interface since this interface was last enabled
`Rx LSRs`	The total number of Link-State Requests (LSRs) received on this interface since this interface was last enabled
`Tx LSRs`	The total number of Link-State Requests (LSRs) transmitted on this interface since this interface was last enabled
`Rx LSUs`	The total number of Link-State Updates (LSUs) received on this interface since this interface was last enabled
`Tx LSUs`	The total number of Link-State Updates (LSUs) transmitted on this interface since this interface was last enabled
`Rx LS Acks`	The total number of Link-State Acknowledgements received on this interface since this interface was last enabled
`Tx LS Acks`	The total number of Link-State Acknowledgements transmitted on this interface since this interface was last enabled
`Retransmits`	The total number of OSPF retransmits sent on this interface since this interface was last enabled
`Discards`	The total number of OSPF packets discarded on this interface since this interface was last enabled

Table 9: Show Detailed Interface Output Fields (Continued)

Label	Description
`Bad Networks`	The total number of OSPF packets received with invalid network or mask since this interface was last enabled
`Bad Virt Links`	The total number of OSPF packets received on this interface that are destined for a virtual link that does not exist since this interface was last enabled
`Bad Areas`	The total number of OSPF packets received with an area mismatch since this interface was last enabled
`Bad Dest Addrs`	The total number of OSPF packets received with the incorrect IP destination address since this interface was last enabled
`Bad Auth Types`	The total number of OSPF packets received with an invalid authorization type since this interface was last enabled
`Auth Failures`	The total number of OSPF packets received with an invalid authorization key since this interface was last enabled
`Bad Neighbors`	The total number of OSPF packets received where the neighbor information does not match the information this router has for the neighbor since this interface was last enabled
`Bad Pkt Types`	The total number of OSPF packets received with an invalid OSPF packet type since this interface was last enabled
`Bad Lengths`	The total number of OSPF packets received on this interface with a total length not equal to the length given in the packet itself since this interface was last enabled
`Bad Hello Int.`	The total number of OSPF packets received where the hello interval given in the packet was not equal to that configured on this interface since this interface was last enabled
`Bad Dead Int.`	The total number of OSPF packets received where the dead interval given in the packet was not equal to that configured on this interface since this interface was last enabled
`Bad Options`	The total number of OSPF packets received with an option that does not match those configured for this interface or area since this interface was last enabled
`Bad Versions`	The total number of OSPF packets received with bad OSPF version numbers since this interface was last enabled
`Bad Checksums`	The total number of OSPF packets received with bad checksums since this interface was last enabled

Table 9: Show Detailed Interface Output Fields (Continued)

Label	Description
`LSA Count`	The total number of link-state advertisements in this area's link-state database, excluding AS External LSAs
`LSA Checksum`	The 32-bit unsigned sum of the link-state database advertisements' LS checksums contained in this area's link-state database. This checksum excludes AS External LSAs (type 5).

neighbor

Syntax **neighbor** [*ip-int-name*] [*router-id*] [**detail**]
neighbor [**remote** *ip-address*] [**detail**]

Context show>router>ospf

Description This command displays all neighbor information or all information on neighbors of a router identified by interface name or router ID.

The **detail** option produces a large amount of data. It is recommended that this option be used only when requesting a specific neighbor.

Parameters *ip-int-name* — displays neighbor information only for neighbors of the interface identified by the interface name

router-id — displays neighbor information for the neighbor identified by the the specified router ID

detail — displays detailed information on the interface

ip-address — displays information for a far-end neighbor, identified by IP address

Output The following outputs are examples of OSPF neighbor information:

- OSPF standard neighbor information (Sample Output, Table 10)
- OSPF detailed neighbor information (Sample Output, Table 11)

Sample Output

```
A:ALU-A# show router ospf neighbor
===============================================================================
OSPF Neighbors
===============================================================================
Interface-Name                   Rtr Id          State      Pri  RetxQ   TTL
-------------------------------------------------------------------------------
pc157-2/1                        10.13.8.158     Full       1    0       37
pc157-2/2                        10.13.7.165     Full       100  0       33
pc157-2/3                        10.13.6.188     Full       1    0       38
-------------------------------------------------------------------------------
No. of Neighbors: 3
===============================================================================
A:ALU-A#
```

Table 10: Show Neighbor Output Fields

Label	Description
Interface Name	The interface name or IP address this neighbor is using in its IP source address. Note that, on links with no address, this will not be 0.0.0.0, but the address of another of the neighbor's interfaces.
Rtr Id	A 32-bit integer uniquely identifying the neighboring router in the Autonomous System
State	Down — the initial state of a neighbor conversation. It indicates that there has been no recent information received from the neighbor.
	Attempt — this state is only valid for neighbors attached to NBMA networks. It indicates that no recent information has been received from the neighbor, but that a more concerted effort should be made to contact the neighbor.
	Init — in this state, a Hello packet has recently been seen from the neighbor. However, bidirectional communication has not yet been established with the neighbor (that is, the router itself did not appear in the neighbor's Hello packet).
	Two Way — in this state, communication between the two routers is bidirectional
	ExchStart — the first step in creating an adjacency between the two neighboring routers. The goal of this step is to decide which router is the master, and to decide upon the initial database descriptor sequence number.
	Exchange — in this state, the router is describing its entire link-state database by sending database description packets to the neighbor

Table 10: Show Neighbor Output Fields (Continued)

Label	Description
`State (cont.)`	`Loading` — in this state, Link-State Request packets are sent to the neighbor asking for the more recent LSAs that have been discovered (but not yet received) in the Exchange state
	`Full` — in this state, the neighboring routers are fully adjacent. These adjacencies will now appear in router-LSAs and network-LSAs.
`Pri`	The priority of this neighbor in the designated router election algorithm. The value 0 signifies that the neighbor is not eligible to become the designated router on this particular network.
`RetxQ`	The current length of the retransmission queue
`TTL`	The time until this neighbor is declared down; this timer is set to the dead router interval when a valid Hello packet is received from the neighbor
`No. of Neighbors`	The number of adjacent OSPF neighbors on this interface

Sample Output

```
A:ALU-A# show router ospf neighbor 10.13.8.150 detail
===============================================================================
OSPF Neighbors
-------------------------------------------------------------------------------
Neighbor Rtr Id  : 10.13.8.158         Interface: pc157-2/1
-------------------------------------------------------------------------------
Neighbor IP Addr : 10.16.1.8
Local IF IP Addr : 10.16.1.7
Area Id          : 0.0.0.0
Designated Rtr   : 0.0.0.0             Backup Desig Rtr : 0.0.0.0
Neighbor State   : Full                Priority         : 1
Retrans Q Length : 0                   Options          : -E--O-
Events           : 4                   Last Event Time  : 05/06/2008 00:11:16
Up Time          : 1d 18:20:20         Time Before Dead : 38 sec
Bad Nbr States   : 1                   LSA Inst fails   : 0
Bad Seq Nums     : 0                   Bad MTUs         : 0
Bad Packets      : 0                   LSA not in LSDB  : 0
Option Mismatches: 0                   Nbr Duplicates   : 0
Num Restarts     : 0                   Last Restart at  : Never
-------------------------------------------------------------------------------
A:ALU-A#
```

Table 11: Show Detailed Neighbor Output Fields

Label	Description
`Neighbor IP Addr`	The IP address this neighbor is using in its IP source address. Note that, on links with no IP address, this will not be 0.0.0.0, but the address of another of the neighbor's interfaces.
`Local IF IP Addr`	The IP address of this OSPF interface
`Area Id`	A 32-bit integer uniquely identifying the area to which this interface is connected. Area ID 0.0.0.0 is used for the OSPF backbone.
`Designated Rtr`	The IP interface address of the router identified as the designated router for the network in which this interface is configured Set to 0.0.0.0 if there is no designated router
`Neighbor Rtr Id`	A 32-bit integer uniquely identifying the neighboring router in the AS
`Neighbor State`	`Down` — the initial state of a neighbor conversation. It indicates that there has been no recent information received from the neighbor.
	`Attempt` — this state is only valid for neighbors attached to NBMA networks. It indicates that no recent information has been received from the neighbor, but that a more concerted effort should be made to contact the neighbor.
	`Init` — in this state, a Hello packet has recently been seen from the neighbor. However, bidirectional communication has not yet been established with the neighbor (that is, the router itself did not appear in the neighbor's Hello packet).
	`Two Way` — in this state, communication between the two routers is bidirectional
	`Exchange start` — the first step in creating an adjacency between the two neighboring routers. The goal of this step is to decide which router is the master, and to decide upon the initial database descriptor sequence number.
	`Exchange` — in this state, the router is describing its entire link-state database by sending database description packets to the neighbor
	`Loading` — in this state, Link-State Request packets are sent to the neighbor asking for the more recent LSAs that have been discovered (but not yet received) in the Exchange state
	`Full` — in this state, the neighboring routers are fully adjacent. These adjacencies will now appear in router-LSAs and network-LSAs.
`Priority`	The priority of this neighbor in the designated router election algorithm. The value 0 signifies that the neighbor is not eligible to become the designated router on this particular network.

Table 11: Show Detailed Neighbor Output Fields (Continued)

Label	Description
`Retrans Q Length`	The current length of the retransmission queue
`Options`	`E` — external routes support
	`MC` — multicast support (not applicable)
	`N/P` — type 7 LSA support
	`EA` — external attribute LSA support
	`DC` — demand circuit support
	`O` — opaque LSA support
`Backup Desig Rtr`	The IP interface address of the router identified as the backup designated router for the network in which this interface is configured Set to 0.0.0.0 if there is no backup designated router
`Events`	The number of times this neighbor relationship has changed state, or an error has occurred
`Last Event Time`	The time that the last event occurred that affected the adjacency to the neighbor
`Up Time`	The uninterrupted time, in hundredths of seconds, that the adjacency to this neighbor has been up. To evaluate when the last state change occurred, see last event time.
`Time Before Dead`	The time until this neighbor is declared down; this timer is set to the dead router interval when a valid Hello packet is received from the neighbor
`Bad Nbr States`	The total number of OSPF packets received when the neighbor state was not expecting to receive this packet type since this interface was last enabled
`LSA Inst fails`	The total number of times that an LSA could not be installed into the link-state database due to a resource allocation issue since this interface was last enabled
`Bad Seq Nums`	The total number of times that a database description packet was received with a sequence number mismatch since this interface was last enabled
`Bad MTUs`	The total number of times that the MTU in a received database description packet was larger than the MTU of the receiving interface since this interface was last enabled

Table 11: Show Detailed Neighbor Output Fields (Continued)

Label	Description
`Bad Packets`	The total number of times that an LS update was received with an illegal LS type or an option mismatch since this interface was last enabled
`LSA not in LSDB`	The total number of times that an LS request was received for an LSA not installed in the LSDB of this router since this interface was last enabled
`Option Mismatches`	The total number of times that an LS update was received with an option mismatch since this interface was last enabled
`Nbr Duplicates`	The total number of times that a duplicate database description packet was received during the exchange state since this interface was last enabled

range

Syntax **range** [*area-id*]

Context show>router>ospf

Description This command displays ranges of addresses on an ABR for the purpose of route summarization or suppression.

Parameters *area-id* — displays the configured ranges for the specified area

Output The following output is an example of OSPF range information, and Table 12 describes the fields.

Sample Output

```
A:ALU-A# show router ospf range
==========================================================
OSPF Ranges
==========================================================
Area Id         Prefix       Advertise   LSDB Type
----------------------------------------------------------
No. of Ranges: 0
==========================================================
A:ALU-A#
```

Table 12: Show Area Range Output Fields

Label	Description
Area Id	A 32-bit integer uniquely identifying an area. Area ID 0.0.0.0 is used for the OSPF backbone.
Prefix	The mask for the range expressed as a decimal integer mask length or in dotted-decimal notation
Advertise	False — the specified address/mask is not advertised outside the area
	True — the specified address/mask is advertised outside the area
LSDB Type	NSSA — this range was specified in the NSSA context, and specifies that the range applies to external routes (via type 7 LSAs) learned within the NSSA when the routes are advertised to other areas as type 5 LSAs
	Summary — this range was not specified in the NSSA context; the range applies to summary LSAs even if the area is an NSSA

spf

Syntax **spf**

Context show>router>ospf

Description This command displays statistics of shortest path first (SPF) calculations.

Output The following output is an example of SPF information, and Table 13 describes the fields.

Sample Output

```
A:ALU-A# show router ospf spf
===============================================================================
OSPF SPF Statistics
===============================================================================
 Total SPF Runs          :  109
 Last Full SPF run @     :  11/07/2008 18:43:07
 Last Full SPF Time      :  < 0.01 secs
      Intra SPF Time     :  < 0.01 secs
      Inter SPF Time     :  < 0.01 secs
      Extern SPF Time    :  < 0.01 secs
      RTM Updt Time      :  < 0.01 secs

 Min/Avg/Max Full SPF Times   :  0.02/0.00/0.06 secs
 Min/Avg/Max RTM Updt Times   :  0.02/0.00/0.06 secs

 Total Sum Incr SPF Runs :  333
 Total Ext Incr SPF Runs :  0
===============================================================================
```

Table 13: Show SPF Output Fields

Label	Description
`Total SPF Runs`	The total number of incremental SPF runs triggered by new or updated LSAs
`Last Full SPF run @`	The date and time that the external OSPF SPF was last run
`Last Full SPF Time`	The length of time, in seconds, when the last full SPF was run
`Intra SPF Time`	The time that intra-area SPF was last run on this area
`Inter SPF Time`	The total number of incremental SPF runs triggered by new or updated type 3 and type 4 summary LSAs
`Extern SPF Time`	The total number of incremental SPF runs triggered by new or updated type 5 external LSAs
`RTM Updt Time`	The time, in hundredths of seconds, used to perform a total SPF calculation
`Min/Avg/Max Full SPF Times`	`Min` — the minimum time, in hundredths of seconds, used to perform a total SPF calculation
	`Avg` — the average time, in hundredths of seconds, of all the SPF calculations performed by this OSPF router
	`Max` — the maximum time, in hundredths of seconds, used to perform a total SPF calculation
`Min/Avg/Max RTM Updt Times`	`Min` — the minimum time, in hundredths of seconds, used to perform an RTM update **Note**: the RTM update is performed after the SPF calculation. The update is used to inform the routing table manager of any route or cost changes from the latest SPF calculation.
	`Avg` — the average time, in hundredths of seconds, of all the RTM updates performed by this OSPF router
	`Max` — the maximum time, in hundredths of seconds, used to perform an RTM update
`Total Sum Incr SPF Runs`	The total number of incremental SPF runs triggered by new or updated type 3 and type 4 summary LSAs
`Total Ext Incr SPF Runs`	The total number of incremental SPF runs triggered by new or updated type 5 external LSAs

statistics

Syntax **statistics**

Context show>router>ospf

Description This command displays the global OSPF statistics.

Output The following output is an example of OSPF statistical information, and Table 14 describes the fields.

Sample Output

```
A:ALU-A# show router ospf statistics
===============================================================================
OSPF Statistics
===============================================================================
Rx Packets         : 308462           Tx Packets         : 246800
Rx Hellos          : 173796           Tx Hellos          : 149062
Rx DBDs            : 67               Tx DBDs            : 48
Rx LSRs            : 21               Tx LSRs            : 19
Rx LSUs            : 105672           Tx LSUs            : 65530
Rx LS Acks         : 28906            Tx LS Acks         : 32141
New LSAs Recvd     : 38113            New LSAs Orig      : 21067
Ext LSAs Count     : 17               No of Areas        : 3
No of Interfaces   : 327              No of Neighbors    : 0
Retransmits        : 46               Discards           : 0
Bad Networks       : 0                Bad Virt Links     : 0
Bad Areas          : 0                Bad Dest Addrs     : 0
Bad Auth Types     : 0                Auth Failures      : 0
Bad Neighbors      : 0                Bad Pkt Types      : 0
Bad Lengths        : 0                Bad Hello Int.     : 0
Bad Dead Int.      : 0                Bad Options        : 0
Bad Versions       : 0                Bad Checksums      : 0
Failed SPF Attempts: 0
CSPF Requests      : 0                CSPF Request Drops : 0
CSPF Path Found    : 0                CSPF Path Not Found: 0
===============================================================================
A:ALU-A#
```

Table 14: Show OSPF Statistics Output Fields

Label	Description
`Rx Packets`	The total number of OSPF packets received on all OSPF enabled interfaces
`Tx Packets`	The total number of OSPF packets transmitted on all OSPF enabled interfaces
`Rx Hellos`	The total number of OSPF Hello packets received on all OSPF enabled interfaces
`Tx Hellos`	The total number of OSPF Hello packets transmitted on all OSPF enabled interfaces
`Rx DBDs`	The total number of OSPF database description packets received on all OSPF enabled interfaces
`Tx DBDs`	The total number of OSPF database description packets transmitted on all OSPF enabled interfaces
`Rx LSRs`	The total number of OSPF Link-State Requests (LSRs) received on all OSPF enabled interfaces
`Tx LSRs`	The total number of OSPF Link-State Requests (LSRs) transmitted on all OSPF enabled interfaces
`Rx LSUs`	The total number of OSPF Link-State Updates (LSUs) received on all OSPF enabled interfaces
`Tx LSUs`	The total number of OSPF Link-State Updates (LSUs) transmitted on all OSPF enabled interfaces
`Rx LS Acks`	The total number of OSPF Link-State Acknowledgements received on all OSPF enabled interfaces
`New LSAs Recvd`	The total number of new OSPF Link-State Advertisements received on all OSPF enabled interfaces
`New LSAs Orig`	The total number of new OSPF Link-State Advertisements originated on all OSPF enabled interfaces
`Ext LSAs Count`	The total number of OSPF External Link-State Advertisements
`No of Areas`	The number of areas configured for OSPF (maximum 4)
`No of Interfaces`	The number of interfaces configured for OSPF on the router
`No of Neighbors`	The number of adjacent OSPF neighbors on this interface
`Retransmits`	The total number of OSPF Retransmits transmitted on all OSPF enabled interfaces

Table 14: Show OSPF Statistics Output Fields (Continued)

Label	Description
`Discards`	The total number of OSPF packets discarded on all OSPF enabled interfaces
`Bad Networks`	The total number of OSPF packets received on all OSPF enabled interfaces with invalid network or mask
`Bad Virt Links`	The total number of OSPF packets received on all OSPF enabled interfaces that are destined for a virtual link that does not exist
`Bad Areas`	The total number of OSPF packets received on all OSPF enabled interfaces with an area mismatch
`Bad Dest Addrs`	The total number of OSPF packets received on all OSPF enabled interfaces with the incorrect IP destination address
`Bad Auth Types`	The total number of OSPF packets received on all OSPF enabled interfaces with an invalid authorization type
`Auth Failures`	The total number of OSPF packets received on all OSPF enabled interfaces with an invalid authorization key
`Bad Neighbors`	The total number of OSPF packets received on all OSPF enabled interfaces where the neighbor information does not match the information this router has for the neighbor
`Bad Pkt Types`	The total number of OSPF packets received on all OSPF enabled interfaces with an invalid OSPF packet type
`Bad Lengths`	The total number of OSPF packets received on all OSPF enabled interfaces with a total length not equal to the length given in the packet itself
`Bad Hello Int.`	The total number of OSPF packets received on all OSPF enabled interfaces where the hello interval given in the packet was not equal to that configured for the respective interface
`Bad Dead Int.`	The total number of OSPF packets received on all OSPF enabled interfaces where the dead interval given in the packet was not equal to that configured for the respective interface
`Bad Options`	The total number of OSPF packets received on all OSPF enabled interfaces with an option that does not match those configured for the respective interface or area
`Bad Versions`	The total number of OSPF packets received on all OSPF enabled interfaces with bad OSPF version numbers
`Bad Checksums`	The total number of OSPF packets received with bad checksums since this interface was last enabled

Table 14: Show OSPF Statistics Output Fields (Continued)

Label	Description
`Failed SPF Attempts`	The total number of failed SPF calculation attempts
`CSPF Requests`	The total number of constraint-based SPF requests
`CSPF Request Drops`	The total number of constraint-based SPF requests dropped
`CSPF Path Found`	A path that fulfills the set of constraints defined in MPLS traffic engineering
`CSPF Path Not Found`	A path that does not fulfill the set of constraints defined in MPLS traffic engineering

status

Syntax **status**

Context show>router>ospf

Description This command displays the general status of OSPF.

Output The following output is an example of OSPF status information, and Table 15 describes the fields.

Sample Output

```
A:ALU-A# show router ospf status
===============================================================================
OSPF Status
===============================================================================
OSPF Cfg Router Id           : 0.0.0.0
OSPF Oper Router Id          : 10.10.10.104
OSPF Version                 : 2
OSPF Admin Status            : Enabled
OSPF Oper Status             : Enabled
Preference                   : 10
External Preference          : 150
Backbone Router              : True
Area Border Router           : False
AS Border Router             : False
Opaque LSA Support           : True
Traffic Engineering Support  : False
RFC 1583 Compatible          : True
Demand Exts Support          : False
In Overload State            : False
In External Overflow State   : False
Exit Overflow Interval       : 0
Last Overflow Entered        : Never
Last Overflow Exit           : Never
External LSA Limit           : -1
```

```
Reference Bandwidth          : 100,000,000 Kbps
Init SPF Delay               : 1000 msec
Sec SPF Delay                : 1000 msec
Max SPF Delay                : 10000 msec
Min LS Arrival Interval      : 1000 msec
Init LSA Gen Delay           : 5000 msec
Sec LSA Gen Delay            : 5000 msec
Max LSA Gen Delay            : 5000 msec
Last Ext SPF Run             : Never
Ext LSA Cksum Sum            : 0x0
OSPF Last Enabled            : 01/12/2009 15:32:11
Export Policies              : None
OSPF Ldp Sync Admin Status   : Enabled
===============================================================================
A:ALU-A#
```

Table 15: Show OSPF Status Output Fields

Label	Description
`OSPF Cfg Router Id`	The router ID configured for the router
`OSPF Oper Router ID`	The operational router ID. The 7705 SAR defaults to the system IP address or, if not configured, the last 4 bytes of the system MAC address.
`OSPF Version`	The current version number of the OSPF protocol: 2
`OSPF Admin Status`	`Disabled` — the OSPF process is disabled on all interfaces
	`Enabled` — the OSPF process is active on at least one interface
`OSPF Oper Status`	`Disabled` — the OSPF process is not operational on all interfaces
	`Enabled` — the OSPF process is operational on at least one interface
`Preference`	The route preference for OSPF internal routes
`External Preference`	The route preference for OSPF external routes
`Backbone Router`	`False` — this router is not configured as an OSPF backbone router
	`True` — this router is configured as an OSPF backbone router
`Area Border Router`	`False` — this router is not configured as an area border router
	`True` — this router is configured as an area border router
`AS Border Router`	`False` — this router is not configured as an Autonomous System border (boundary) router
	`True` — this router is configured as an Autonomous System border (boundary) router

Table 15: Show OSPF Status Output Fields (Continued)

Label	Description
Opaque LSA Support	False — this router does not support opaque LSAs
	True — this router supports opaque LSAs
Traffic Engineering Support	False — this router does not support traffic engineering
	True — this router supports traffic engineering
RFC 1583 Compatible	False — this router is not RFC 1583 compatible
	True — this router is RFC 1583 compatible
Demand Exts Support	False — this router does not demand external route support
	True — this router does demand external route support
In Overload State	False — this router is not in an overload state
	True — this router is in an overload state
In External Overflow State	False — this router is not in an external overflow state
	True — this router is in an external overflow state
Exit Overflow Interval	The time to wait before the router exits the overflow state
Last Overflow Entered	Indicates when the router last entered an overflow state
Last Overflow Exit	Indicates when the router last exited an overflow state
External LSA limit	The number of external LSAs allowed
Reference bandwidth	The configured reference bandwidth, in kilobits per second
Init SPF Delay	The initial SPF calculation delay
Sec SPF Delay	The SPF calculation delay between the first and second calculations
Max SPF Delay	The maximum interval between two consecutive SPF calculations
Min LS Arrival Interval	The minimum interval between LSAs
Init LSA Gen Delay	The initial LSA generation delay

Table 15: Show OSPF Status Output Fields (Continued)

Label	Description
`Sec LSA Gen Delay`	The delay between the generation of the first and second LSAs
`Max LSA Gen Delay`	The maximum interval between two consecutive LSAs
`Last Ext SPF Run`	The time that the last external SPF calculation was run
`Ext LSA Cksum Sum`	The 32-bit unsigned sum of the LS checksums of the external LSAs contained in this area's link-state database
`OSPF Last Enabled`	The time that OSPF was last enabled on the interface
`Export Policies`	Indicates whether any routing policies have been applied to the OSPF interface
`OSPF Ldp Sync Admin Status`	Indicates whether the IGP-LDP synchronization feature is enabled or disabled on all interfaces participating in the OSPF routing protocol

virtual-link

Syntax **virtual-link** [**detail**]

Context show>router>ospf

Description This command displays information for OSPF virtual links.

Parameters **detail —** provides operational and statistical information about virtual links associated with this router

Output The following output is an example of OSPF virtual link information, and Table 16 describes the fields.

Sample Output

```
A:ALU-A# show router ospf virtual-link
=================================================================
OSPF Virtual Links
=================================================================
Nbr Rtr Id      Area Id         Local Interface     Metric State
-----------------------------------------------------------------
180.0.0.10      0.0.0.1         180.1.7.12          300    PToP
180.0.0.10      0.0.0.2         180.2.7.12          300    PToP
-----------------------------------------------------------------
No. of OSPF Virtual Links: 2
=================================================================
A:ALU-A#
```

```
A:ALU-A# show router ospf virtual-link detail
===============================================================================
OSPF Virtual Links (detailed)
===============================================================================
Neighbor Router Id :  180.0.0.10
-------------------------------------------------------------------------------
Nbr Router Id  : 180.0.0.10             Area Id        : 0.0.0.1
Local Interface: 180.1.7.12             Metric         : 300
State          : Point To Point         Admin State    : Up
Hello Intrvl   : 10 sec                 Rtr Dead Intrvl: 60 sec
Tot Rx Packets : 43022                  Tot Tx Packets : 42964
Rx Hellos      : 24834                  Tx Hellos      : 24853
Rx DBDs        : 3                      Tx DBDs        : 2
Rx LSRs        : 0                      Tx LSRs        : 0
Rx LSUs        : 15966                  Tx LSUs        : 16352
Rx LS Acks     : 2219                   Tx LS Acks     : 1757
Retransmits    : 0                      Discards       : 0
Bad Networks   : 0                      Bad Versions   : 0
Bad Areas      : 0                      Bad Dest Addrs : 0
Bad Auth Types : 0                      Auth Failures  : 0
Bad Neighbors  : 0                      Bad Pkt Types  : 0
Bad Lengths    : 0                      Bad Hello Int. : 0
Bad Dead Int.  : 0                      Bad Options    : 0
Retrans Intrvl : 5 sec                  Transit Delay  : 1 sec
Last Event     : 11/07/2008 17:11:56    Authentication : None
-------------------------------------------------------------------------------
Neighbor Router Id : 180.0.0.10
-------------------------------------------------------------------------------
Nbr Router Id  : 180.0.0.10             Area Id        : 0.0.0.2
Local Interface: 180.2.7.12             Metric         : 300
State          : Point To Point         Admin State    : Up
Hello Intrvl   : 10 sec                 Rtr Dead Intrvl: 60 sec
Tot Rx Packets : 43073                  Tot Tx Packets : 43034
Rx Hellos      : 24851                  Tx Hellos      : 24844
Rx DBDs        : 3                      Tx DBDs        : 2
Rx LSRs        : 1                      Tx LSRs        : 1
Rx LSUs        : 18071                  Tx LSUs        : 17853
Rx LS Acks     : 147                    Tx LS Acks     : 334
Retransmits    : 0                      Discards       : 0
Bad Networks   : 0                      Bad Versions   : 0
Bad Areas      : 0                      Bad Dest Addrs : 0
Bad Auth Types : 0                      Auth Failures  : 0
Bad Neighbors  : 0                      Bad Pkt Types  : 0
Bad Lengths    : 0                      Bad Hello Int. : 0
Bad Dead Int.  : 0                      Bad Options    : 0
Retrans Intrvl : 5 sec                  Transit Delay  : 1 sec
Last Event     : 11/07/2008 17:12:00    Authentication : None
===============================================================================
A:ALU-A#
```

Table 16: Show Virtual Link Output Fields

Label	Description
`Nbr Rtr ID`	The router ID(s) of neighboring routers
`Area Id`	A 32-bit integer that identifies an area
`Local Interface`	The IP address of the local egress interface used to maintain the adjacency to reach this virtual neighbor
`Metric`	The metric value associated with the route. This value is used when importing this static route into other protocols. When the metric is configured as 0, then the metric configured in OSPF, default-metric, applies. This value is also used to determine which static route to install in the forwarding table.
`State`	The operational state of the virtual link to the neighboring router
`Authentication`	Specifies whether authentication is enabled for the interface or virtual link
`Hello Intrvl`	The length of time, in seconds, between the Hello packets that the router sends on the interface
`Rtr Dead Intrvl`	The total number of OSPF packets received where the dead interval given in the packet was not equal to that configured on this interface since the OSPF admin status was enabled
`Tot Rx Packets`	The total number of OSPF packets received on this interface since the OSPF admin status was enabled
`Rx Hellos`	The total number of OSPF Hello packets received on this interface since the OSPF admin status was enabled
`Rx DBDs`	The total number of OSPF database description packets received on this interface since the OSPF admin status was enabled
`Rx LSRs`	The total number of Link-State Requests (LSRs) received on this interface since the OSPF admin status was enabled
`Rx LSUs`	The total number of Link-State Updates (LSUs) received on this interface since the OSPF admin status was enabled
`Rx LS Acks`	The total number of Link-State Acknowledgements received on this interface since the OSPF admin status was enabled
`Tot Tx Packets`	The total number of OSPF packets transmitted on this interface since the OSPF admin status was enabled
`Tx Hellos`	The total number of OSPF Hello packets transmitted on this interface since the OSPF admin status was enabled
`Tx DBDs`	The total number of OSPF database description packets transmitted on this interface since the OSPF admin status was enabled

Table 16: Show Virtual Link Output Fields (Continued)

Label	Description
Tx LSRs	The total number of OSPF Link-State Requests (LSRs) transmitted on this interface since the OSPF admin status was enabled
Tx LSUs	The total number of OSPF Hello packets transmitted on this interface since the OSPF admin status was enabled
Tx LS Acks	The total number of OSPF Link-State Acknowledgements transmitted on this interface since the OSPF admin status was enabled
Retransmits	The total number of OSPF retransmits sent on this interface since the OSPF admin status was last enabled
Discards	The total number of OSPF packets discarded on this interface since the OSPF admin status was last enabled
Bad Networks	The total number of OSPF packets received with invalid network or mask since the OSPF admin status was last enabled
Bad Versions	The total number of OSPF packets received with bad OSPF version numbers since the OSPF admin status was last enabled
Bad Areas	The total number of OSPF packets received with an area mismatch since the OSPF admin status was last enabled
Bad Dest Addrs	The total number of OSPF packets received with the incorrect IP destination address since the OSPF admin status was last enabled
Bad Auth Types	The total number of OSPF packets received with an invalid authorization type since the OSPF admin status was last enabled
Auth Failures	The total number of OSPF packets received with an invalid authorization key since the OSPF admin status was last enabled
Bad Neighbors	The total number of OSPF packets received where the neighbor information does not match the information this router has for the neighbor since the OSPF admin status was last enabled
Bad Pkt Types	The total number of OSPF packets received with an invalid OSPF packet type since the OSPF admin status was last enabled
Bad Lengths	The total number of OSPF packets received on this interface with a total length not equal to the length given in the packet itself since the OSPF admin status was last enabled
Bad Hello Int.	The total number of OSPF packets received where the hello interval given in the packet was not equal to that configured on this interface since the OSPF admin status was last enabled

Table 16: Show Virtual Link Output Fields (Continued)

Label	Description
`Bad Dead Int.`	The total number of OSPF packets received where the dead interval given in the packet was not equal to that configured on this interface since the OSPF admin status was last enabled
`Bad Options`	The total number of OSPF packets received with an option that does not match those configured for this interface or area since the OSPF admin status was last enabled
`Retrans Intrvl`	The length of time, in seconds, that OSPF waits before retransmitting an unacknowledged link-state advertisement (LSA) to an OSPF neighbor
`Transit Delay`	The time, in seconds, that it takes to transmit a link-state advertisement (LSA) on the interface or virtual link
`Last Event`	The date and time that an event was last associated with this OSPF interface

virtual-neighbor

Syntax **virtual-neighbor** [**remote** *ip-address*] [**detail**]

Context show>router>ospf

Description This command displays virtual neighbor information.

The **detail** option produces a large amount of data. It is recommended that this option be used only when requesting information on a specific neighbor.

Parameters *ip-address* — displays the specified router. This reduces the amount of output displayed.

detail — displays detailed information on the virtual neighbor

Output The following output is an example of OSPF virtual neighbor information, and Table 17 describes the fields.

Sample Output

```
A:ALU-A# show router ospf virtual-neighbor
===============================================================================
OSPF Virtual Neighbors
===============================================================================
Nbr IP Addr     Nbr Rtr Id      Nbr State Transit Area    RetxQ Len  Dead Time
-------------------------------------------------------------------------------
180.1.6.10      180.0.0.10      Full      0.0.0.1         0          58
180.2.9.10      180.0.0.10      Full      0.0.0.2         0          52
-------------------------------------------------------------------------------
No. of Neighbors: 2
===============================================================================
A:ALU-A#
```

```
A:ALU-A# show router ospf virtual-neighbor detail
===============================================================================
OSPF Virtual Neighbors
===============================================================================
Virtual Neighbor Router Id : 180.0.0.10
-------------------------------------------------------------------------------
Neighbor IP Addr : 180.1.6.10          Neighbor Rtr Id  : 180.0.0.10
Neighbor State   : Full                Transit Area     : 0.0.0.1
Retrans Q Length : 0                   Options          : -E--
Events           : 4                   Last Event Time  : 11/07/2008 17:11:56
Up Time          : 2d 17:47:17         Time Before Dead : 57 sec
Bad Nbr States   : 1                   LSA Inst fails   : 0
Bad Seq Nums     : 0                   Bad MTUs         : 0
Bad Packets      : 0                   LSA not in LSDB  : 0
Option Mismatches: 0                   Nbr Duplicates   : 0
-------------------------------------------------------------------------------
Virtual Neighbor Router Id : 180.0.0.10
-------------------------------------------------------------------------------
Neighbor IP Addr : 180.2.9.10          Neighbor Rtr Id  : 180.0.0.10
Neighbor State   : Full                Transit Area     : 0.0.0.2
Retrans Q Length : 0                   Options          : -E--
Events           : 4                   Last Event Time  : 11/07/2008 17:11:59
Up Time          : 2d 17:47:14         Time Before Dead : 59 sec
Bad Nbr States   : 1                   LSA Inst fails   : 0
Bad Seq Nums     : 0                   Bad MTUs         : 0
Bad Packets      : 0                   LSA not in LSDB  : 0
Option Mismatches: 0                   Nbr Duplicates   : 0
===============================================================================
A:ALU-A#
```

Table 17: Show Virtual Neighbor Output Fields

Label	Description
`Nbr IP Addr`	The IP address this neighbor is using in its IP source address. Note that, on links with no address, this will not be 0.0.0.0, but the address of another of the neighbor's interfaces.
`Nbr Rtr ID`	The router ID(s) of neighboring routers
`Transit Area`	The transit area ID that links the backbone area with the area that has no physical connection with the backbone
`RetxQ Len/ Retrans Q Length`	The current length of the retransmission queue
`No. of Neighbors`	The total number of OSPF neighbors adjacent on this interface, in a state of INIT or greater, since the OSPF admin status was enabled
`Nbr State`	The operational state of the virtual link to the neighboring router
`Options`	The total number of OSPF packets received with an option that does not match those configured for this virtual interface or transit area since the OSPF admin status was enabled
`Events`	The total number of events that have occurred since the OSPF admin status was enabled
`Last Event Time`	The date and time that an event was last associated with this OSPF interface
`Up Time`	The uninterrupted time, in hundredths of seconds, that the adjacency to this neighbor has been up
`Dead Time/Time Before Dead`	The amount of time, in seconds, until the dead router interval expires
`Bad Nbr States`	The total number of OSPF packets received where the neighbor information does not match the information this router has for the neighbor since the OSPF admin status was last enabled
`LSA Inst fails`	The total number of times an LSA could not be installed into the LSDB due to a resource allocation issue since the OSPF admin status was last enabled
`Bad Seq Nums`	The total number of times that a database description packet was received with a sequence number mismatch since the OSPF admin status was last enabled
`Bad MTUs`	The total number of times that the MTU in a received database description packet was larger than the MTU of the receiving interface since the OSPF admin status was enabled

Table 17: Show Virtual Neighbor Output Fields (Continued)

Label	Description
Bad Packets	The total number of times that an LS update was received with an illegal LS type or an option mismatch since the OSPF admin status was enabled
LSA not in LSDB	The total number of times that an LS request was received for an LSA not installed in the LSDB of this router since the OSPF admin status was enabled
Option Mismatches	The total number of times that an LS update was received with an option mismatch since the OSPF admin status was enabled
Nbr Duplicates	The total number of times that a duplicate database description packet was received during the Exchange state since the OSPF admin status was enabled

Clear Commands

ospf

Syntax **ospf**

Context clear>router

Description This command clears and resets OSPF protocol entities.

database

Syntax **database** [**purge**]

Context clear>router>ospf

Description This command clears all LSAs received from other nodes and refreshes all self-originated LSAs.

purge — clears all self-originated LSAs and reoriginates all self-originated LSAs

export

Syntax **export**

Context clear>router>ospf

Description This command re-evaluates all effective export route policies.

neighbor

Syntax **neighbor** [*ip-int-name* | *ip-address*]

Context clear>router>ospf

Description This command marks the neighbor as dead and reinitiates the affected adjacencies.

Parameters *ip-int-name* — clears all neighbors for the interface specified by this interface name

ip-address — clears all neighbors for the interface specified by this IP address

statistics

Syntax **statistics**

Context clear>router>ospf

Description This command clears all neighbor, router, interface, SPF, and global statistics for OSPF.

Debug Commands

ospf

Syntax **ospf**

Context debug>router

Description This command enables the context for OSPF debugging purposes.

area

Syntax **area** [*area-id*]
no area

Context debug>router>ospf

Description This command enables or disables debugging for an OSPF area.

Parameters *area-id* — the OSPF area ID expressed in dotted-decimal notation or as a 32-bit decimal integer

area-range

Syntax **area-range** [*ip-address*]
no area-range

Context debug>router>ospf

Description This command enables or disables debugging for an OSPF area range.

Parameters *ip-address* — the IP address for the range used by the ABR to advertise into another area

cspf

Syntax **cspf** [*ip-address*]
no cspf

Context debug>router>ospf

Description This command enables or disables debugging for an OSPF constraint-based shortest path first (CSPF).

Parameters *ip-address* — the IP address for the range used for CSPF

interface

Syntax **interface** [*ip-int-name* | *ip-address*]
no interface

Context debug>router>ospf

Description This command enables or disables debugging for an OSPF interface.

Parameters *ip-int-name* — the IP interface name. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

ip-address — the interface's IP address

leak

Syntax **leak** [*ip-address*]
no leak

Context debug>router>ospf

Description This command enables or disables debugging for OSPF leaks.

Parameters *ip-address* — the IP address to debug OSPF leaks

lsdb

Syntax **lsdb** [**type**] [*ls-id*] [*adv-rtr-id*] [**area** *area-id*]
no lsdb

Context debug>router>ospf

Description This command enables or disables debugging for an OSPF link-state database.

Parameters **type —** the OSPF link-state database type

Values router, network, summary, asbr, extern, nssa

ls-id — an LSA type-specific field containing either a router ID or an IP address. It identifies the piece of the routing domain being described by the advertisement.

adv-rtr-id — the router identifier of the router advertising the LSA

area-id — the OSPF area ID expressed in dotted-decimal notation or as a 32-bit decimal integer

misc

Syntax [**no**] **misc**

Context debug>router>ospf

Description This command enables or disables debugging for miscellaneous OSPF events.

neighbor

Syntax **neighbor** [*ip-int-name* | *router-id*]
no neighbor

Context debug>router>ospf

Description This command enables or disables debugging for an OSPF neighbor.

Parameters *ip-int-name* — the neighbor interface name

router-id — neighbor information for the neighbor identified by the the specified router ID

nssa-range

Syntax **nssa-range** [*ip-address*]
no nssa-range

Context debug>router>ospf

Description This command enables or disables debugging for an NSSA range.

Parameters *ip-address* — the IP address range to debug

packet

Syntax **packet** [*packet-type*] [*ip-address*]
no packet

Context debug>router>ospf

Description This command enables or disables debugging for OSPF packets.

Parameters *packet-type* — the OSPF packet type to debug

Values hello, dbdescr, lsrequest, lsupdate, lsack

ip-address — the IP address to debug

rtm

Syntax **rtm** [*ip-address*]
no rtm

Context debug>router>ospf

Description This command enables or disables debugging for the OSPF routing table manager.

Parameters *ip-address* — the IP address to debug

spf

Syntax **spf** [*type*] [*dest-addr*]
no spf

Context debug>router>ospf

Description This command enables or disables debugging for OSPF SPF. Information regarding overall SPF start and stop times are shown. To see detailed information regarding the SPF calculation of a given route, the route must be specified as an optional argument.

Parameters *type* — the area to debug

Values intra-area, inter-area, external

dest-addr — the destination IP address to debug

virtual-neighbor

Syntax **virtual-neighbor** [*ip-address*]
no virtual-neighbor

Context debug>router>ospf

Description This command enables or disables debugging for an OSPF virtual neighbor.

Parameters *ip-address* — the IP address of the virtual neighbor

IS-IS

In This Chapter

This chapter provides information about configuring the Intermediate System-to-Intermediate System (IS-IS) protocol.

Topics in this chapter include:

Overview of IS-IS

IS-IS is an interior gateway protocol (IGP), similar to OSPF, that is used within large autonomous systems (ASs). IS-IS is a link-state protocol. Each IS-IS router maintains an identical database (called the link-state database, topological database, or routing information database [RIB]) of the AS, including information about the local state of each router (for example, its usable interfaces and reachable neighbors).

IS-IS-TE (IS-IS with traffic engineering extensions) is used to advertise reachability information and traffic engineering information such as available bandwidth.

Entities within IS-IS include networks, intermediate systems, and end systems. In IS-IS, a network is an autonomous system (AS), or routing domain, with intermediate systems and end systems. A router, such as the 7705 SAR, is an intermediate system. Intermediate systems send, receive, and forward protocol data units (PDUs). End systems are network devices (or hosts) that send and receive PDUs but do not forward them.

Intermediate system and end system protocols allow routers and nodes to identify each other. IS-IS sends out link-state updates (called link-state PDUs, or LSPs) periodically throughout the network so that each router can maintain current network topology information.

IS-IS uses a cost metric that represents the status of a link (and optionally, the bandwidth of the interface) in an algorithm to determine the best route to a destination. This algorithm is called the Shortest Path First (SPF), or Dijkstra, algorithm. Routing decisions are made using the link-state information. IS-IS evaluates topology changes and, if necessary, performs SPF recalculations.

When the best route to a particular destination is determined, the route information is sent to the routing table manager (RTM). The RTM may contain more than one best route to a destination from multiple protocols. Because metrics from different protocols are not comparable, the RTM uses preference to select the best route. The route with the lowest preference value is selected.

The best routes from the RTM are then added to the forwarding table (also known as the forwarding information base [or FIB]). All forwarding decisions are based on the information in the forwarding database.

The forwarding (or dropping) of packets is controlled by filters applied to the interface and route policies applied to the IS-IS protocol. Refer to the 7705 SAR OS Router Configuration Guide for information on filters and route policies.

In Release 3.0, the following major IS-IS features are supported:

- IS-IS areas (two-level hierarchy)
- ISO network addressing
- neighbors and adjacencies
- metrics
- authentication
- route redistribution and summarization
- IS-IS traffic engineering (TE) extensions (to track and advertise available bandwidth, link administration groups (or link colors), SRLGs, and TE metrics – used by MPLS traffic engineering; that is, RSVP-TE)

IS-IS Areas (Two-level Hierarchy)

IS-IS can subdivide an autonomous system into areas to simplify the calculation of routes and minimize the size of IP routing tables. When an AS is divided into areas, each IS-IS router in an area must maintain an identical link-state database of the area topology, but routes from other areas can be summarized. Sometimes one ‘default’ route can be used to represent many different routes. The topology is hidden from routing devices in other areas, which minimizes the size of the link-state database and reduces IS-IS link-state PDUs.

IS-IS uses a two-level hierarchy when dividing an AS into smaller areas. A system logically belongs to one area. Level 1 routing is performed within an area. Level 2 routing is performed between areas. A 7705 SAR can be configured as a level 1 router, level 2 router, or level 1/2 router.

Level 1 routers know the topology in their area, including all routers and end systems in their area, but do not know the identity of routers or destinations outside of their area. Level 1 routers forward traffic with destinations outside of their area to a level 1/2 router in their area.

Sometimes, the shortest path to an outside destination is not through the closest level 1/2 router, or the only level 1/2 router to forward packets out of an area is not operational. In order to avoid such cases, route leaking provides a mechanism to leak level 2 routes to level 1 routers to provide routing information regarding inter-area routes. Therefore, a level 1 router has more options to forward packets. For more information about route leaking, see Configuring Leaking on page 165.

Level 2 routers know the level 2 topology and know which addresses are reachable by each level 2 router. Level 2 routers do not need to know the topology within any level 1 area, except if the level 2 router is also a level 1 router within a single area. By default, only level 2 routers can exchange PDUs or routing information directly with external routers located outside the routing domain.

Table 18 describes the router types (or intermediate systems) within IS-IS.

Table 18: IS-IS Intermediate Systems

Intermediate System	Description
Level 1	Maintains a link-state database of other routers that reside in the same area (i.e. local area) Exchanges topology information for the local area Routing is performed within the area, based on the area ID portion of the ISO address (see ISO Network Addressing) If the destination address is in the area (area ID is equal), routers forward the packets to the level 1 router that is advertising the destination address, based on the system ID If the destination address is not in the area (area ID is not equal), routers forward the packet to the nearest level 1/2 router in the local area
Level 2	Resides within an area but connects to other level 2 routers in multiple areas in a backbone mesh Maintains a link-state database of other level 2 routers and of the level 1/2 routers in each local area Exchanges topology information between areas Routing is performed between areas based on the area address
Level 1/2	Acts as an area border router with links to the level 2 backbone as well as to the level 1 routers within its area Maintains two link-state databases – a level 1 link-state database of the routers in the local area and a level 2 link-state database of the backbone and any level 1/2 routers Exchanges topology information within the local area and between areas Routing is performed within and between areas If the destination address is in the area (area ID is equal), routers use the level 1 database to forward the packets to the level 1 router that is advertising the destination address, based on the system ID If the destination address is not in the area (area ID is not equal), routers use the level 2 database to forward the packet based on the area ID

Figure 3 shows an example of an IS-IS topology.

Figure 3: IS-IS Topology

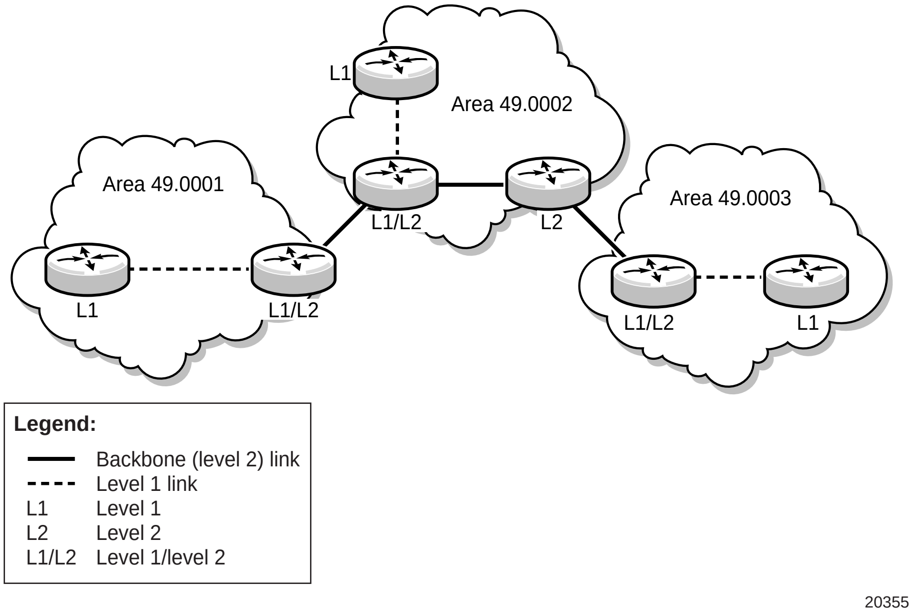

ISO Network Addressing

IS-IS uses ISO network addresses. There are two types of network addresses:

- Network Service Access Point (NSAP)

  NSAP addresses identify a point of connection to the network, such as a router interface. Each NSAP represents a service that is available at that node. An end system can have multiple NSAP addresses; the addresses differ only by the last byte (called the *n*-selector). In addition to having multiple services, a single node can belong to multiple areas.

- Network Entity Title (NET)

  NET addresses identify network layer entities or processes instead of services. Structurally, an NET is identical to an NSAP address but has an *n*-selector of 00. Most end systems have one NET. Intermediate systems (routers) can have up to three NETs, differentiated by the area ID.

NSAP addresses are divided into three parts. Only the area ID portion is configurable:

- area ID – a variable-length field between 1 and 13 bytes that identifies the area to which the router belongs. This field includes the Authority and Format Identifier (AFI) as the first (most significant byte) and the area identifier.
- system ID – A 6-byte system identifier. This value is not configurable. The system ID is derived from the system or router ID and uniquely identifies the router.
- selector ID – A 1-byte selector identifier that is always 00 for an NET. This value is not configurable.

The area ID portion of the NET can be manually configured with 1 to 13 bytes. If fewer than 13 bytes are entered, the rest of the field is padded with zeros.

Neighbors and Adjacencies

IS-IS routers discover their neighbors by exchanging Hello PDUs. Neighbors are routers that have an interface to a common network/area. In a broadcast-supported topology, one router sends Hello packets to a multicast address and receives Hello packets in return. Unicast Hello packets are used in non-broadcast topologies.

Because all routing devices on a common network must agree on certain parameters, these parameters are included in Hello packets. Differences in these parameters can prevent neighbor relationships from forming.

A level 1 router will not become a neighbor with a node that does not have a common area address. However, if a level 1 router has area addresses A, B, and C, and a neighbor has area addresses B and D, the level 1 router will accept the other node as a neighbor because address B is common to both routers.

When Hello packets have been successfully exchanged, the neighbors are considered to be adjacent.

Within an area, level 1 routers exchange LSPs that identify the IP addresses reachable by each router. Each router has one LSP that contains information about that router; zero or more IP addresses, subnet masks, and metric combinations can be included in each LSP. Each level 1 router is manually configured with the IP address, subnet mask, and metric combinations that are reachable on each interface.

Level 2 routers exchange LSPs that include a complete list of IP addresses, subnet masks, and metrics specifying all the IP addresses that are reachable in their area. Level 2 routers can also report external reachability information, corresponding to addresses reachable by routers in other routing domains or autonomous systems.

Routers with common area addresses form level 1 adjacencies. Routers with no common NET addresses form level 2 adjacencies, if they are capable. See Figure 4.

Figure 4: Using Area Addresses to Form Adjacencies

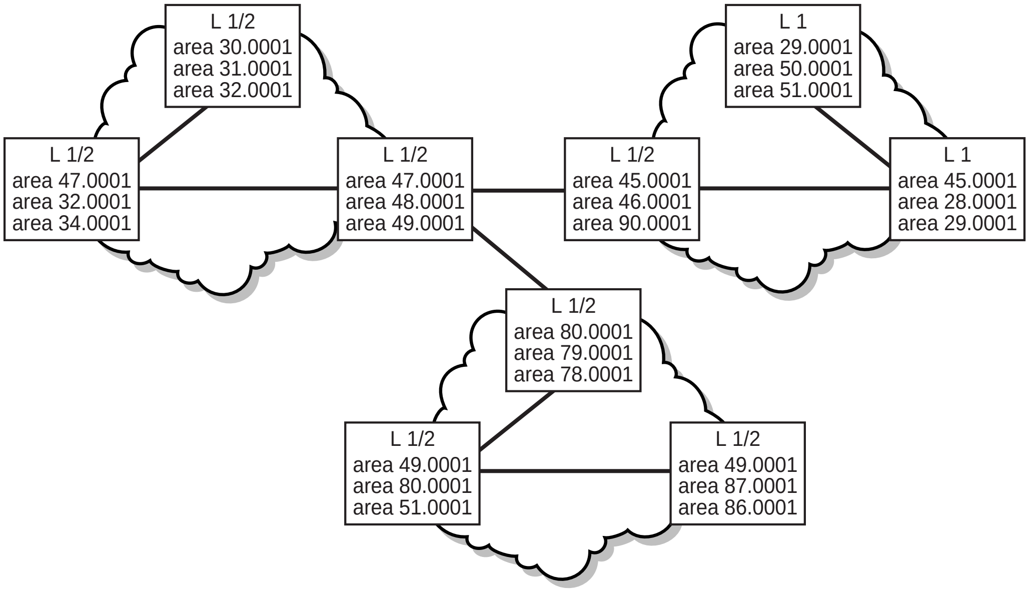

20108

Level 2 adjacencies are formed with other level 2 nodes whose area addresses do not overlap. If the area addresses do not overlap, the link is considered by both routers to be level 2 and only level 2 LSPs flow on the link.

Designated Routers

In multi-access broadcast networks, such as Ethernet networks, with at least two attached routers, a designated router can be elected. The IS-IS protocol refers to the designated router as the designated intermediate system (DIS).

The concept of a designated router was developed in order to avoid the formation of adjacencies between all attached routers. Without a designated router, the area would be flooded with link-state PDUs (LSPs) – a router would send LSPs to all its adjacent neighbors, and each in turn would send LSPs to all their neighbors, and so on. This would create multiple copies of the same LSP on the same link.

The designated router reduces the number of adjacencies required because each router forms an adjacency only with the designated router. Only the designated router sends LSPs in multicast format to the rest of the network, reducing the amount of routing protocol traffic.

In IS-IS, a broadcast subnetwork with multiple connected routers is considered to be a pseudonode. The pseudonode has links to each of the routers and each of the routers has a single link to the pseudonode (rather than links to each of the other routers). LSPs are generated on behalf of the pseudonode by the DIS.

The DIS has two tasks:

- create and update the pseudonode LSP
- flood the LSP over the LAN

The DIS is automatically elected based on the interface priority of the router and/or if it has the highest MAC address of all routers in the LAN. If all interface priorities are the same, the router with the highest subnetwork point of attachment (SNPA) is selected. The SNPA is the MAC address on a LAN.

Every IS-IS router interface is assigned both a level 1 priority and a level 2 priority. If a new router starts up in the LAN and has a higher interface priority, the new router preempts the original DIS and becomes the new DIS. The new DIS purges the old pseudonode LSP and floods a new set of LSPs.

Because different priorities can be set according to level 1 or level 2 routing, there can be two different routers in an Ethernet LAN that are DIS-designated. One DIS supports all level 1 routers, and the other DIS supports all level 2 routers on that segment.

The DIS generates the pseudonode LSP. The DIS reports all LAN neighbors (including itself) in the pseudonode LSP. All LAN routers communicate with the pseudonode via their LSPs. The pseudonode reduces the number of adjacencies by having all physical devices exchange information only with the pseudonode. Each router listens for updates to the pseudonode and updates its individual topology according to those updates.

Notes:

- In point-to-point networks, where a single pair of routers is connected, no designated router is elected. An adjacency must be formed with the neighbor router.
- To significantly improve adjacency forming and network convergence, a network should be configured as point-to-point if only two routers are connected, even if the network is a broadcast media such as Ethernet.

IS-IS Packet Types

Table 19 describes the packet types used by IS-IS to exchange protocol information.

Table 19: IS-IS Packet Types

Packet Type	Description
Hello PDUs	Routers with IS-IS enabled send Hello PDUs to IS-IS-enabled interfaces to discover neighbors and establish adjacencies.
Link-state PDUs (LSPs)	LSPs contain information about the state of adjacencies to neighboring IS-IS systems and are used to build the link-state database. LSPs are flooded periodically throughout an area. Level 1 and level 2 LSPs are supported.
Complete sequence number PDUs (CSNPs)	In order for all routers to maintain the same information (synchronize), CSNPs inform other routers that some LSPs might be outdated or missing from their database. CSNPs contain a complete list of all LSPs in the current IS-IS database. Level 1 and level 2 CSNPs are supported.
Partial sequence number PDUs (PSNPs)	PSNPs are used to request missing LSPs and acknowledge that an LSP was received. Level 1 and level 2 PSNPs are supported.

IS-IS sends only the changed information, not the whole topology information or whole link-state database, when a change takes place. From the topological database, each router constructs a tree of shortest paths with itself as root (that is, runs the Dijkstra algorithm). IS-IS distributes routing information between routers belonging to a single AS.

To summarize:

- Hello PDUs are sent over the IS-IS-enabled interfaces to discover neighbors and establish adjacencies.
- IS-IS neighbor relationships are formed if the Hello PDUs contain information that meets the criteria for forming an adjacency.
- Routers can build a link-state PDU based upon their local interfaces that are configured for IS-IS and prefixes learned from other adjacent routers.
- Routers flood LSPs to the adjacent neighbors except the neighbor from which they received the same LSP. The link-state database is constructed from these LSPs.
- A Shortest Path Tree (SPT) is calculated by each router, and from this SPT the routing table is built.

Metrics

IS-IS uses a cost metric that represents the status of a link in an algorithm to determine the best route to a destination. This algorithm is called the Shortest Path First (SPF), or Dijkstra, algorithm. Routing decisions are made using the link-state information. IS-IS evaluates topology changes and, if necessary, performs SPF recalculations.

To calculate the lowest cost to reach a destination, each configured level on each interface must have a cost. The costs for each level on an interface may be different.

In IS-IS, if the metric is not configured, the default cost 10 is used, regardless of the actual capacity of the link. By default, IS-IS does not use reference bandwidth in the calculation, unlike OSPF.

Authentication

All IS-IS protocol exchanges can be authenticated. This guarantees that only trusted routers can participate in autonomous system routing. The Alcatel-Lucent implementation of IS-IS supports plain text (simple password) and Message Digest 5 (MD5) authentication.

When authentication is enabled on a link, a text string password must be configured. Neighbor IS-IS routers must supply the password in all IS-IS packets they send to an interface.

Plain text authentication includes the password in each IS-IS packet sent on a link.

MD5 authentication is more secure than plain text authentication. MD5 authentication uses the password as an encryption key. Routers in the same routing domain must be configured with the same key. When the MD5 hashing algorithm is used for authentication, MD5 is used to verify data integrity by creating a 128-bit message digest from the data input that is included in each packet. The packet is transmitted to the router neighbor and can only be decrypted if the neighbor has the correct password.

By default, authentication is not enabled on an interface.

Route Redistribution and Summarization

Route redistribution is the taking of routes from one protocol and sending them to another protocol. In Release 3.0, the 7705 SAR supports the redistribution of static routes into OSPF and IS-IS and the redistribution of routes between IS-IS levels. The routes can be redistributed as level 1, level 2, or level 1/2 routes, depending on the level capability of the IS-IS router.

Route redistribution involves the use of routing policies. For information on routing policies, refer to the 7705 SAR OS Router Configuration Guide, “Route Policies”.

IS-IS route summarization allows users to create aggregate IPv4 addresses that include multiple groups of IPv4 addresses for a given IS-IS level. Routes redistributed from other routing protocols can also be summarized.

IS-IS route summarization helps to reduce the size of the link-state database and the routing table. It also helps to reduce the chance of route flapping, when a router alternately advertises a destination network via one route then another (or advertises a route as unavailable then available again) in quick sequence.

IS-IS-TE Extensions

IS-IS traffic engineering (TE) extensions enable the 7705 SAR to include traffic engineering information in the algorithm in order to calculate the best route to a destination. The traffic information includes:

- maximum reservable bandwidth
- unreserved bandwidth
- available bandwidth
- link administration groups (or link colors)
- SRLGs
- TE metrics

Bidirectional Forwarding Detection (BFD) for IS-IS

BFD is a simple protocol for detecting failures in a network. BFD uses a “hello” mechanism that sends control messages periodically to the far end and receives periodic control messages from the far end. BFD can detect device, link, and protocol failures.

When BFD is enabled on an IS-IS interface, the state of the interface is tied to the state of the BFD session between the local node and remote (far-end) node. BFD is implemented in asynchronous mode only (similar to a heartbeat message), meaning that neither end responds to control messages; rather, the messages are sent in the interval configured at each end.

If the configured number of consecutive BFD missed messages is reached, the link is declared down and IS-IS takes the appropriate action (for example, generates an LSP against the failed link or reroutes around the failed link).

Due to the lightweight nature of BFD, frequency of BFD packets can be relatively high (up to 10 per second), hence it can detect failures faster than other detection protocols, making it ideal for use in applications such as mobile transport.

IS-IS Configuration Process Overview

Figure 5 shows the process to provision basic IS-IS parameters.

Figure 5: IS-IS Configuration Process

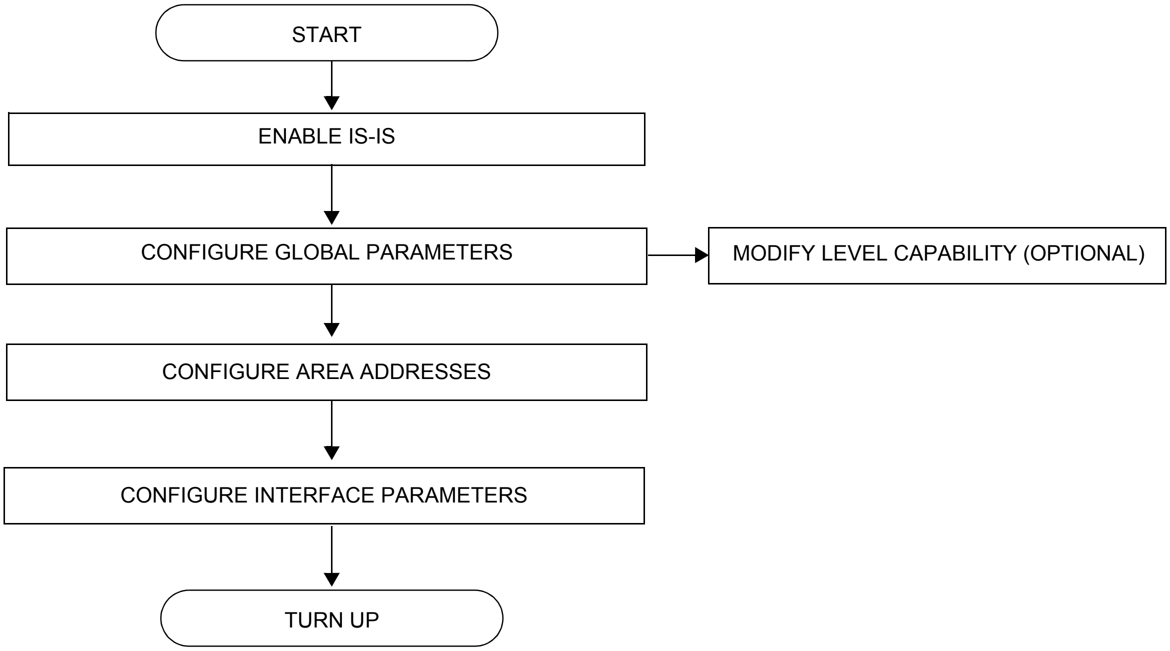
START
ENABLE IS-IS
CONFIGURE GLOBAL PARAMETERS
MODIFY LEVEL CAPABILITY (OPTIONAL)
CONFIGURE AREA ADDRESSES
CONFIGURE INTERFACE PARAMETERS
TURN UP

Configuration Notes

General

- IS-IS must be enabled on each participating 7705 SAR.
- There are no default network entity titles.
- There are no default interfaces.
- By default, 7705 SAR routers are assigned a level 1/2 capability.

Reference Sources

For information on supported IETF drafts and standards, as well as standard and proprietary MIBs, refer to Standards and Protocol Support on page 347.

Configuring IS-IS with CLI

This section provides information to configure the Intermediate System-to-Intermediate System (IS-IS) protocol using the command line interface.

Topics in this section include:

IS-IS Configuration Overview

For IS-IS to operate on 7705 SAR routers, IS-IS must be explicitly enabled, and at least one area address and interface must be configured. If IS-IS is enabled but no area address or interface is configured, no routes are exchanged. When at least one area address and interface are configured, adjacencies can be formed and routes exchanged.

Router Levels

The router's level capability can be configured globally and on a per-interface basis. The interface level parameters specify the interface's routing level. The neighbor capability and parameters define the adjacencies that are established.

When IS-IS is enabled, the global default level capability is level 1/2, which enables the router to operate as either a level 1 and/or a level 2 router with the associated databases. The router runs separate shortest path first (SPF) calculations for the level 1 area routing and for the level 2 multi-area routing to create the IS-IS routing table.

The level value can be modified on both or either of the global and interface levels to be only level 1-capable, only level 2-capable, or both level 1 and level 2-capable.

If the default value is not modified on any routers in the area, the routers try to form both level 1 and level 2 adjacencies on all IS-IS interfaces. If the default values are modified to level 1 or level 2, the number of adjacencies formed are limited to that level only.

Area Addresses

The `area-id` command specifies the area address portion of the NET, which is used to define the IS-IS area to which the router will belong. At least one area ID must be configured on each router participating in IS-IS. A maximum of three area IDs can be configured per router.

The area address identifies a point of connection to the network, such as a router interface, and is called a network service access point (NSAP). The routers in an area manage routing tables of destinations within the area. The Network Entity Title (NET) value is used to identify the IS-IS area to which the router belongs.

NSAP addresses are divided into three parts. Only the area ID portion is configurable:

- area ID – a variable-length field between 1 and 13 bytes that identifies the area to which the router belongs. This field includes the Authority and Format Identifier (AFI) as the first (most significant byte) and the area identifier.
- system ID – A 6-byte system identifier. This value is not configurable. The system ID is derived from the system or router ID and uniquely identifies the router.
- selector ID – A 1-byte selector identifier that is always 00 for an NET. This value is not configurable.

The area ID portion of the NET can be manually configured with 1 to 13 bytes. If fewer than 13 bytes are entered, the rest of the field is padded with zeros.

Interface Level Capability

The level capability value configured on the interface level is compared to the level capability value configured on the global level to determine the type of adjacencies that can be established. The default value for 7705 SAR routers and interfaces is level 1/2.

Table 20 lists capability combinations and the potential adjacencies that can be formed.

Table 20: Potential Adjacency Capabilities

Global Level	Interface Level	Potential Adjacency
Level 1/2	Level 1/2	Level 1 and/or level 2
Level 1/2	Level 1	Level 1 only
Level 1/2	Level 2	Level 2 only
Level 2	Level 1/2	Level 2 only
Level 2	Level 2	Level 2 only
Level 2	Level 1	None
Level 1	Level 1/2	Level 1 only
Level 1	Level 2	None
Level 1	Level 1	Level 1 only

Route Leaking

An autonomous system running IS-IS can be divided into level 1 areas with a level 2-connected subset (backbone) of the topology that interconnects all of the level 1 areas. Within each level 1 area, the routers exchange link-state information. Level 2 routers also exchange level 2 link-state information to compute routes between areas.

Routers in a level 1 area typically only exchange information within the level 1 area. For IP destinations not found in the prefixes in the level 1 database, the level 1 router forwards PDUs to the nearest level 1/2 router with the attachment bit set in its level 1 link-state PDU.

Routing to the closest level 1/2 router may lead to sub-optimal routing, because the shortest path to the destination is not always through the closest router. To reduce sub-optimal routing, route leaking provides a mechanism to leak (or redistribute) level 2 information into level 1 areas. By distributing more detailed information into the level 1 area, a level 1 router is able to make a better decision as to which level 1/2 router should forward the packet.

The Alcatel-Lucent implementation of IS-IS route leaking is in compliance with RFC 2966, *Domain-wide Prefix Distribution with Two-Level IS-IS*.

Basic IS-IS Configuration

The basic IS-IS configuration tasks that must be performed are:

- enable IS-IS
- modify the level capability on the global level from the default level 1/2 (if required)
- define area addresses
- configure IS-IS interfaces

The following output displays IS-IS default values:

```
ALU-A>config>router>isis# info detail
----------------------------------------------
     level-capability level-1/2
     no authentication-key
     no authentication-type
     authentication-check
     csnp-authentication
     lsp-lifetime 1200
     no export
     hello-authentication
     psnp-authentication
     no traffic-engineering
     no reference-bandwidth
     no disable-ldp-sync
     spf-wait 10 1000 1000
     lsp-wait 5 0 1
     level 1
         no authentication-key
         no authentication-type
         csnp-authentication
         external-preference 160
         hello-authentication
         preference 15
         psnp-authentication
         no wide-metrics-only
     exit
     level 2
         no authentication-key
         no authentication-type
         csnp-authentication
         external-preference 165
         hello-authentication
         preference 18
         psnp-authentication
         no wide-metrics-only
     exit
    interface "system"
         level-capability level-1/2
         csnp-interval 10
         no hello-authentication-key
         no hello-authentication-type
         no interface-type
         lsp-pacing-interval 100
         retransmit-interval 5
         no bfd-enable ipv4
         no mesh-group
         no passive
         level 1
              no hello-authentication-key
              no hello-authentication-type
              hello-interval 9
              hello-multiplier 3
              no metric
              no passive
              priority 64
          exit
          level 2
              no hello-authentication-key
              no hello-authentication-type
              hello-interval 9
              hello-multiplier 3
              no metric
              no passive
              priority 64
          exit
          no shutdown
      exit
      no shutdown
----------------------------------------------
ALU-A>config>router>isis#
```

Configuring IS-IS Components

The following sections show the CLI syntax for:

- Enabling IS-IS
- Configuring the Router Level
- Configuring Area Addresses
- Configuring Global IS-IS Parameters
- Configuring Interface Parameters
- Configuring Leaking
- Redistributing External IS-IS Routers

Enabling IS-IS

IS-IS must be enabled in order for the protocol to be active.

Note: Careful planning is essential when implementing commands that can affect the behavior of global and interface levels.

To configure IS-IS on a router, enter the following command:

CLI Syntax:
```
config
     router <router-name>
        isis
```

Configuring the Router Level

When IS-IS is enabled, the default `level-capability` is level 1/2. This means that the router operates with both level 1 and level 2 routing capabilities. To change the default value in order for the router to operate as a level 1 router or a level 2 router only, you must explicitly modify the `level` value.

Select level 1 to route only within an area. Select level 2 to route to destinations outside an area, toward other eligible level 2 routers.

If the level is modified, the protocol shuts down and restarts, which may affect existing adjacencies and routes.

The `level-capability` value can be configured on the global level and on the interface level. The `level-capability` value determines which level values can be assigned on the router level or on an interface level.

To configure the router level, enter the following command:

CLI Syntax:
```
config>router# isis
          level-capability {level-1 | level-2 | level-1/2}
```

The following example displays a level configuration:

```
A:ALU-A>config>router>isis# info
----------------------------------------------
     level-capability level-1/2
     ----------------------------------------------
A:ALU-A>config>router>isis#
```

Configuring Area Addresses

Use the following syntax to configure an area address. A maximum of three area addresses can be configured.

CLI Syntax:
```
config>router# isis
          area-id area-address
```

The following example displays an area ID configuration:

```
A:ALU-A>config>router>isis# info
----------------------------------------------
     area-id 49.0180.0001
     area-id 49.0180.0002
     area-id 49.0180.0003
----------------------------------------------
A:ALU-A>config>router>isis#
```

Configuring Global IS-IS Parameters

Commands and parameters configured on the global level are inherited by the interface levels. Parameters specified in the interface and interface level configurations override the global configuration.

Use the following syntax to configure global IS-IS parameters:

CLI Syntax:
```
config>router# isis
      level-capability {level-1 | level-2 | level-1/2}
      [no] authentication-check
      authentication-type {password | message-digest}
      authentication-key {authentication-key | hash-key}
         [hash | hash2]
      overload [timeout seconds]
      traffic-engineering
```

The following example displays a global level configuration:

```
A:ALU-A>config>router>isis# info
----------------------------------------------
     level-capability level-2
     area-id 49.0180.0001
     area-id 49.0180.0002
     area-id 49.0180.0003
     authentication-key "H5KBAWrAAQU" hash
     authentication-type password
     overload timeout 90
     traffic-engineering
----------------------------------------------
A:ALU-A>config>router>isis#
```

Configuring Interface Parameters

By default, there are no interfaces associated with IS-IS. An interface belongs to all areas configured on a router. Interfaces cannot belong to separate areas. There are no default interfaces applied to the router’s IS-IS instance. You must configure at least one IS-IS interface in order for IS-IS to work.

To enable IS-IS on an interface, first configure an IP interface in the `config>router>interface` context. Then, apply the interface in the `config>router>isis>interface` context.

You can configure both level 1 parameters and level 2 parameters on an interface. The `level-capability` value determines which level values are used.

Note: For point-to-point interfaces, only the values configured under level 1 are used, regardless of the operational level of the interface.

Use the following syntax to configure interface parameters:

CLI Syntax:

```
config>router# isis
     level {1 | 2}
        [no] wide-metrics-only
     interface ip-int-name
        level-capability {level-1 | level-2 | level-1/2}
        mesh-group [value | blocked]
        interface-type {broadcast | point-to-point}
```

The following example displays a global level and interface configuration:

```
A:ALU-A>config>router>isis# info
----------------------------------------------
     level-capability level-2
     area-id 49.0180.0001
     area-id 49.0180.0002
     area-id 49.0180.0003
     authentication-key "H5KBAWrAAQU" hash
     authentication-type password
     traffic-engineering
     level 1
         wide-metrics-only
     exit
     level 2
         wide-metrics-only
     exit
     interface "system"
     exit
     interface "ALU-1-2"
         level-capability level-2
         mesh-group 85
     exit
     interface "ALU-1-3"
         level-capability level-1
         interface-type point-to-point
         mesh-group 101
     exit
     interface "ALU-1-5"
         level-capability level-1
         interface-type point-to-point
         mesh-group 85
     exit
     interface "to-103"
         mesh-group 101
     exit
----------------------------------------------
A:ALU-A>config>router>isis#
```

Configuring Leaking

IS-IS allows a two-level hierarchy to route PDUs. Level 1 areas can be interconnected by a contiguous level 2 backbone.

The level 1 link-state database contains information only about that area. The level 2 link-state database contains information about the level 2 system and each of the level 1 systems in the area. A level 1/2 router contains information about both level 1 and level 2 databases. A level 1/2 router advertises information about its level 1 area toward the other level 1/2 or level 2 routers.

Packets with destinations outside the level 1 area are forwarded toward the closest level 1/2 router which, in turn, forwards the packets to the destination area.

Sometimes, the shortest path to an outside destination is not through the closest level 1/2 router, or the only level 1/2 router to forward packets out of an area is not operational. Route leaking provides a mechanism to leak level 2 information to level 1 routers to provide routing information regarding inter-area routes. Therefore, a level 1 router has more options to forward packets.

Configure a route policy to leak routes from level 2 into level 1 areas in the `config>router>policy-options>policy-statement` context. For more information on creating route policies, refer to the 7705 SAR OS Router Configuration Guide.

The following example shows the commands to configure prefix list and policy statement parameters in the `config>router` context.

Example:

```
config>router>policy-options# prefix-list loops
..>policy-options>prefix-list# prefix 10.1.1.0/24 longer
..>policy-options>prefix-list# exit
..>policy-options# policy-statement leak
..>policy-options>policy-statement# entry 10
..>policy-options>policy-statement>entry# from
..>policy-options>policy-statement>entry>from# prefix-
   list loops
..>policy-options>policy-statement>entry>from# level 2
..>policy-options>policy-statement>entry>from# exit
..>policy-options>policy-statement>entry# to
..>policy-options>policy-statement>entry>to# level 1
..>policy-options>policy-statement>entry>to# exit
..>policy-options>policy-statement>entry# action accept
..>policy-options>policy-statement>entry>action# exit
..>policy-options>policy-statement>entry# exit
..>policy-options>policy-statement#exit
..>policy-options# commit
..>policy-options#
```

```
A:ALU-A>config>router>policy-options# info
----------------------------------------------
     prefix-list "loops"
          prefix 10.1.1.0/24 longer
     exit
     policy-statement "leak"
          entry 10
               from
                    prefix-list "loop"
                    level 2
               exit
               to
                    level 1
               exit
               action accept
               exit
          exit
     exit
----------------------------------------------
A:ALU-A>config>router>policy-options#
```

Next, apply the policy to leak routes from level 2 into level 1 routers on ALU-A:

CLI Syntax:

```
config>router# isis
     export leak
```

```
A:ALU-A>config>router>isis# info
----------------------------------------------
     area-id 49.0180.0001
     area-id 49.0180.0002
     area-id 49.0180.0003
     authentication-key "//oZrvtvFPn06S42lRIJsE" hash
     authentication-type password
     no authentication-check
     export "leak"
...
----------------------------------------------
A:ALU-A>config>router>isis#
```

After the policy is applied, create a policy to redistribute external IS-IS routes from level 1 routers into the level 2 backbone (see Redistributing External IS-IS Routers). In the `config>router` context, configure the following policy statement parameters:

Example:

```
config>router>policy-options# begin
..>policy-options# policy-statement "isis-ext"
..>policy-options>policy-statement# entry 10
..>policy-options>policy-statement>entry$ from
..>policy-options>policy-statement>entry>from$ external
..>policy-options>policy-statement>entry>from# exit
..>policy-options>policy-statement>entry# to
```

```
..>policy-options>policy-statement>entry>to$ level 2
..>policy-options>policy-statement>entry>to# exit
..>policy-options>policy-statement>entry# action accept
..>policy-options>policy-statement>entry>action# exit
..>policy-options>policy-statement>entry# exit
..>policy-options>policy-statement# exit
..>policy-options# commit
```

Redistributing External IS-IS Routers

By default, IS-IS does not redistribute level 1 external routes into level 2. The policy to redistribute external IS-IS routes must be explicitly applied. Policies are created in the `config>router>policy-options` context. Refer to the 7705 SAR OS Router Configuration Guide for information on creating policies.

The following example displays the policy statement configuration:

```
A:ALU-A>config>router>policy-options# info
-----------------------------------------------
     prefix-list "loops"
          prefix 10.1.1.0/24 longer
     exit
     policy-statement "leak"
          entry 10
               from
                    prefix-list "loop"
                    level 2
               exit
               to
                    level 1
               exit
               action accept
               exit
          exit
     exit
     policy-statement "isis-ext"
          entry 10
               from
                    external
               exit
               to
                    level 2
               exit
               action accept
               exit
          exit
     exit
-----------------------------------------------
A:ALU-A>config>router>policy-options#
```

IS-IS Configuration Management Tasks

This section discusses the following IS-IS configuration management tasks:

- Disabling IS-IS
- Removing IS-IS
- Modifying Global IS-IS Parameters
- Modifying IS-IS Interface Parameters

Disabling IS-IS

The `shutdown` command disables IS-IS on the router. The configuration settings are not changed, reset, or removed.

Use the following CLI syntax to disable IS-IS on a router:

CLI Syntax:
```
config>router# isis
        shutdown
```

Removing IS-IS

The `no isis` command deletes the IS-IS protocol instance. The IS-IS configuration reverts to the default settings.

Use the following CLI syntax to remove IS-IS:

CLI Syntax:
```
config>router#
        no isis
```

Modifying Global IS-IS Parameters

You can modify, disable, or remove global IS-IS parameters without shutting down entities. The changes are applied immediately. Modifying the level capability on the global level causes the IS-IS protocol to restart.

The following example displays an IS-IS global parameter modification.

Example:

```
config>router>isis# overload timeout 500
config>router>isis# level-capability level-1/2
config>router>isis# no authentication-check
config>router>isis# authentication-key raiderslost
```

The following example displays the IS-IS configuration with the modifications entered in the previous example:

```
A:ALU-A>config>router>isis# info
----------------------------------------------
     area-id 49.0180.0001
     area-id 49.0180.0002
     area-id 49.0180.0003
     authentication-key "//oZrvtvFPn06S42lRIJsE" hash
     authentication-type password
     no authentication-check
     overload timeout 500
     level 1
         wide-metrics-only
     exit
     level 2
         wide-metrics-only
     exit
     interface "system"
     exit
     interface "ALU-1-2"
         level-capability level-2
         mesh-group 85
     exit
     interface "ALU-1-3"
         level-capability level-1
         interface-type point-to-point
         mesh-group 101
     exit
     interface "ALU-1-5"
         level-capability level-1
         interface-type point-to-point
         mesh-group 85
     exit
     interface "to-103"
         mesh-group 101
     exit
     interface "A-B"
     exit
     interface "A-C"
     exit
---------------------------------------------------
```

Modifying IS-IS Interface Parameters

You can modify, disable, or remove interface level IS-IS parameters without shutting down entities. Changes take effect immediately. Modifying the level capability on the interface causes the IS-IS protocol on the interface to restart.

To remove an interface, use the `no interface` *`ip-int-name`* command.

To disable an interface, use the `shutdown` command in the interface context.

The following example displays an IS-IS interface parameter modification.

Example:

```
config>router# isis
config>router>isis# interface ALU-1-3
config>router>isis>if# mesh-group 85
config>router>isis>if# passive
config>router>isis>if# lsp-pacing-interval 5000
config>router>isis>if# exit
config>router>isis# interface to-103
config>router>isis>if# hello-authentication-type message-
digest
config>router>isis>if# hello-authentication-key 49ersrule
config>router>isis>if# exit
```

The following example displays the IS-IS configuration with the modifications entered in the previous example:

```
A:ALU-A>config>router>isis# info
----------------------------------------------
     area-id 49.0180.0001
     area-id 49.0180.0002
     area-id 49.0180.0003
     authentication-key "//oZrvtvFPn06S42lRIJsE" hash
     authentication-type password
     no authentication-check
     overload timeout 500
     level 1
          wide-metrics-only
     exit
     level 2
          wide-metrics-only
     exit
     interface "system"
     exit
     interface "ALU-1-2"
          level-capability level-2
          mesh-group 85
     exit
     interface "ALU-1-3"
          level-capability level-1
          interface-type point-to-point
          lsp-pacing-interval 5000
          mesh-group 85
          passive
     exit
     interface "ALU-1-5"
          level-capability level-1
          interface-type point-to-point
          mesh-group 85
     exit
     interface "to-103"
          hello-authentication-key "DvR3l264KQ6vXMTvbAZ1mE" hash
          hello-authentication-type message-digest
          mesh-group 101
     exit
     interface "A-B"
     exit
     ----------------------------------------------
     A:ALU-A>config>router>isis#
```

IS-IS Command Reference

Command Hierarchies

- Configuration Commands
- Show Commands
- Clear Commands
- Debug Commands
- Tools Commands (refer to the Tools chapter in the 7705 SAR OS Services Guide)

Configuration Commands

```
config
    — router
        — [no] isis
            — [no] area-id area-address
            — [no] authentication-check
            — authentication-key {authentication-key | hash-key} [hash | hash2]
            — no authentication-key
            — authentication-type {password | message-digest}
            — no authentication-type
            — [no] csnp-authentication
            — [no] disable-ldp-sync
            — export policy-name [ policy-name...(up to 5 max)]
            — no export
            — [no] hello-authentication
            — [no] interface ip-int-name
                — [no] bfd-enable ipv4
                — csnp-interval seconds
                — no csnp-interval
                — hello-authentication-key [authentication-key | hash-key] [hash | hash2]
                — no hello-authentication-key
                — hello-authentication-type {password | message-digest}
                — no hello-authentication-type
                — interface-type {broadcast | point-to-point}
                — no interface-type
                — level {1 | 2}
                    — hello-authentication-key {authentication-key | hash-key}
                      [hash | hash2]
                    — no hello-authentication-key
                    — hello-authentication-type {password | message-digest}
                    — no hello-authentication-type
                    — hello-interval seconds
                    — no hello-interval
                    — hello-multiplier multiplier
                    — no hello-multiplier
                    — metric metric
                    — no metric
                    — [no] passive
                    — priority number
                    — no priority
                — level-capability {level-1 | level-2 | level-1/2}
                — no level-capability
                — lsp-pacing-interval milliseconds
                — no lsp-pacing-interval
                — mesh-group [value | blocked]
                — no mesh-group
                — [no] passive
                — retransmit-interval seconds
                — no retransmit-interval
                — [no] shutdown
            — level {1 | 2}
                — authentication-key [authentication-key | hash-key] [hash | hash2]
                — no authentication-key
                — authentication-type {password | message-digest}
                — no authentication-type
                — [no] csnp-authentication
                — external-preference external-preference
                — no external-preference
                — [no] hello-authentication
                — preference preference
                — no preference
                — [no] psnp-authentication
                — [no] wide-metrics-only
            — level-capability {level-1 | level-2 | level-1/2}
            — no level-capability
            — lsp-lifetime seconds
            — no lsp-lifetime
            — lsp-wait lsp-wait [lsp-initial-wait [lsp-second-wait]]
            — no lsp-wait
            — overload [timeout seconds]
            — no overload
            — overload-on-boot [timeout seconds]
            — no overload-on-boot
            — [no] psnp-authentication
            — reference-bandwidth bandwidth-in-kbps
            — no reference-bandwidth
            — [no] shutdown
            — spf-wait spf-wait [spf-initial-wait [spf-second-wait]
            — no spf-wait
            — summary-address {ip-prefix/prefix-length | ip-prefix [netmask]} level
            — no summary-address {ip-prefix/prefix-length | ip-prefix [netmask]}
            — [no] traffic-engineering
```

Show Commands

```
show
     — router
              — isis
                      — adjacency [ip-int-name | ip-address | nbr-system-id] [detail]
                      — database [system-id | lsp-id] [detail] [level level]
                      — hostname
                      — interface [ip-int-name | ip-address] [detail]
                      — routes [ipv4-unicast]
                      — spf [detail]
                      — spf-log [detail]
                      — statistics
                      — status
                      — summary-address [ip-prefix [/prefix-length]]
                      — topology [ipv4-unicast] [detail]
```

Clear Commands

```
clear
     — router
              — isis
                      — adjacency [system-id]
                      — database [system-id]
                      — export
                      — spf-log
                      — statistics
```

Debug Commands

```
debug
     — router
              — isis
                      — [no] adjacency [ip-int-name | ip-address | nbr-system-id]
                      — [no] cspf
                      — interface [ip-int-name | ip-address]
                      — no interface
                      — leak [ip-address]
                      — no leak
                      — [no] lsdb [level-number] [system-id | lsp-id]
                      — [no] misc
                      — packet [packet-type] [ip-int-name | ip-address] [detail]
                      — no packet
                      — rtm [ip-address]
                      — no rtm
                      — [no] spf [level-number] [system-id]
```

Command Descriptions

Configuration Commands

Generic Commands

shutdown

Syntax [**no**] **shutdown**

Context config>router>isis
config>router>isis>interface

Description This command administratively disables the entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many entities must be explicitly enabled using the **no shutdown** command.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they can be deleted.

Unlike other commands and parameters where the default state is not indicated in the configuration file, **shutdown** and **no shutdown** are always indicated in system-generated configuration files.

The **no** form of the command puts an entity into the administratively enabled state.

Default **IS-IS Global —** the IS-IS protocol is created in the **no shutdown** state

IS-IS Interface — when an IP interface is configured as an IS-IS interface, IS-IS on the interface is in the **no shutdown** state by default

Global Commands

isis

Syntax [**no**] **isis**

Context config>router

Description This command activates IS-IS on the router and enables access to the context to define IS-IS parameters.

The **no** form of the command deletes the IS-IS protocol instance and removes all configuration parameters.

Default **no isis**

area-id

Syntax [**no**] **area-id** *area-address*

Context config>router>isis

Description This command configures the area ID portion of the Network Service Access Point (NSAP) address, which identifies a point of connection to the network, such as a router interface.

Addresses in the IS-IS protocol are based on the ISO NSAP addresses and Network Entity Titles (NETs), not IP addresses. NET addresses are constructed similarly to NSAPs with the exception that the selector ID is always 00. NET addresses are exchanged in Hello and LSP PDUs. All NET addresses configured on the node are advertised to its neighbors.

Up to three area addresses can be configured.

NSAP addresses are divided into three parts. Only the area ID portion is configurable:

- area ID – a variable-length field between 1 and 13 bytes that identifies the area to which the router belongs. This field includes the Authority and Format Identifier (AFI) as the first (most significant) byte and the area identifier.
- system ID – A 6-byte system identifier. This value is not configurable. The system ID is derived from the system or router ID and uniquely identifies the router.
- selector ID – A 1-byte selector identifier that is always 00 for an NET. This value is not configurable.

For level 1 interfaces, neighbors can have different area IDs, but they must have at least one area ID (AFI + area) in common. Sharing a common area ID, they become neighbors and area merging between the potentially different areas can occur.

For level 2 interfaces, neighbors can have different area IDs. However, if they have no area IDs in common, they become only level 2 neighbors and only level 2 LSPs are exchanged.

For level 1/2 interfaces, neighbors can have different area IDs. If they have at least one area ID (AFI + area) in common, they become neighbors. In addition to exchanging level 2 LSPs, area merging between potentially different areas can occur.

If multiple **area-id** commands are entered, the system ID of all subsequent entries must match the system ID of the first area address.

The **no** form of the command removes the area address.

Default **no area-id —** no area address is assigned

Parameters *area-address* — the area ID, from 1 to 13 bytes (if fewer than 13 bytes are entered, the rest of the field is padded with zeros)

authentication-check

Syntax [no] **authentication-check**

Context config>router>isis

Description This command sets an authentication check to reject PDUs that do not match the type or key requirements.

The default behavior when authentication is configured is to reject all IS-IS protocol PDUs that have a mismatch in either the authentication type or authentication key.

When **no authentication-check** is configured, authentication PDUs are generated and IS-IS PDUs are authenticated on receipt. However, although mismatches cause an event to be generated, the mismatches will not be rejected.

Default **authentication-check**

authentication-key

Syntax **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
no authentication-key

Context config>router>isis
config>router>isis>level

Description This command sets the authentication key used to verify PDUs sent by neighboring routers on the interface. Neighboring routers use passwords to authenticate PDUs sent from an interface. For authentication to work, both the authentication key and the authentication type on a segment must match. The authentication-type command must also be entered.

To configure authentication on the global level, configure this command in the **config>router>isis** context. When this parameter is configured on the global level, all PDUs are authenticated, including the Hello PDU.

To override the global setting for a specific level, configure the **authentication-key** command in the **config>router>isis>level** context. When configured within the specific level, Hello PDUs are not authenticated.

By default, no authentication key is configured.

The **no** form of the command removes the authentication key.

Default **no authentication-key**

Parameters *authentication-key* — the authentication key can be any combination of ASCII characters up to 254 characters in length (unencrypted). If spaces are used in the string, the entire string must be enclosed in double quotes (“ ”).

hash-key — the hash key can be any combination of ASCII characters up to 352 characters in length (encrypted) or 407 characters in length (if the **hash2** parameter is used). If spaces are used in the string, the entire string must be enclosed in double quotes (“ ”).

This is useful when a user must configure the parameter, but for security purposes, the actual unencrypted key value is not provided.

hash — specifies that the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2 — specifies that the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted hash form is assumed.

authentication-type

Syntax **authentication-type** {**password** | **message-digest**}
no authentication-type

Context config>router>isis
config>router>isis>level

Description This command enables either simple password or message-digest authentication in the global IS-IS or IS-IS level context. Both the authentication key and the authentication type on a segment must match. The authentication-key command must also be entered.

Configure the authentication type at the global level in the **config>router>isis** context. Configure or override the global setting by configuring the authentication type in the **config> router>isis>level** context.

The **no** form of the command disables authentication.

Default **no authentication-type**

Parameters **password** — enables simple password (plain text) authentication. If authentication is enabled and no authentication type is specified in the command, simple password authentication is enabled.

message-digest — enables message-digest MD5 authentication in accordance with RFC 1321. If this option is configured, at least one message-digest-key must be configured.

csnp-authentication

Syntax [**no**] **csnp-authentication**

Context config>router>isis
config>router>isis>level

Description This command enables authentication of individual IS-IS packets of complete sequence number PDUs (CSNPs).

The **no** form of the command suppresses authentication of CSNP packets.

Default **csnp-authentication**

disable-ldp-sync

Syntax [**no**] **disable-ldp-sync**

Context config>router>isis

Description This command disables the IGP-LDP synchronization feature on all interfaces participating in the OSPF or IS-IS routing protocol. When this command is executed, IGP immediately advertises the actual value of the link cost for all interfaces that have the IGP-LDP synchronization enabled if the currently advertised cost is different. IGP-LDP synchronization will then be disabled for all interfaces. This command does not delete the interface configuration.

The **no** form of this command restores the default settings and re-enables IGP-LDP synchronization on all interfaces participating in the OSPF or IS-IS routing protocol and for which the **ldp-sync-timer** is configured (refer to the 7705 SAR OS Router Configuration Guide for information on configuring the **ldp-sync-timer**).

Default **no disable-ldp-sync**

export

Syntax **export** *policy-name* [*policy-name…(*up to 5 max)]
no export

Context config>router>isis

Description This command associates export route policies to determine which routes are exported from the route table to IS-IS.

If no export policy is specified, non-IS-IS routes are not exported from the routing table manager to IS-IS.

If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered will override the previous command. A maximum of five policy names can be specified.

The **no** form of the command removes all policies from the configuration.

Refer to the 7705 SAR OS Router Configuration Guide for information on defining route policies.

Default **no export —** no export route policies specified

Parameters *policy-name* — the export route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

The specified names must already be defined.

external-preference

Syntax **external-preference** *external-preference*
no external-preference

Context config>router>isis>level

Description This command configures the preference for IS-IS external routes for the IS-IS level. The preference for internal routes is set with the preference command.

The command configures the preference level for either level 1 or level 2 external routes. The default preferences are listed in Table 21.

A route can be learned by the router from different protocols, in which case, the costs are not comparable. When this occurs, the preference is used to decide which route will be used.

Different protocols should not be configured with the same preference. If this occurs, the tiebreaker is based on the default preferences as listed in Table 21.

Table 21: Route Preference Defaults by Route Type

Route Type	Preference	Configurable
Direct attached	0	No
Static routes	5	Yes
OSPF internal	10	Yes
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes

If multiple routes are learned with an identical preference using the same protocol, the lowest-cost route is used. If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, the decision of which route to use is determined by the configuration of ECMP in the **config>router** context. Refer to the 7705 SAR OS Router Configuration Guide for information on ECMP.

Note: To configure a preference for static routes, use the config>router>static-route command. Refer to the 7705 SAR OS Router Configuration Guide for information.

The **no** form of the command reverts to the default value.

Default **external-preference 160 —** for IS-IS level 1 external routes

external-preference 165 — for IS-IS level 2 external routes

Parameters *external-preference* — the preference for external routes at this level, expressed as a decimal integer

Values 1 to 255

hello-authentication

Syntax [**no**] **hello-authentication**

Context config>router>isis
config>router>isis>level

Description This command enables authentication of individual IS-IS Hello PDUs.

The **no** form of the command suppresses authentication of Hello PDUs.

Default **hello-authentication**

level

Syntax **level** {**1** | **2**}

Context config>router>isis
config>router>isis>interface

Description This command creates the context to configure IS-IS level 1 or level 2 area attributes.

To reset global and/or interface level parameters to the default, the following commands must be entered independently:

level> no hello-authentication-key
level> no hello-authentication-type
level> no hello-interval
level> no hello-multiplier
level> no metric
level> no passive
level> no priority

Special Cases

Global IS-IS Level — the **config>router>isis** context configures default global parameters for both level 1 and level 2 interfaces

IS-IS Interface Level — the **config>router>isis>interface** context configures IS-IS operational characteristics of the interfaces at level 1 and/or level 2. A logical interface can be configured on one level 1 and one level 2 interface. In this case, each level can be configured independently and parameters must be removed independently.

Default **level 1 or level 2**

Parameters **1 —** specifies that the router or interface is a level 1 router or interface

2 — specifies that the router or interface is a level 2 router or interface

level-capability

Syntax **level-capability** {**level-1** | **level-2** | **level-1/2**}
no level-capability

Context config>router>isis
config>router>isis>interface

Description This command configures the routing level for the IS-IS instance.

An IS-IS router and IS-IS interface can operate at level 1, level 2, or both level 1 and level 2.

A level 1 adjacency can be established if there is at least one area address shared by this router and a neighbor. A level 2 adjacency cannot be established over this interface.

A level 1/2 adjacency is created if the neighbor is also configured as a level 1/2 router and has at least one area address in common. A level 2 adjacency is established if there are no common area IDs.

A level 2 adjacency is established if another router is configured as a level 2 or level 1/2 router with interfaces configured as level 1/2 or level 2. Level 1 adjacencies will not established over this interface.

Table 22 lists capability combinations and the potential adjacencies that can be formed.

Table 22: Potential Adjacency Capabilities

Global Level	Interface Level	Potential Adjacency
Level 1/2	Level 1/2	Level 1 and/or level 2
Level 1/2	Level 1	Level 1 only
Level 1/2	Level 2	Level 2 only
Level 2	Level 1/2	Level 2 only
Level 2	Level 2	Level 2 only
Level 2	Level 1	None
Level 1	Level 1/2	Level 1 only
Level 1	Level 2	None
Level 1	Level 1	Level 1 only

The **no** form of the command removes the level capability from the configuration.

Special Cases

IS-IS Router — in the **config>router>isis** context, changing the level capability performs a restart on the IS-IS protocol instance

IS-IS Interface — in the **config>router>isis>interface** context, changing the level capability performs a restart of IS-IS on the interface

Default **level-1/2**

Parameters **level-1** — specifies that the router or interface can operate at level 1 only

level-2 — specifies that the router or interface can operate at level 2 only

level-1/2 — specifies that the router or interface can operate at both level 1 and level 2

lsp-lifetime

Syntax **lsp-lifetime** *seconds*
no lsp-lifetime

Context config>router>isis

Description This command sets the time interval for LSPs originated by the router to be considered valid by other routers in the domain.

Each LSP received is maintained in an LSP database until the LSP lifetime expires, unless the originating router refreshes the LSP. By default, each router refreshes its LSPs every 20 min (1200 s) so that other routers will not age out the LSP.

The LSP lifetime value should be greater than the refresh value; otherwise, the LSP will be aged out before being refreshed.

The **no** form of the command reverts to the default value.

Default **1200**

Parameters *seconds* — the interval for LSPs originated by the route to be considered valid by other routers in the domain

Values 350 to 65335

lsp-wait

Syntax **lsp-wait** *lsp-wait* [*lsp-initial-wait* [*lsp-second-wait*]]
no lsp-wait

Context config>router>isis

Description This command is used to customize the throttling of IS-IS LSP generation. Timers that determine when to generate the first, second, and subsequent LSPs can be controlled with this command. Subsequent LSPs are generated at increasing intervals of the **lsp-second-wait** timer until a maximum value is reached.

Parameters *lsp-wait* — the maximum interval, in seconds, between two consecutive LSPs being generated

Values 1 to 120

Default 5

lsp-initial-wait — the initial LSP generation delay, in seconds

Values 0 to 100

Default 0

lsp-second-wait — the hold time, in seconds, between the generation of the first and second LSPs

Values 1 to 100

Default 1

overload

Syntax **overload** [**timeout** *seconds*]
no overload

Context config>router>isis

Description This command administratively sets the IS-IS router to operate in the overload state for a specific time period, in seconds, or indefinitely.

During normal operation, the router may be forced to enter an overload state due to a lack of resources. When in the overload state, the router is only used if the destination is reachable by the router and will not be used for other transit traffic.

If a time period is specified, the overload state persists for the configured length of time. If no time is specified, the overload state operation is maintained indefinitely.

The **overload** command can be useful in circumstances where the router is overloaded or used prior to executing a **shutdown** command to divert traffic around the router.

The **no** form of the command causes the router to exit the overload state.

Default **no overload**

Parameters *seconds* — the number of seconds that the router remains in the overload state

Values 60 to 1800

Default infinity (overload state maintained indefinitely)

overload-on-boot

Syntax **overload-on-boot** [**timeout** *seconds*]
no overload-on-boot

Context config>router>isis

Description When the router is in an overload state, the router is used only if there is no other router to reach the destination. This command configures IS-IS in the overload state upon bootup until one of the following events occurs:

- the timeout timer expires
- the current overload state is manually overridden with the **no overload** command

The **no overload** command does not affect the **overload-on-boot** function. If the overload state is cleared with the **no overload** command, the router will still re-enter the overload state after rebooting.

If no timeout is specified, IS-IS will go into the overload state indefinitely after a reboot. After the reboot, the IS-IS status will display a permanent overload state:

- L1 LSDB Overload : Manual on boot (Indefinitely in overload)
- L2 LSDB Overload : Manual on boot (Indefinitely in overload)

This state can be cleared with the **no overload** command.

If a timeout value is specified, IS-IS will go into the overload state for the configured timeout after a reboot. After the reboot, the IS-IS status will display the remaining time that the system stays in overload:

- L1 LSDB Overload : Manual on boot (Overload Time Left : 17)
- L2 LSDB Overload : Manual on boot (Overload Time Left : 17)

The overload state can be cleared before the timeout expires with the **no overload** command.

The **no** form of the command removes the overload-on-boot functionality from the configuration.

Default **no overload-on-boot**

Parameters *seconds* — the number of seconds that the router remains in the overload state after rebooting

Values 60 to 1800

Default 60

preference

Syntax **preference** *preference*
no preference

Context config>router>isis>level

This command configures the preference for IS-IS level 1 or level 2 internal routes.

A route can be learned by the router from different protocols, in which case, the costs are not comparable. When this occurs, the preference is used to decide which route will be used.

Different protocols should not be configured with the same preference. If this occurs, the tiebreaker is based on the default preferences as listed in Table 21. If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, the decision of which route to use is determined by the configuration of ECMP in the **config>router** context. Refer to the 7705 SAR OS Router Configuration Guide for information on ECMP.

The **no** form of the command reverts to the default value.

Default **preference 15** — for IS-IS level 1 internal routes

preference 18 — for IS-IS level 2 internal routes

Parameters *preference* — the preference for internal routes expressed as a decimal integer

Values 1 to 255

psnp-authentication

Syntax [**no**] **psnp-authentication**

Context config>router>isis
config>router>isis>level

Description This command enables authentication of individual IS-IS packets of partial sequence number PDUs (PSNPs).

The **no** form of the command suppresses authentication of PSNP packets.

Default **psnp-authentication**

reference-bandwidth

Syntax **reference-bandwidth** *bandwidth-in-kbps*
no reference-bandwidth

Context config>router>isis

Description This command configures the reference bandwidth for the costing of interfaces based on their underlying link speed.

In order to calculate the lowest cost to reach a specific destination, each configured level on each interface must have a cost. If the reference bandwidth is defined, the cost is calculated using the following formula:

cost = reference bandwidth/bandwidth

If the reference bandwidth is configured as 10 Gbytes (10 000 000 000), a 100 Mb/s interface has a default metric of 100. In order for metrics in excess of 63 to be configured, wide metrics must be deployed (see the wide-metrics-only command).

If the reference bandwidth is not configured, all interfaces have a default metric of 10.

The **no** form of the command returns the reference bandwidth to the default value.

Default **no reference-bandwidth** (all interfaces have a metric of 10)

Parameters *bandwidth-in-kbps* — the reference bandwidth in kilobits per second expressed as a decimal integer

Values 1 to 100000000

spf-wait

Syntax **spf-wait** *spf-wait* [*spf-initial-wait* [*spf-second-wait*]]
no spf-wait

Context config>router>isis

Description This command defines the maximum interval, in seconds and milliseconds, between two consecutive SPF calculations. Timers that determine when to initiate the first, second, and subsequent SPF calculations after a topology change occurs can be controlled with this command. Subsequent SPF runs (if required) will occur at doubling intervals of the *spf-second-wait* interval. For example, if the *spf-second-wait* interval is 1000, the next SPF will run after 2000 ms, and the next SPF after that will run after 4000 ms, and so on, until it reaches the **spf-wait** value. The SPF interval will stay at the **spf-wait** value until there are no more SPF runs scheduled in that interval. After a full interval without any SPF runs, the SPF interval will drop back to *spf-initial-wait*.

The timer must be entered in increments of 100 ms. Values entered that do not match this requirement will be rejected.

The **no** form of this command returns to the default.

Default **no spf-wait**

Parameters *spf-wait* — the maximum interval, in seconds, between two consecutive SPF calculations

Values 1 to 120

Default 10

spf-initial-wait — the initial SPF calculation delay in milliseconds after a topology change

Values 10 to 100000

Default 1000

spf-second-wait — the hold time in milliseconds between the first and second SPF calculation

Values 10 to 100000

Default 1000

summary-address

Syntax **summary-address** {*ip-prefix/prefix-length* | *ip-prefix* [*netmask*]} **level**
no summary-address {*ip-prefix/prefix-length* | *ip-prefix* [*netmask*]}

Context config>router>isis

Description This command creates summary addresses.

Default **no summary-address**

Parameters *ip-prefix/prefix-length* — IP prefix and mask length

Values	*ip-prefix*	a.b.c.d (host bits must be 0)
	prefix-length	0 to 32

netmask — the subnet mask in dotted-decimal notation

Values 0.0.0.0 to 255.255.255.255 (network bits all 1 and host bits all 0)

level — the IS-IS level

Values level-1, level-2, level-1/2

traffic-engineering

Syntax [**no**] **traffic-engineering**

Context config>router>isis

Description This command enables traffic engineering and determines if IGP shortcuts are required.

The **no** form of the command disables traffic-engineered route calculations.

Default **no traffic-engineering**

wide-metrics-only

Syntax [**no**] **wide-metrics-only**

Context config>router>isis>level

Description This command enables the exclusive use of wide metrics in the LSPs for the level number. Narrow metrics can have values between 1 and 63. IS-IS can generate two TLVs, one for the adjacency and one for the IP prefix. In order to support traffic engineering, wider metrics are required. When wide metrics are used, a second pair of TLVs are added for the adjacency and the IP prefix.

By default, both sets of TLVs are generated. When **wide-metrics-only** is configured, IS-IS only generates the pair of TLVs with wide metrics for that level.

The **no** form of the command reverts to the default value.

Default **no wide-metrics-only**

Interface Commands

interface

Syntax [**no**] **interface** *ip-int-name*

Context config>router>isis

Description This command creates the context to configure an IS-IS interface.

When an area is defined, the interfaces belong to that area. Interfaces cannot belong to other areas.

If the interface is a POS channel, the OSI Network Layer Control Protocol (OSINLCP) is enabled when the interface is created and removed when the interface is deleted.

The **no** form of the command deletes the IS-IS interface configuration for this interface. The **shutdown** command in the **config>router>isis>interface** context can be used to disable an interface without removing the configuration for the interface.

Default **no interface**

Parameters *ip-int-name* — the IP interface name. Interface names must be unique within the group of defined IP interfaces for the **config>router>interface** command. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

If the IP interface name does not exist or does not have an IP address configured, an error message will be returned.

bfd-enable

Syntax [**no**] **bfd-enable ipv4**

Context config>router>isis>interface

Description This command enables the use of bidirectional forwarding (BFD) to control IPv4 adjacencies. By enabling BFD on a given IS-IS interface, the state of the interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for BFD are set via the BFD command under the IP interface.

The **no** form of this command removes BFD from the associated IPv4 adjacency.

Default **no bfd-enable ipv4**

csnp-interval

Syntax **csnp-interval** *seconds*
no csnp-interval

Context config>router>isis>interface

Description This command configures the interval, in seconds, to send complete sequence number PDUs (CSNPs) from the interface. IS-IS must send CSNPs periodically.

The **no** form of the command reverts to the default value.

Default **csnp-interval 10** – CSN PDUs are sent every 10 s for LAN interfaces

csnp-interval 5 – CSN PDUs are sent every 5 s for point-to-point interfaces

Parameters *seconds* — the CSNP interval expressed in seconds

Values 1 to 65535

hello-authentication-key

Syntax **hello-authentication-key** {*authentication-key* | *hash-key*} [**hash** | **hash2**]
no hello-authentication-key

Context config>router>isis>interface
config>router>isis>interface>level

Description This command configures the authentication key (password) for Hello PDUs. Neighboring routers use the password to verify the authenticity of Hello PDUs sent from this interface. Both the hello authentication key and the hello authentication type on a segment must match. The hello-authentication-type command must also be entered.

To configure the hello authentication key for all levels configured for the interface, use the **hello-authentication-key** command in the **config>router>isis>interface** context.

To configure or override the hello authentication key for a specific level, use the **hello-authentication-key** command in the **config>router>isis>interface>level** context.

If both IS-IS authentication and hello-authentication are configured, Hello messages are validated using hello authentication. If only IS-IS authentication is configured, it will be used to authenticate all IS-IS protocol PDUs, including Hello PDUs.

The **no** form of the command removes the hello authentication key from the configuration.

Default **no hello-authentication-key**

Parameters

authentication-key — the authentication key can be any combination of ASCII characters up to 254 characters in length (unencrypted). If spaces are used in the string, the entire string must be enclosed within double quotes (“ ”).

hash-key — the hash key can be any combination of ASCII characters up to 352 characters in length (encrypted) or 451 characters in length (if the **hash2** parameter is used). If spaces are used in the string, the entire string must be enclosed within double quotes (“ ”).

This is useful when a user must configure the parameter, but for security purposes, the actual unencrypted key value is not provided.

hash — specifies that the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2 — specifies that the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted hash form is assumed.

hello-authentication-type

Syntax

hello-authentication-type {**password** | **message-digest**}
no hello-authentication-type

Context

config>router>isis>interface
config>router>isis>interface>level

Description

This command enables Hello authentication at either the interface or level context. Both the hello authentication key and the hello authentication type on a segment must match. The hello-authentication-key command must also be entered.

To configure the hello authentication type for all levels configured for the interface, use the **hello-authentication-type** command in the **config>router>isis>interface** context.

To configure or override the hello authentication type for a specific level, use the **hello-authentication-type** command in the **config>router>isis>interface>level** context.

The **no** form of the command disables Hello PDU authentication.

Default

no hello-authentication-type

Parameters

password — enables simple password (plain text) authentication. If authentication is enabled and no authentication type is specified in the command, simple password authentication is enabled.

message-digest — enables message-digest MD5 authentication in accordance with RFC 1321. If this option is configured, at least one message-digest-key must be configured.

hello-interval

Syntax **hello-interval** *seconds*
no hello-interval

Context config>router>isis>interface>level

Description This command configures the interval between IS-IS Hello PDUs issued on the interface at this level.

The **no** form of this command reverts to the default value.

Default **3** – for designated inter-systems

9 – for non-designated inter-systems

Parameters *seconds* — the hello interval, in seconds, expressed as a decimal integer

Values 1 to 20000

hello-multiplier

Syntax **hello-multiplier** *multiplier*
no hello-multiplier

Context config>router>isis>interface>level

Description This command configures the number of missing Hello PDUs from a neighbor after which the router declares the adjacency down.

The **no** form of this command reverts to the default value.

Default **3**

Parameters *multiplier* — the multiplier for the hello interval, in seconds, expressed as a decimal integer

Values 2 to 100

interface-type

Syntax **interface-type** {**broadcast** | **point-to-point**}
no interface-type

Context config>router>isis>interface

Description This command configures the interface type to be either broadcast or point-to-point.

Use this command to set the interface type of an Ethernet link to point-to-point to avoid having to carry the broadcast adjacency maintenance overhead of the Ethernet link, provided the link is used as a point-to-point link.

If the interface type is not known when the interface is added to IS-IS, and the IP interface is subsequently bound (or moved) to a different interface type, this command must be entered manually.

The **no** form of the command reverts to the default value.

Default **broadcast** – if the physical interface is Ethernet or unknown

point-to-point – if the physical interface is T1, E1, or SONET/SDH

Parameters **broadcast —** configures the interface to maintain this link as a broadcast network. To significantly improve adjacency forming and network convergence, a network should be configured as point-to-point if only two routers are connected, even if the network is a broadcast media such as Ethernet.

point-to-point — configures the interface to maintain this link as a point-to-point link

lsp-pacing-interval

Syntax **lsp-pacing-interval** *milliseconds*
no lsp-pacing-interval

Context config>router>isis>interface

Description This command configures the interval between link-state PDUs (LSPs) sent from this interface. Controlling the time between LSPs ensures that adjacent neighbors are not being bombarded with excessive data.

A value of 0 means that no LSPs are sent from the interface.

The **no** form of the command reverts to the default value.

Default **100**

Parameters *milliseconds* — the interval that LSPs can be sent from the interface, expressed as a decimal integer

Values 0 to 65335

mesh-group

Syntax **mesh-group** [*value* | **blocked**]
no mesh-group

Context config>router>isis>interface

Description This command assigns an interface to a mesh group. Mesh groups limit the amount of flooding that occurs when a new or changed LSP is advertised throughout an area.

All routers in a mesh group should be fully meshed. When LSPs need to be flooded, only a single copy is received rather than one copy per neighbor.

To create a mesh group, configure the same mesh group value for each interface that is part of the mesh group. All routers must have the same mesh group value configured for all interfaces that are part of the mesh group.

To prevent an interface from flooding LSPs, the optional **blocked** parameter can be specified.

Caution: Configure mesh groups carefully. It is easy to create isolated islands that will not receive updates if other links fail.

The **no** form of the command removes the interface from the mesh group.

Default **no mesh-group**

Parameters *value* — the unique decimal integer that distinguishes this mesh group from other mesh groups on this router or on other routers

Values 1 to 2000000000

blocked — prevents an interface from flooding LSPs

metric

Syntax **metric** *metric*
no metric

Context config>router>isis>interface>level

Description This command configures the metric used for the level on this IS-IS interface.

To calculate the lowest cost to reach a given destination, each configured level on each interface must have a cost. The costs for each level on an interface may be different.

If the metric is not configured, the default value of 10 is used unless the reference-bandwidth is configured.

The **no** form of the command reverts to the default value.

Default **no metric (10)**

Parameters *metric* — the metric assigned to this level on this interface, expressed as a decimal integer

Values 1 to 16777215

passive

Syntax [**no**] **passive**

Context config>router>isis>interface
config>router>isis>interface>level

Description This command adds the passive attribute to the IS-IS interface, which causes the interface to be advertised as an IS-IS interface without running the IS-IS protocol. Normally, only interface addresses that are configured for IS-IS are advertised as IS-IS interfaces at the level that they are configured.

If the passive mode is enabled, the interface or the interface at the specified level ignores ingress IS-IS protocol PDUs and will not transmit IS-IS protocol PDUs.

The **no** form of the command removes the passive attribute.

Default Service interfaces defined with the **config**>**router**>**service-prefix** command are passive. All other interfaces are not passive.

priority

Syntax **priority** *number*
no priority

Context config>router>isis>interface>level

Description This command configures the priority of the IS-IS interface that is used in an election of the designated router (DIS) on a multi-access network.

This parameter is only used if the interface is a broadcast type.

The priority is included in Hello PDUs transmitted by the interface on a multi-access network. The router with the highest priority becomes the designated router. The designated router is responsible for sending LSPs about the network and the routers attached to it.

The **no** form of the command reverts to the default value.

Default **64**

Parameters *number* — the priority for this interface at this level, expressed as a decimal integer

Values 0 to 127

retransmit-interval

Syntax **retransmit-interval** *seconds*
no retransmit-interval

Context config>router>isis>interface

Description This command specifies the interval, in seconds, that IS-IS will wait before retransmitting an unacknowledged LSP to an IS-IS neighbor.

If the retransmit interval expires and no acknowledgment has been received, the LSP will be retransmitted.

The **no** form of this command reverts to the default interval.

Default **5**

Parameters *seconds* — the retransmit interval, in seconds, expressed as a decimal integer

Values 1 to 65335

Show Commands

isis

Syntax **isis**

Context show>router

Description This command enables the context to display IS-IS information.

adjacency

Syntax **area** [*ip-int-name* | *ip-address* | *nbr-system-id*] [**detail**]

Context show>router>isis

Description This command displays information about IS-IS neighbors. If no parameters are specified, all adjacencies are displayed. If **detail** is specified, operational and statistical information is displayed.

Parameters *ip-int-name* — displays only adjacencies with the specified interface

ip-address — displays only adjacencies with the specified IP address

nbr-system-id — displays only the adjacency with the specified system ID

detail — displays detailed information about the adjacency

Output The following output is an example of IS-IS adjacency information, and Table 23 describes the fields for both summary and detailed outputs.

Sample Output

```
A:ALU-A# show router isis adjacency
===============================================================================
ISIS Adjacency
===============================================================================
System ID           Usage State  Hold Interface                 MT  Enab
-------------------------------------------------------------------------------
Dut-B                L1    Up     2    ip-to1                    Yes
Dut-B                L2    Up     2    ip-to2                    Yes
Dut-F                L1L2  Up     5    ip-303                    Yes
-------------------------------------------------------------------------------
Adjacencies : 3
===============================================================================
A:ALU-A#
```

Table 23: Show Adjacency Output Fields

Label	Description
`System ID`	System ID of the neighbor
`SNPA`	Subnetwork point of attachment (MAC address of next hop)
`Usage/L. Circ Typ`	Level on the interface: L1, L2, or L1/2
`Interface`	Interface name associated with the neighbor
`Up Time`	Length of time that the interface has been up
`State`	State of the adjacency: up, down, new, one-way, initializing, or rejected
`Priority`	Priority to become the designated router
`Nbr Sys Typ`	Level of the neighbor router: L1, L2, or L1/2
`Hold/Hold Time`	Hold time remaining for the adjacency
`Max Hold`	Maximum hold time for the adjacency
`Adj Level`	Level of the adjacent router
`MT Enab/MT Enabled`	Not applicable
`Topology`	Unicast

database

Syntax **database** [*system-id* | *lsp-id*] [**detail**] [**level** *level*]

Context show>router>isis

Description This command displays information about the IS-IS link-state database.

If the system ID and LSP ID are not specified, all database entries are listed.

Parameters *system-id* — displays only the LSPs related to the specified system ID

lsp-id — displays only the specified LSP (hostname)

detail — displays detailed information on the link-state database entries

level — displays information only for the specified level

Output The following outputs are examples of IS-IS database information:

- IS-IS summary database information (Sample Output, Table 24)
- IS-IS detailed database information (Sample Output, Table 25)

Sample Output

```
A:ALU-A# show router isis database
==============================================================================
ISIS Database
==============================================================================
LSP ID                                  Sequence Checksum Lifetime Attributes
------------------------------------------------------------------------------

Displaying Level 1 database
------------------------------------------------------------------------------
ALU-A.00-00                             0x7      0x51b4   1177     L1L2
Level (1) LSP Count : 1

Displaying Level 2 database
------------------------------------------------------------------------------
ALU-A.00-00                             0x7      0x51b4   1014     L1L2
Level (2) LSP Count : 1
==============================================================================
A:ALU-A#
```

Table 24: Show Database Summary Output Fields

Label	Description
LSP ID	LSP IDs are auto-assigned by the originating IS-IS node. The LSP ID consists of three sections: the first 6 bytes are the system ID for that node, followed by a single byte value for the pseudonode generated by that router, followed by a fragment byte that starts at 0. For example, if a router's system ID is 1800.0000.0029, the first LSP ID is 1800.0000.0029.00-00. If there are too many routes, LSP ID 1800.0000.0029.00-01 is created to contain the excess routes. If the router is the designated router (or designated intermediate system ([DIS]) on a broadcast network, a pseudonode LSP is created. Usually the internal circuit ID is used to determine the ID assigned to the pseudonode. For instance, for circuit 4, an LSP pseudonode with ID 1800.0000.0029.04-00 is created. The 7705 SAR learns hostnames and uses the hostname in place of the system ID.
Sequence	The sequence number of the LSP that allows other systems to determine if they have received the latest information from the source
Checksum	The checksum of the entire LSP packet
Lifetime	Length of time, in seconds, that the LSP remains valid

Table 24: Show Database Summary Output Fields (Continued)

Label	Description
Attributes	OV — the overload bit is set
	L1 — specifies a level 1 router
	L2 — specifies a level 2 router
	L1L2 — specifies a level 1/2 router
	ATT — the attachment bit is set; when set, the router can act as a level 2 router and can reach other areas

Sample Output

```
*A:ALU-A# show router isis database detail
==============================================================================
ISIS Database
==============================================================================

Displaying Level 1 database
------------------------------------------------------------------------------
LSP ID    : ALU-A.00-00                                 Level     : L1
Sequence  : 0x7                    Checksum  : 0x51b4   Lifetime  : 1079
Version   : 1                      Pkt Type  : 18       Pkt Ver   : 1
Attributes: L1L2                   Max Area  : 3
SysID Len : 6                      Used Len  : 50       Alloc Len : 1492

TLVs :
  Supp Protocols:
    Protocols     : IPv4
  IS-Hostname   : ALU-A
  Router ID   :
    Router ID   : 255.0.0.0

Level (1) LSP Count : 1

Displaying Level 2 database
------------------------------------------------------------------------------
LSP ID    : ALU-A.00-00                                 Level     : L2
Sequence  : 0x7                    Checksum  : 0x51b4   Lifetime  : 900
Version   : 1                      Pkt Type  : 20       Pkt Ver   : 1
Attributes: L1L2                   Max Area  : 3
SysID Len : 6                      Used Len  : 50       Alloc Len : 1492

TLVs :
  Supp Protocols:
    Protocols     : IPv4
  IS-Hostname   : ALU-A
  Router ID   :
    Router ID   : 255.0.0.0

Level (2) LSP Count : 1
==============================================================================
A:ALU-A#
```

Table 25: Show Database Detailed Output Fields

Label	Description
LSP ID	LSP IDs are auto-assigned by the originating IS-IS node. The LSP ID consists of three sections: the first 6 bytes are the system ID for that node, followed by a single byte value for the pseudonode generated by that router, followed by a fragment byte that starts at 0. For example, if a router's system ID is 1800.0000.0029, the first LSP ID is 1800.0000.0029.00-00. If there are too many routes, LSP ID 1800.0000.0029.00-01 is created to contain the excess routes. If the router is the designated router (or designated intermediate system ([DIS]) on a broadcast network, a pseudonode LSP is created. Usually the internal circuit ID is used to determine the ID assigned to the pseudonode. For instance, for circuit 4, an LSP pseudonode with ID 1800.0000.0029.04-00 is created. The 7705 SAR learns hostnames and uses the hostname in place of the system ID.
Sequence	The sequence number of the LSP that allows other systems to determine if they have received the latest information from the source
Checksum	The checksum of the entire LSP packet
Lifetime	Length of time, in seconds, that the LSP remains valid
Attributes	OV — the overload bit is set
	L1 — specifies a level 1 router
	L2 — specifies a level 2 router
	L1L2 — specifies a level 1/2 router
	ATT — the attachment bit is set; when set, the router can act as a level 2 router and can reach other areas
LSP Count	A sum of all the configured level 1 and level 2 LSPs
LSP ID	A unique identifier for each LSP, consisting of the system ID, pseudonode ID, and LSP name
Version	The version protocol ID extension – always set to 1
Pkt Type	The PDU type number
PkT Ver	The version protocol ID extension – always set to 1
Max Area	The maximum number of area addresses supported
SysID Len	The length of the system ID field (0 or 6)
Used Len	The actual length of the PDU

Table 25: Show Database Detailed Output Fields (Continued)

Label	Description
`Alloc Len`	The amount of memory space allocated for the LSP
`Area Address`	The area addresses to which the router is connected
`Supp Protocols`	The supported data protocols
`IS-Hostname`	The name of the router from which the LSP originated
`Virtual Flag`	`0` — level 1 routers report this octet as 0 to all neighbors
	`1` — indicates that the path to a neighbor is a level 2 virtual path used to repair an area partition
`Neighbor`	The routers running interfaces to which the router is connected
`Internal Reach`	A 32-bit metric A bit is added for the up/down transitions resulting from level 2 to level 1 route leaking
`IP Prefix`	The IP addresses that the router knows about by externally originated interfaces
`Metrics`	The routing metric used in the IS-IS link-state calculations

hostname

Syntax **hostname**

Context show>router>isis

Description This command displays the hostname database.

Output The following output is an example of hostname database information, and Table 26 describes the fields.

Sample Output

```
*A:ALU-A show router isis hostname

===============================================================================
Hosts
===============================================================================
System Id                Hostname
-------------------------------------------------------------------------------
2550.0000.0000           7705_custDoc
===============================================================================
```

Table 26: Show Hostname Database Output Fields

Label	Description
`System ID`	The system ID mapped to the hostname
`Hostname`	The hostname for the specified system ID

interface

Syntax **interface** [*ip-int-name* | *ip-address*] [**detail**]

Context show>router>isis

Description This command displays the details of the IS-IS interface, which can be identified by IP address or IP interface name. If neither is specified, all in-service interfaces are displayed.

Parameters *ip-int-name* — displays only the interface identified by this interface name

ip-address — displays only the interface identified by this IP address

detail — displays detailed information on the interface

Output The following outputs are examples of IS-IS interface information:

- IS-IS summary interface information (Sample Output, Table 27)
- IS-IS detailed interface information (Sample Output, Table 28)

Sample Output

```
A:ALU-A# show router isis interface system
==============================================================================
ISIS Interfaces
==============================================================================
Interface                        Level CircID  Oper State   L1/L2 Metric
------------------------------------------------------------------------------
system                           L1L2  1       Up           0/0
------------------------------------------------------------------------------
Interfaces : 1
==============================================================================
A:ALU-A#
```

Table 27: Show Interface Output Fields

Label	Description
`Interface`	The interface name
`Level`	The interface level: L1, L2, or L1L2
`CircID`	The circuit identifier
`Oper State`	`Up` — the interface is operationally up
	`Down` — the interface is operationally down
`L1/L2 Metric`	Interface metric for level 1 and level 2, if none are set to 0

Sample Output

```
A:ALU-A# show router isis interface detail
==============================================================================
ISIS Interfaces
==============================================================================
------------------------------------------------------------------------------
Interface      : system                          Level Capability: L1L2
Oper State     : Up                              Admin State     : Up
Auth Type      : None
Circuit Id     : 1                               Retransmit Int. : 5
Type           : Pt-to-Pt                        LSP Pacing Int. : 100
Mesh Group     : Inactive                        CSNP Int.       : 10
Bfd Enabled    : No
Te Metric      : 0                               Te State        : Down
Admin Groups   : None
Ldp Sync       : outOfService                    Ldp Sync Wait   : Disabled
Ldp Timer State: Disabled                        Ldp Tm Left     : 0

  Level        : 1                               Adjacencies     : 0
  Auth Type    : None                            Metric          : 0
  Hello Timer  : 9                               Hello Mult.     : 3
  Priority     : 64
  Passive      : No


  Level        : 2                               Adjacencies     : 0
  Auth Type    : None                            Metric          : 0
  Hello Timer  : 9                               Hello Mult.     : 3
  Priority     : 64
  Passive      : No

==============================================================================
A:ALU-A#
```

Table 28: Show Interface Detailed Output Fields

Label	Description
`Interface`	The interface name
`Level Capability`	The routing level for the IS-IS routing process
`Oper State`	`Up` — the interface is operationally up
	`Down` — the interface is operationally down
`Admin State`	`Up` — the interface is administratively up
	`Down` — the interface is administratively down
`Auth Type`	The authentication type for the interface
`Circuit Id`	The circuit identifier
`Retransmit Int.`	The length of time, in seconds, that IS-IS will wait before retransmitting an unacknowledged LSP to an IS-IS neighbor

Table 28: Show Interface Detailed Output Fields (Continued)

Label	Description
`Type`	The interface type: point-to-point or broadcast
`LSP Pacing Int.`	The interval between LSPs sent from this interface
`Mesh Group`	Indicates whether a mesh group has been configured
`CSNP Int.`	The time, in seconds, that complete sequence number PDUs (CSNPs) are sent from the interface
`BFD Enabled`	Indicates whether BFD is enabled or disabled
`TE Metric`	The TE metric configured for this interface. This metric is flooded out in the TE metric sub-TLV in the IS-IS-TE LSPs. Depending on the configuration, either the TE metric value or the native IS-IS metric value is used in CSPF computations.
`TE State`	The MPLS interface TE status from the IS-IS standpoint
`Admin Groups`	The bitmap inherited from the MPLS interface that identifies the admin groups to which this interface belongs
`Ldp Sync`	Specifies whether the IGP-LDP synchronization feature is enabled or disabled on all interfaces participating in the IS-IS routing protocol
`Ldp Sync Wait`	The time to wait for the LDP adjacency to come up
`Ldp Timer State`	The state of the LDP sync time left on the IS-IS interface
`LDP TM Left`	The time left before IS-IS reverts back to advertising normal metrics for this interface
`Level`	The interface level
`Adjacencies`	The number of adjacencies for this interface
`Auth Type`	The authentication type for the interface level
`Metric`	Indicates whether a metric has been configured for the interface level
`Hello Timer`	The interval between IS-IS Hello PDUs issued on the interface at this level
`Hello Mult.`	Not applicable
`Priority`	The priority of the IS-IS interface that is used in an election of the designated router on a multi-access network
`Passive`	Indicates if passive mode is enabled or disabled; if enabled, the interface is advertised as an IS-IS interface without running the IS-IS protocol

routes

Syntax **routes [ipv4-unicast]**

Context show>router>isis

Description This command displays the routes in the IS-IS routing table.

Parameters **ipv4-unicast** — displays IPv4 unicast parameters

Output The following output is an example of IS-IS route information, and Table 29 describes the fields.

Sample Output

```
A:ALU-A# show router isis routes
===============================================================================
Route Table
===============================================================================
Prefix                             Metric       Lvl/Typ Ver.   SysID/Hostname
  NextHop                           MT
-------------------------------------------------------------------------------
10.10.1.0/24                       10           1/Int.  5       ALU-A
   10.10.1.2
===============================================================================
A:ALU-A#
```

Table 29: Show Routing Table Output Fields

Label	Description
`Prefix`	The route prefix and mask
`Metric MT`	The metric of the route
`Lvl/Typ`	The level (1 or 2) and the route type (internal or external)
`Ver.`	The SPF version that generated the route
`SysID/Hostname`	The hostname for the specific system ID
`NextHop`	The system ID of the next hop (or the hostname, if possible)

spf

Syntax **spf [detail]**

Context show>router>isis

Description This command displays information about shortest path first (SPF) calculations.

Parameters **detail** — displays detailed SPF information

Output The following output is an example of SPF information, and Table 30 describes the fields for both summary and detailed outputs.

Sample Output

```
A:ALU-A# show router isis spf
===============================================================================
Path Table
===============================================================================
Node                                    Interface                   Nexthop
-------------------------------------------------------------------------------
================================================================================
A:ALU-A#
```

Table 30: Show SPF Output Fields

Label	Description
Node	The route node and mask
Interface	The outgoing interface name for the route
Nexthop	The system ID next hop or hostname
Metric	The metric of the route
SNPA	The subnetwork point of attachment where a router is physically connected to a subnetwork

spf-log

Syntax **spf-log** [**detail**]

Context show>router>isis

Description This command displays the last 20 SPF events.

Parameters **detail** — displays detailed information about the SPF events

Output The following output is an example of SPF events, and Table 31 describes the fields.

Sample Output

```
A:ALU-A# show router isis spf-log

===============================================================================
ISIS SPF Log
===============================================================================
When                      Duration       L1 Nodes   L2 Nodes   Event Count
-------------------------------------------------------------------------------
01/30/2009 11:01:54       <0.01s         1          1          3
-------------------------------------------------------------------------------
```

Table 31: Show SPF Log Output Fields

Label	Description
`When`	The timestamp when the SPF run started on the system
`Duration`	The time (in hundredths of seconds) required to complete the SPF run
`L1 Nodes`	The number of level 1 nodes involved in the SPF run
`L2 Nodes`	The number of level 2 nodes involved in the SPF run
`Event Count`	The number of SPF events that triggered the SPF calculation
`Log Entries`	The total number of log entries

statistics

Syntax **statistics**

Context show>router>isis

Description This command displays information about IS-IS traffic statistics.

Output The following output is an example of IS-IS statistical information, and Table 32 describes the fields.

Sample Output

```
A:ALU-A# show router isis statistics
===============================================================================
ISIS Statistics
===============================================================================
ISIS Instance     : 1                       SPF Runs       : 0
Purge Initiated   : 0                       LSP Regens.    : 39

CSPF Statistics
Requests          : 0                       Request Drops  : 0
Paths Found       : 0                       Paths Not Found: 0

-------------------------------------------------------------------------------
```

```
PDU Type   Received   Processed  Dropped    Sent       Retransmitted
-------------------------------------------------------------------------------
LSP        0          0          0          0          0
IIH        0          0          0          0          0
CSNP       0          0          0          0          0
PSNP       0          0          0          0          0
Unknown    0          0          0          0          0
===============================================================================
A:ALU-A#
```

Table 32: Show IS-IS Statistics Output Fields

Label	Description
`Purge Initiated`	The number of times that purges have been initiated
`SPF Runs`	The number of times that SPF calculations have been made
`LSP Regens`	The number of LSP regenerations
`Requests`	The number of CSPF requests made to the protocol
`Paths Found`	The number of responses to CSPF requests for which paths satisfying the constraints were found
`PDU Type`	The PDU (packet) type
`Received`	The number of LSPs received by this instance of the protocol
`Processed`	The number of LSPs processed by this instance of the protocol
`Dropped`	The number of LSPs dropped by this instance of the protocol
`Sent`	The number of LSPs sent out by this instance of the protocol
`Retransmitted`	The number of LSPs that had to be retransmitted by this instance of the protocol

status

Syntax **status**

Context show>router>isis

Description This command displays the general status of IS-IS.

Output The following output is an example of IS-IS status information, and Table 33 describes the fields.

Sample Output

```
A:ALU-A# show router isis status
=============================================================================
ISIS Status
=============================================================================
System Id            : 2550.0000.0000
Admin State          : Up
Ipv4 Routing         : Enabled
Last Enabled         : 05/21/2009 19:44:42
Level Capability     : L1L2
Authentication Check : True
Authentication Type  : None
CSNP-Authentication  : Enabled
HELLO-Authentication : Enabled
PSNP-Authentication  : Enabled
Traffic Engineering  : Disabled
LSP Lifetime         : 1200
LSP Wait             : 5 sec (Max)   0 sec (Initial)   1 sec (Second)
Adjacency Check      : loose
L1 Auth Type         : none
L2 Auth Type         : none
L1 CSNP-Authenticati*: Enabled
L1 HELLO-Authenticat*: Enabled
L1 PSNP-Authenticati*: Enabled
L1 Preference        : 15
L2 Preference        : 18
L1 Ext. Preference   : 160
L2 Ext. Preference   : 165
L1 Wide Metrics      : Disabled
L2 Wide Metrics      : Disabled
L1 LSDB Overload     : Disabled
L2 LSDB Overload     : Disabled
L1 LSPs              : 1
L2 LSPs              : 1
Last SPF             : 05/20/2009 14:21:53
SPF Wait             : 10 sec (Max)   1000 ms (Initial)   1000 ms (Second)
Export Policies      : None
Area Addresses       : 49.0001
Ldp Sync Admin State : Up
===============================================================================
A:ALU-A#
```

Table 33: Show IS-IS Status Output Fields

Label	Description
System Id	The system ID mapped to the hostname
Admin State	Up — IS-IS is administratively up
	Down — IS-IS is administratively down
IPv4 Routing	Enabled — IPv4 routing is enabled
	Disabled — IPv4 routing is disabled
Last Enabled	The date and time that IS-IS was last enabled on the router
Level Capability	The routing level for the IS-IS routing process
Authentication Check	True — all IS-IS mismatched packets are rejected
	False — authentication is performed on received IS-IS protocol packets but mismatched packets are not rejected
Authentication Type	The method of authentication used to verify the authenticity of packets sent by neighboring routers on an IS-IS interface
CSNP-Authentication	Indicates whether authentication of CSNP packets is enabled
HELLO-Authentication	Indicates whether authentication of Hello packets is enabled
PSNP Authentication	Indicates whether authentication of PSNP packets is enabled
Traffic Engineering	Enabled — TE is enabled for the router
	Disabled — TE is disabled; TE metrics are not generated and are ignored when received by this node
LSP Lifetime	Length of time that the LSPs originated by the router are to be considered valid by other routers in the domain
LSP Wait	Length of time that the router will generate the first, second, and subsequent LSPs
Adjacency Check	Type of adjacency check – always loose
L1 Auth Type	The method of authentication used to verify the authenticity of packets sent by neighboring routers to an IS-IS level 1 router
L2 Auth Type	The method of authentication used to verify the authenticity of packets sent by neighboring routers to an IS-IS level 2 router

Table 33: Show IS-IS Status Output Fields (Continued)

Label	Description
L1 CSNP-Authentication	Indicates whether authentication of CSNP packets is enabled on the level 1 router
L1 HELLO-Authentication	Indicates whether authentication of Hello packets is enabled on the level 1 router
L1 PSNP Authentication	Indicates whether authentication of PSNP packets is enabled on the level 1 router
L1 Preference	The preference level for level 1 internal routes
L2 Preference	The preference level for level 2 internal routes
L1 Ext. Preference	The preference level for level 1 external routes
L2 Ext. Preference	The preference level for level 2 external routes
L1 Wide Metrics	Indicates whether wide metrics are enabled or disabled for level 1 routers
L2 Wide Metrics	Indicates whether wide metrics are enabled or disabled for level 2 routers
L1 LSDB Overload	Indicates whether link-state database overload is enabled or disabled for level 1 routers
L2 LSDB Overload	Indicates whether link-state database overload is enabled or disabled for level 2 routers
L1 LSPs	Number of LSPs sent on the level 1 router interface
L2 LSPs	Number of LSPs sent on the level 2 router interface
Last SPF	Date and time that the last SPF calculation was performed
SPF Wait	Length of time that the first, second, and subsequent SPF calculations are initiated after a topology change occurs
Export Policies	Indicates if export policies are applied to the router
Area Addresses	The number of area addresses (area IDs) configured for the router
LDP Sync Admin State	Indicates whether the IGP-LDP synchronization feature is enabled or disabled on all interfaces participating in the IS-IS routing protocol

summary-address

Syntax **summary-address** [*ip-prefix*[/*prefix-length*]]

Context show>router>isis

Description This command displays IS-IS summary addresses.

Parameters *ip-prefix/prefix-length* — IP prefix and mask length

Values

ip-prefix	a.b.c.d (host bits must be 0)
prefix-length	0 to 32

Output The following output is an example of IS-IS summary address information, and Table 34 describes the fields.

Sample Output

```
A:ALU-A# show router isis summary-address
===============================================================================
ISIS Summary Address
===============================================================================
Address                                                  Level
1.0.0.0/8                                                L1
3.3.3.3/32                                               L2
-------------------------------------------------------------------------------
A:ALU-A#
```

Table 34: Show Summary Address Output Fields

Label	Description
`Address`	The IP address
`Level`	The IS-IS level from which the prefix should be summarized

topology

Syntax **topology** [**ipv4-unicast**] [**detail**]

Context show>router>isis

Description This command displays IS-IS topology information.

Parameters **ipv4-unicast** — displays IPv4 unicast parameters

detail — displays detailed topology information

Output The following output is an example of IS-IS topology information, and Table 35 describes the fields.

Sample Output

```
A:ALU-A# show router isis topology

===============================================================================
Topology Table
===============================================================================
Node                                Interface                  Nexthop
-------------------------------------------------------------------------------
ALU-A.00                            ip-ser01                   ALU-A
===============================================================================
A:ALU-A#
```

Table 35: Show IS-IS Topology Output Fields

Label	Description
`Node`	The IP address
`Interface`	The interface name
`Nexthop`	The next hop IP address

Clear Commands

adjacency

Syntax **adjacency** [*system-id*]

Context clear>router>isis

Description This command clears and resets the entries from the IS-IS adjacency database.

Parameters *system-id* — 6-octet system identifier in the form xxxx.xxxx.xxxx

database

Syntax **database** [*system-id*]

Context clear>router>isis

Description This command removes the entries from the IS-IS link-state database that contains information about PDUs.

Parameters *system-id* — 6-octet system identifier in the form xxxx.xxxx.xxxx

export

Syntax **export**

Context clear>router>isis

Description This command re-evaluates the route policies for IS-IS.

spf-log

Syntax **spf-log**

Context clear>router>isis

Description This command clears the SPF log.

statistics

Syntax **statistics**

Context clear>router>isis

Description This command clears and resets all IS-IS statistics.

Debug Commands

adjacency

Syntax [**no**] **adjacency** [*ip-int-name* | *ip-address* | *nbr-system-id*]

Context debug>router>isis

Description This command enables or disables debugging for IS-IS adjacency.

Parameters *ip-int-name* — debugs only adjacencies with the specified interface

ip-address — debugs only adjacencies with the specified IP address

nbr-system-id — debugs only the adjacency with the specified system ID

cspf

Syntax [**no**] **cspf**

Context debug>router>isis

Description This command enables or disables debugging for an IS-IS constraint-based shortest path first (CSPF).

interface

Syntax **interface** [*ip-int-name* | *ip-address*]
no interface

Context debug>router>isis

Description This command enables or disables debugging for an IS-IS interface.

Parameters *ip-int-name* — the IP interface name. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

ip-address — the interface's IP address

Values a.b.c.d

leak

Syntax **leak** [*ip-address*]
no leak

Context debug>router>isis

Description This command enables or disables debugging for IS-IS leaks.

Parameters *ip-address* — the IP address to debug IS-IS leaks

Values a.b.c.d

lsdb

Syntax [**no**] **lsdb** [*level-number*] [*system-id* **|** *lsp-id*]

Context debug>router>isis

Description This command enables or disables debugging for the IS-IS link-state database.

Parameters *level-number* — 1 or 2

system-id — 6-octet system identifier in the form xxxx.xxxx.xxxx

lsp-id — the hostname (38 characters maximum)

misc

Syntax [**no**] **misc**

Context debug>router>isis

Description This command enables or disables debugging for miscellaneous IS-IS events.

packet

Syntax **packet** [*packet-type*] [*ip-int-name* **|** *ip-address*] [**detail**]
no packet

Context debug>router>isis

Description This command enables or disables debugging for IS-IS packets.

Parameters *packet-type* — the IS-IS packet type to debug

Values ptop-hello | l1-hello | l2-hello | l1-psnp | l2-psnp | l1-csnp | l2-csnp | l1-lsp | l2-lsp

ip-int-name — the IP interface name. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

ip-address — the IP address to debug

Values a.b.c.d

detail — provides detailed debugging information

rtm

Syntax **rtm** [*ip-address*]
no rtm

Context debug>router>isis

Description This command enables or disables debugging for the IS-IS routing table manager.

Parameters *ip-address* — the IP address to debug

Values a.b.c.d

spf

Syntax [**no**] **spf** [*level-number*] [*system-id*]

Context debug>router>isis

Description This command enables or disables debugging for IS-IS SPF.

Parameters *level-number* — 1 or 2

system-id — 6-octet system identifier in the form xxxx.xxxx.xxxx

BGP

In This Chapter

This chapter provides information to configure BGP.

Topics in this chapter include:

BGP Overview

Border Gateway Protocol (BGP) is an inter-autonomous system routing protocol. An autonomous system (AS) is a network or a group of routers logically organized and controlled by a common network administration. BGP enables routers to exchange network reachability information, including information about other ASs that traffic must traverse to reach other routers in other ASs. In order to implement BGP, the AS number must be specified in the `config>router` context. A 7705 SAR OS BGP configuration must contain at least one group and include information about at least one neighbor (peer).

AS paths are the routes to each destination. Other attributes, such as the path's origin, the system's route preference, aggregation, route reflection, and communities included in the AS path are called path attributes. When BGP interprets routing and topology information, loops can be detected and eliminated. Route preference for routes learned from the configured peer(s) can be enabled among groups of routes to enforce administrative preferences and routing policy decisions.

BGP Communication

There are two types of BGP peers: internal BGP (IBGP) peers and external BGP (EBGP) peers (see Figure 6). The 7705 SAR, Release 3.0, supports only IBGP; consequently, a 7705 SAR is intended to peer with other 7705 SARs, or with 7750 (or 7710) SRs within the same AS.

- Within an AS, IBGP is used to communicate with peers.
- Outside of an AS or between ASs, EBGP is used to communicate with peers in different autonomous systems. Routes received from a router in a different AS can be advertised to both EBGP and IBGP peers. For ASs with 7705 SARs, the EBGP communication must be handled by another service router, such as a 7750 SR or a 7710 SR.

Autonomous systems use BGP to share routing information — such as routes to each destination and information about the route or AS path — with other ASs. Routing tables contain lists of known routers, reachable addresses, and associated path cost metrics to each router. BGP uses the information and path attributes to compile a network topology.

Figure 6: BGP Configuration

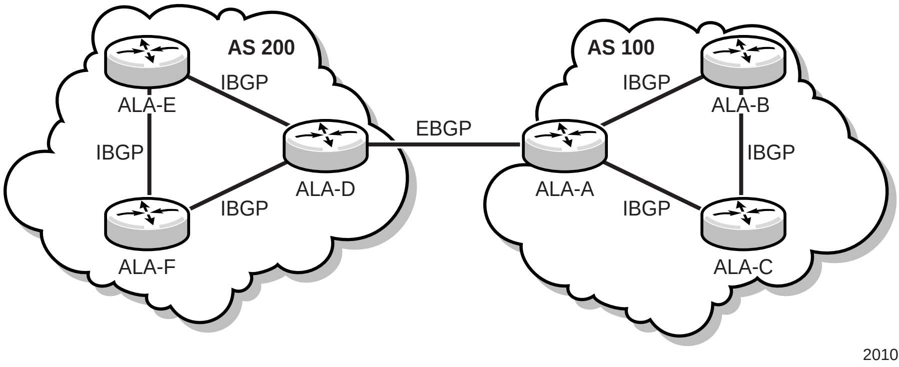

Message Types

Four message types are used by BGP to negotiate parameters, exchange routing information and indicate errors. They are:

- Open message — after a transport protocol connection is established, the first message sent by each side is an Open message. If the Open message is acceptable, a Keepalive message confirming the Open message is sent back. Once the Open message is confirmed, Update, Keepalive, and Notification messages can be exchanged.

  Open messages consist of the BGP header and the following fields:

  → version — the current BGP version number is 4

  → local AS number — the autonomous system number is configured in the `config>router` context

  → hold time — the maximum time BGP will wait between successive messages (either Keepalive or Update) from its peer, before closing the connection. Configure the local hold time within the `config>router>bgp` context.

  → BGP identifier — IP address of the BGP system or the router ID. The router ID must be a valid host address.

- Update message — Update messages are used to transfer routing information between BGP peers. The information contained in the packet can be used to construct a graph describing the relationships of the various autonomous systems. By applying rules, routing information loops and some other anomalies can be detected and removed from the inter-AS routing.

The Update messages consist of a BGP header and the following optional fields:

  - → unfeasible routes length — the length of the field that lists the routes being withdrawn from service because they are considered unreachable
  - → withdrawn routes — the associated IP address prefixes for the routes withdrawn from service
  - → total path attribute length — the total length of the path field that provides the attributes for a possible route to a destination
  - → path attributes — the path attributes presented in variable-length TLV format
  - → network layer reachability information (NLRI) — IP address prefixes of reachability information

- Keepalive message — Keepalive messages, consisting of only a 19-octet message header, are exchanged between peers frequently so hold timers do not expire. The Keepalive messages determine if a link is unavailable.
- Notification message — A Notification message is sent when an error condition is detected. The peering session is terminated and the BGP connection (TCP connection) is closed immediately after sending it.

Group Configuration and Peers

To enable BGP routing, participating routers must have BGP enabled and be assigned to an autonomous system, and the neighbor (peer) relationships must be specified. A router can belong to only one AS. TCP connections must be established in order for neighbors to exchange routing information and updates. Neighbors exchange BGP Open messages that include information such as AS numbers, BGP versions, router IDs, and hold-time values. Keepalive messages determine if a connection is established and operational. The hold-time value specifies the maximum time BGP will wait between successive messages (either Keepalive or Update) from its peer, before closing the connection.

In BGP, peers are arranged into groups. A group must contain at least one neighbor. A neighbor must belong to a group. Groups allow multiple peers to share similar configuration attributes.

Although neighbors do not have to belong to the same AS, they must be able to communicate with each other. If TCP connections are not established between two neighbors, the BGP peering session will not be established and updates will not be exchanged.

Peer relationships are defined by configuring the IP address of the routers that are peers of the local BGP system. When neighbor and peer relationships are configured, the BGP peers exchange Update messages to advertise network reachability information.

Hierarchical Levels

BGP parameters are initially applied at the global level. These parameters are inherited by the group and neighbor (peer) levels. Parameters can be modified and overridden on a level-specific basis. BGP command hierarchy consists of three levels:

- global level
- group level
- neighbor level

Many of the hierarchical BGP commands can be modified at different levels. The most specific value is used. That is, a BGP group-specific command takes precedence over a global BGP command. A neighbor-specific command takes precedence over a global BGP and group-specific command; for example, if you modify a BGP neighbor-level command default, the new value takes precedence over group- and global-level settings.

➡ **Note:** Careful planning is essential to implement commands that can affect the behavior of global, group, and neighbor levels. Because the BGP commands are hierarchical, analyze the values that can disable features at the global or group levels that must be enabled at the neighbor level. For example, if you enable the damping command at the global level but want it disabled only for a specific neighbor (not for all neighbors within the group), you cannot configure a double-no command (no no damping) to enable the feature.

Route Reflection

In a standard BGP configuration, all BGP speakers within an AS must have full BGP mesh to ensure that all externally learned routes are redistributed through the entire AS. IBGP speakers do not readvertise routes learned from one IBGP peer to another IBGP peer. If a network grows, scaling issues could emerge because of the full mesh configuration requirement. Instead of peering with all other IBGP routers in the network, each IBGP router only peers with a router configured as a route reflector.

Route reflection circumvents the full mesh requirement but maintains the full distribution of external routing information within an AS. Route reflection is effective in large networks because it is manageable, scalable, and easy to implement. Route reflection is implemented in autonomous systems with a large internal BGP mesh in order to reduce the number of IBGP sessions required within an AS.

A large AS can be subdivided into smaller ASs, called clusters. Route clusters are similar to these subautonomous systems and include route reflector(s) and clients. Each cluster contains at least one route reflector, which is responsible for redistributing route updates to all clients.

Route reflector clients do not need to maintain a full peering mesh between each other. They only require a peering to the route reflector(s) in their cluster. The route reflectors must maintain a full peering mesh between all non-clients within the AS.

Each route reflector must be assigned a cluster ID and specify which neighbors are clients and which are non-clients to determine which neighbors should receive reflected routes and which should be treated as a standard IBGP peer. Additional configuration is not required for the route reflector, aside from the typical BGP neighbor parameters.

BGP speakers within the AS that are not peers with the route reflector are called non-clients. Non-clients are peers to a route reflector but do not understand the route reflector attributes. Several BGP-speaking routers can peer with a route reflector. A route reflector forms peer connections to other route reflectors.

Figure 7 displays a simple configuration with several IBGP 7705 SARs.

When SR-A receives a route from SR-1 (an external neighbor), it must advertise route information to SAR-B, SAR-C, SAR-D, SAR-E, and SAR-F. To prevent loops, IBGP learned routes are not readvertised to other IBGP peers.

Figure 7: Fully Meshed BGP Configuration

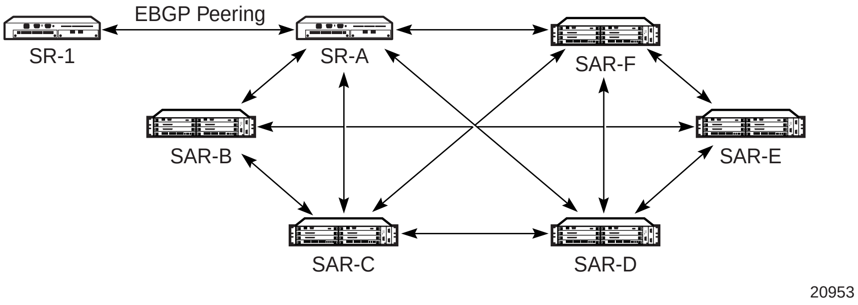

When route reflectors are configured, the routers within a cluster do not need to be fully meshed. Figure 7 depicts a fully meshed network and Figure 8 depicts the same network but with route reflectors configured to minimize the IBGP mesh between SR-A, SAR-B, SAR-C, and SAR-D. SR-A, configured as the route reflector, is responsible for redistributing route updates to clients SAR-B, SAR-C, and SAR-D. IBGP peering between SAR-B, SAR-C and SAR-D is not necessary because even IBGP learned routes are reflected to the route reflector's clients.

In Figure 8, SAR-E and SAR-F are shown as non-clients of the route reflector. As a result, a full mesh of IBGP peerings must be maintained between SR-A, SAR-E, and SAR-F.

Figure 8: BGP Configuration with Route Reflectors

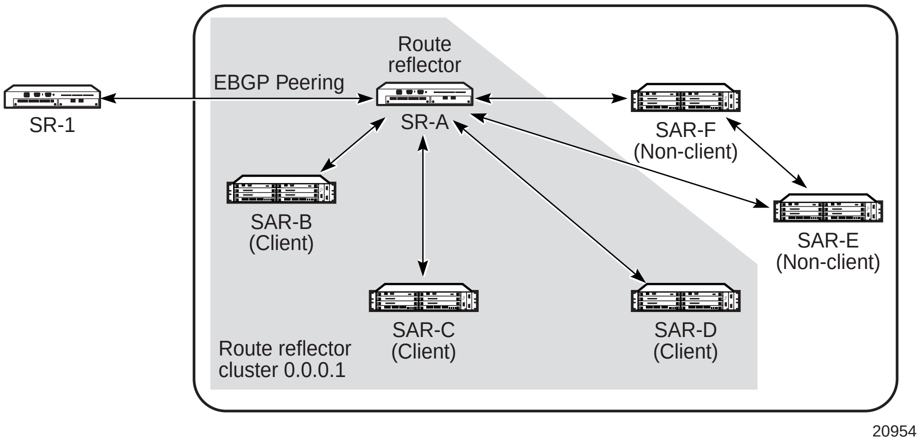

BGP speakers within an AS that are not configured as reflectors are considered to be client peers. Non-client peers are other routers in the AS. A route reflector enables communication between the clients and non-client peers. Route reflector-to-client peer configurations do not need to be fully meshed, but non-client peers need to be fully meshed within an AS.

A grouping, called a cluster, is composed of a route reflector and its client peers. A cluster ID identifies the grouping unless specific BGP peerings are configured. A cluster's clients do not share information messages with other peers outside the cluster. Multiple route reflectors can be configured within a cluster for redundancy. A router assumes the role as a route reflector by configuring the `cluster` *`cluster-id`* command. No other command is required unless you want to disable reflection to specific clients.

When a route reflector receives an advertised route, depending on the sender and neighbors (peers), it selects the best path. Routes received from an EBGP peer are advertised unmodified (to retain next-hop information) to all clients and non-client peers in the AS. Routes received from a non-client peer are advertised to all clients in the AS. Routes received from a client are advertised to all clients and non-client peers.

Sending of BGP Communities

The capability to explicitly enable or disable the sending of the BGP community attribute to BGP neighbors, other than through the use of policy statements, is supported.

This feature allows an administrator to enable or disable the sending of BGP communities to an associated peer. This feature overrides communities that are already associated with a given route or that may have been added via an export route policy. In other words, even if the export policies leave BGP communities attached to a given route, when the disable-communities feature is enabled, no BGP communities are advertised to the associated BGP peers.

Route Selection Criteria

For each prefix in the routing table, the routing protocol selects the best path. Then, the best path is compared to the next path in the list until all paths in the list are exhausted. The following parameters are used to determine the best path.

1. Routes are not considered if they are unreachable.
2. An RTM's preference is lowered as well as the hierarchy of routes from a different protocol. The lower the preference, the higher the chance of the route being the active route.
3. Routes with higher local preference have preference.

4. Routes with the shorter AS path have preference.
5. Routes with the lower origin have preference: IGP = 0, EGP = 1, INCOMPLETE = 2.
6. Routes with the lowest Multi-Exit Discriminator (MED) metric have preference.
7. Routes learned by an EBGP peer rather than those learned from an IBGP peer are preferred.
8. Routes with the lowest IGP cost to the next-hop path attribute are preferred.
9. Routes with the lowest BGP-ID are preferred.
10. Routes with shortest cluster list are preferred.
11. Routes with lowest next-hop IP address are preferred.

Command Interactions and Dependencies

This section highlights the BGP command interactions and dependencies that are important for configuration or operational maintenance of 7705 SAR routers. Topics covered in this section are:

- Changing the Autonomous System Number
- Changing the Local AS Number
- Changing the Router ID at the Configuration Level
- Hold Time and Keepalive Timer Dependencies
- Import and Export Route Policies
- Route Damping and Route Policies

Note that this information can be found in the BGP Command Reference on page 265, which provides detailed descriptions of the configuration commands.

Changing the Autonomous System Number

If the AS number is changed on a router with an active BGP instance, the new AS number will not be used until the BGP instance is restarted either by administratively disabling or enabling the BGP instance or by rebooting the system with the new configuration.

Changing the Local AS Number

Changing the local AS of an active BGP instance:

- at the global level — causes the BGP instance to restart with the new local AS number
- at the group level — causes BGP to re-establish the peer relationships with all peers in the group with the new local AS number
- at the neighbor level — causes BGP to re-establish the peer relationship with the new local AS number

Changing the Router ID at the Configuration Level

If you configure a new router ID in the `config>router-id` context, protocols are not automatically restarted with the new router ID. The next time a protocol is initialized or reinitialized, the new router ID is used. An interim period of time can occur when different protocols use different router IDs.

Hold Time and Keepalive Timer Dependencies

The BGP hold time specifies the maximum time BGP will wait between successive messages (either Keepalive or Update) from its peer, before closing the connection. This configuration parameter can be set at three levels. The most specific value is used:

- global level — applies to all peers
- group level — applies to all peers in the group
- neighbor level — only applies to the specified peer

Although the keepalive time can be user-specified, the configured keepalive timer is overridden by the value of hold time under the following circumstances.

- If the hold time specified is less than the configured keepalive time, then the operational keepalive time is set to one third of the specified hold time; the configured keepalive time is unchanged.
- If the hold time is set to zero, then the operational value of the keepalive time is set to zero; the configured keepalive time is unchanged. This means that the connection with the peer will be up permanently and no keepalive packets are sent to the peer.

If the hold time or keepalive values are changed, the changed timer values take effect when the new peering relationship is established. Changing the values causes the peerings to restart. The changed timer values are used when renegotiating the peer relationship.

Import and Export Route Policies

Import and export route policy statements are specified for BGP at the global, group, and neighbor level. Up to five unique policy statement names can be specified in the command line per level. The most specific command is applied to the peer. Defining the policy statement name is not required before being applied. Policy statements are evaluated in the order in which they are specified within the command context until the first matching policy statement is found.

The import and export policies configured at different levels are not cumulative. The most specific value is used. An import or export policy command specified at the neighbor level takes precedence over the same command specified at the group or global level. An import or export policy command specified at the group level takes precedence over the same command specified at the global level.

Route Damping and Route Policies

To prevent BGP systems from sending excessive route changes to peers, BGP route damping can be implemented. Damping can reduce the number of Update messages sent between BGP peers, to reduce the load on peers, without adversely affecting the route convergence time for stable routes.

The damping profile defined in the policy statement is applied to control route damping parameters. Route damping characteristics are specified in a route damping profile and are referenced in the action for the policy statement or in the action for a policy entry. Damping can be specified at the global, group, or neighbor level with the most specific command applied to the peer.

BGP Configuration Process Overview

Figure 9 displays the process to provision basic BGP parameters.

Figure 9: BGP Configuration and Implementation Flow

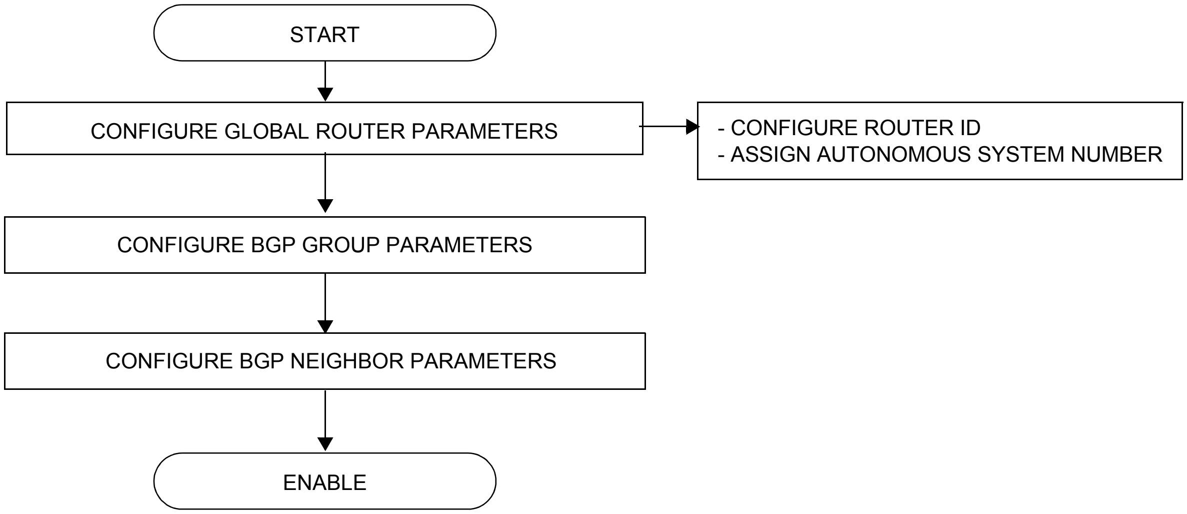
START
CONFIGURE GLOBAL ROUTER PARAMETERS
- CONFIGURE ROUTER ID
- ASSIGN AUTONOMOUS SYSTEM NUMBER
CONFIGURE BGP GROUP PARAMETERS
CONFIGURE BGP NEIGHBOR PARAMETERS
ENABLE

Configuration Notes

This section describes BGP configuration caveats.

General

- Before BGP can be configured, the router ID (a valid host address, not the MAC address default) and autonomous system global parameters must be configured.
- BGP instances must be explicitly created on each BGP peer. There are no default BGP instances on a 7705 SAR.

BGP Defaults

The following list summarizes the BGP configuration defaults.

- By default, the 7705 SAR is not assigned to an AS.
- A BGP instance is created in the administratively enabled state.
- A BGP group is created in the administratively enabled state.
- A BGP neighbor is created in the administratively enabled state.
- No BGP router ID is specified. If no BGP router ID is specified, BGP uses the router system interface address.
- The 7705 SAR OS BGP timer defaults are the values recommended in IETF drafts and RFCs (see BGP MIB Notes).
- If no import route policy statements are specified, then all BGP routes are accepted.
- If no export route policy statements specified, then all BGP routes are advertised and non-BGP routes are not advertised.

BGP MIB Notes

The 7705 SAR OS implementation of the RFC 1657 MIB variables listed in Table 36 differs from the IETF MIB specification.

Table 36: 7705 SAR OS and IETF MIB Variations

MIB Variable	Description	RFC 1657 Allowed Values	7705 SAR OS Allowed Values
bgpPeerMinASOrig-intionInterval	Time interval in seconds for the MinASOriginationInterval timer. The suggested value for this timer is 15 s.	1 to 65535	2 to 255
bgpPeerMinRouteAd-vertisementInterval	Time interval in seconds for the MinRouteAdvertisementInterval timer. The suggested value for this timer is 30 s.	1 to 65535	2 to 255

If SNMP is used to set a value of *X* to the MIB variable in Table 37, there are three possible results:

Table 37: Table 15:MIB Variable with SNMP

Condition	Result
X is within IETF MIB values and *X* is within 7705 SAR OS values	SNMP set operation does not return an error MIB variable set to *X*
X is within IETF MIB values and *X* is outside 7705 SAR OS values	SNMP set operation does not return an error MIB variable set to “nearest” 7705 SAR OS supported value (for example, 7705 SAR OS range is 2 to 255 and *X* = 65535, MIB variable will be set to 255) Log message generated
X is outside IETF MIB values and *X* is outside 7705 SAR OS values	SNMP set operation returns an error

When the value set using SNMP is within the IETF allowed values and outside the 7705 SAR OS values as specified in Table 36 and Table 37, a log message is generated. The log messages that display are similar to the following log messages.

Sample Log Message for setting bgpPeerMinASOriginationInterval to 65535

```
576 2006/11/12 19:45:48 [Snmpd] BGP-4-bgpVariableRangeViolation:
Trying to set bgpPeerMinASOrigInt to 65535 - valid range is [2-255]
- setting to 255
```

Sample Log Message for setting bgpPeerMinASOriginationInterval to 1

```
594 2006/11/12 19:48:05 [Snmpd] BGP-4-bgpVariableRangeViolation:
Trying to set bgpPeerMinASOrigInt to 1 - valid range is [2-255] -
setting to 2
```

Sample Log Message for setting bgpPeerMinRouteAdvertisementInterval to 256

```
535 2006/11/12 19:40:53 [Snmpd] BGP-4-bgpVariableRangeViolation:
Trying to set bgpPeerMinRouteAdvInt to 256 - valid range is [2-255]
- setting to 255
```

Sample Log Message for setting bgpPeerMinRouteAdvertisementInterval to 1

```
566 2006/11/12 19:44:41 [Snmpd] BGP-4-bgpVariableRangeViolation:
Trying to set bgpPeerMinRouteAdvInt to 1 - valid range is [2-255] -
setting to 2
```

Configuring BGP with CLI

This section provides information to configure BGP using the command line interface.

Topics in this section include:

BGP Configuration Overview

Preconfiguration Requirements

Before BGP can be implemented, the following entities must be configured:

- the autonomous system (AS) number for the router

  An AS number is a globally unique value that associates a router with a specific autonomous system. This number is used to exchange exterior routing information with neighboring ASs and as an identifier of the AS itself. Each router participating in BGP must have an AS number specified.

  In order to implement BGP, the AS number must be specified in the `config>router` context.
- the router ID

  The router ID is the IP address of the local router. The router ID identifies a packet's origin. The router ID must be a valid host address.

BGP Hierarchy

BGP is configured in the `config>router>bgp` context. Three hierarchical levels are included in BGP configurations:

- global level
- group level
- neighbor level

Commands and parameters configured at the global level are inherited by the group and neighbor levels, although parameters configured at the group and neighbor levels take precedence over global configurations.

Internal BGP Configurations

A BGP system is composed of ASs that share network reachability information. Network reachability information is shared with adjacent BGP systems' neighbors. Further logical groupings are established within the BGP systems that are within ASs. On the 7705 SAR, Release 3.0, BGP supports IBGP for routing information exchanges.

Internal BGP (IBGP) is used within an AS. An IBGP speaker peers to the same AS and typically does not share a subnet. Neighbors do not have to be directly connected to each other. Since IBGP neighbors are not required to be directly connected, IBGP uses the IGP path (the IP next-hop learned from the IGP) to reach an IBGP neighbor for its peering connection.

BGP Route Reflectors

In a standard BGP configuration, all BGP speakers within an AS must have a full BGP mesh to ensure that all externally learned routes are redistributed through the entire AS. IBGP speakers do not readvertise routes learned from one IBGP peer to another IBGP peer. If a network grows, scaling issues could emerge because of the full mesh configuration requirement. Route reflection circumvents the full mesh requirement but still maintains the full distribution of external routing information within an AS.

Autonomous systems using route reflection arrange BGP routers into groups called clusters. Each cluster contains at least one route reflector that is responsible for redistributing route updates to all clients. Route reflector clients do not need to maintain a full peering mesh between each other. They only require a peering to the route reflector(s) in their cluster. The route reflectors must maintain a full peering mesh between all non-clients within the AS.

Each route reflector must be assigned a cluster ID and specify which neighbors are clients and which are non-clients to determine which neighbors should receive reflected routes and which should be treated as a standard IBGP peer. Additional configuration is not required for the route reflector except for the typical BGP neighbor parameters.

Figure 10 illustrates an autonomous system with clusters.

Figure 10: Route Reflection Network Diagram Example

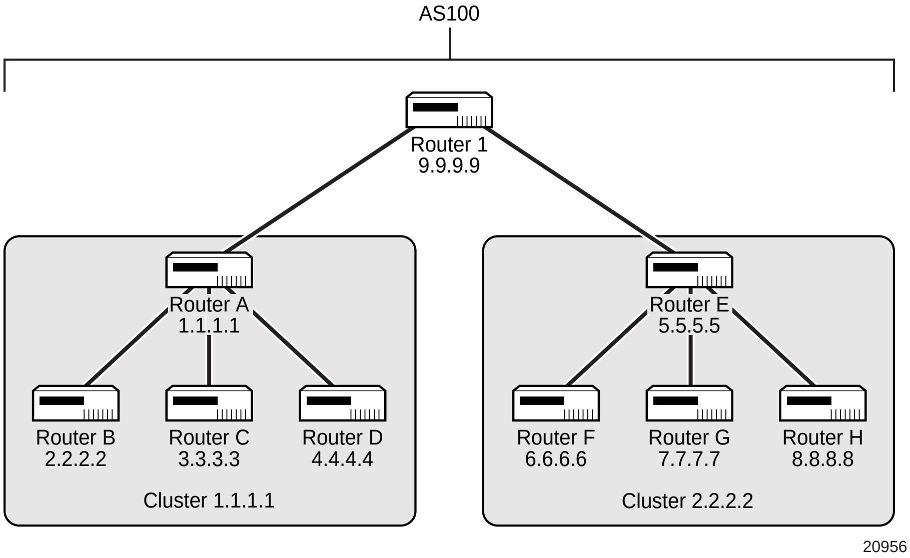

The following configuration example shows the minimum BGP configuration for routers in Cluster 1.1.1.1, shown in Figure 10.

```
ALU-A
     config router bgp
          group cluster1
               peer-as 100
               cluster 1.1.1.1
               neighbor 2.2.2.2
               exit
               neighbor 3.3.3.3
               exit
               neighbor 4.4.4.4
               exit
          exit
          group RRs
               peer-as 100
               neighbor 5.5.5.5
               exit
               neighbor 9.9.9.9
               exit
          exit
     exit
ALU-B
     config router bgp
          group cluster1
               peer-as 100
               neighbor 1.1.1.1
```

```
                exit
            exit
        exit
    ALU-C
        config router bgp
            group cluster1
                peer-as 100
                neighbor 1.1.1.1
                exit
            exit
        exit
    ALU-D
        config router bgp
            group cluster1
                peer-as 100
                neighbor 1.1.1.1
                exit
            exit
        exit
```

Basic BGP Configuration

This section provides information to configure BGP and configuration examples of common configuration tasks. The minimum BGP parameters that must be configured are:

- an autonomous system number for the router
- a router ID

> **Note:** If a new or different router ID value is entered in the BGP context, the new value takes precedence and overwrites the router-level router ID.

- a BGP peer group
- a BGP neighbor with which to peer
- a BGP peer-AS that is associated with the above peer

The BGP configuration commands have three primary configuration levels:

- global configuration (`config>router>bgp`)
- BGP group configuration (`config>router>bgp>group`)
- BGP neighbor configuration (`config>router>bgp>group>neighbor`)

Within the three levels, many of the configuration commands are repeated. For the repeated commands, the command that is most specific to the neighboring router is in effect; that is, neighbor settings have precedence over group settings, which have precedence over BGP global settings.

The following is a sample configuration that includes the parameters in the list above. The other parameters shown below are optional:

```
info
#--------------------------------------------------
echo "IP Configuration"
#--------------------------------------------------
...
     autonomous-system 200
     router-id 10.10.10.103
#--------------------------------------------------
...
#--------------------------------------------------
echo "BGP Configuration"
#--------------------------------------------------
     bgp
          graceful-restart
          exit
          cluster 0.0.0.100
          export "direct2bgp"
          router-id 10.0.0.12
          group "Group1"
               connect-retry 20
               hold-time 90
               keepalive 30
               local-preference 100
               remove-private
               peer-as 200
               neighbor 10.0.0.8
                    description "To_Router B - IBGP Peer"
                    connect-retry 20
                    hold-time 90
                    keepalive 30
                    local-address 10.0.0.12
                    passive
                    preference 99
                    peer-as 200
               exit
          exit
          group "Group2"
               connect-retry 20
               hold-time 90
               keepalive 30
               local-preference 100
               remove-private
               peer-as 200
               neighbor 10.0.3.10
                    description "To_Router C - IBGP Peer"
                    connect-retry 20
                    hold-time 90
                    keepalive 30
                    peer-as 200
               exit
          exit
          group "Group3"
               connect-retry 20
               hold-time 30
               keepalive 30
               local-preference 100
               peer-as 200
               neighbor 10.0.0.15
                    description "To_Router E - IBGP Peer"
                    connect-retry 20
                    hold-time 90
                    keepalive 30
                    local-address 10.0.0.12
                    peer-as 200
               exit
          exit
     exit
#--------------------------------------------------
....
A:ALU-48>config>router#
```

Common Configuration Tasks

This section provides a brief overview of the tasks that must be performed to configure BGP and provides the CLI commands. In order to enable BGP, one AS must be configured and at least one group must be configured that includes neighbor (system or IP address) and peering information (AS number).

BGP is configured hierarchically: the global level (applies to all peers), the group level (applies to all peers in peer-group), and the neighbor level (only applies to specified peer). By default, group members inherit the group's configuration parameters, although a parameter can be modified on a per-member basis without affecting the group-level parameters.

Many of the hierarchical BGP commands can be used at different levels. The most specific value is used. That is, a BGP group-specific command takes precedence over a global BGP command. A neighbor-specific command takes precedence over a global BGP or group-specific command.

All BGP instances must be explicitly created on each 7705 SAR. Once created, BGP is administratively enabled.

Configuration planning is essential to organize ASs and the 7705 SARs within the ASs, and to determine the internal and external BGP peering. To configure a basic autonomous system, perform the following tasks.

1. Prepare a plan detailing the autonomous systems, the 7705 SAR belonging to each group, group names, and peering connections.
2. Associate each 7705 SAR with an autonomous system number.
3. Configure each 7705 SAR with a router ID.
4. Associate each 7705 SAR with a peer group name.
5. Specify the local IP address that will be used by the group or neighbor when communicating with BGP peers.
6. Specify neighbors.
7. Specify the autonomous system number associated with each neighbor.

Creating an Autonomous System

Before BGP can be configured, the autonomous system must be configured. In BGP, routing reachability information is exchanged between autonomous systems (ASs). An AS is a group of networks that share routing information. The `autonomous-system` command associates an autonomous system number with the 7705 SAR being configured. A 7705 SAR can only belong to one AS. The `autonomous-system` command is configured in the `config>router` context.

Use the following CLI syntax to associate a 7705 SAR with an autonomous system:

CLI Syntax: `config>router# autonomous-system` *`as-number`*

The following example displays autonomous system configuration command usage:

Example: `config>router# autonomous-system 100`

The following example displays the autonomous system configuration:

```
ALU-B>config>router# info
#------------------------------------------
# IP Configuration
#------------------------------------------
     interface "system"
          address 10.10.10.104/32
     exit
     interface "to-103"
          address 10.0.0.104/24
          port 1/1/1
     exit
     autonomous-system 100
#------------------------------------------
ALU-B>config>router#
```

Configuring a Router ID

In BGP, routing information is exchanged between autonomous systems. The BGP router ID, expressed as an IP address, uniquely identifies the router. It can be set to be the same as the loopback address.

If a new or different router ID value is entered in the BGP context, the new router ID value is used instead of the router ID configured on the router level, system interface level, or inherited from the MAC address. The router-level router ID value remains intact.

A router ID can be derived by:

- defining the value in the `config>router` context, using the `router-id` command
- defining the system interface in the `config>router>interface` *`ip-int-name`* context
- inheriting the last four bytes of the MAC address
- defining the value within the BGP protocol level. The router ID can be defined in the `config>router>bgp` context, using the `router-id` command, and is only used within BGP.

When configuring a new router ID, protocols are not automatically restarted with the new router ID. The next time a protocol is initialized or reinitialized, the new router ID is used. An interim period of time can occur when different protocols use different router IDs. To force the new router ID, issue the `shutdown` and `no shutdown` commands for each protocol that uses the router ID or restart the entire router.

Use the following CLI syntax to configure the router ID:

CLI Syntax: `config>router# router-id` *`ip-address`*

The following example displays router ID configuration command usage:

Example: `config>router# router-id 10.10.10.104`

The following example displays the router ID configuration:

```
ALU-B>config>router# info
----------------------------------------------
# IP Configuration
#------------------------------------------
     interface "system"
          address 10.10.10.104/32
     exit
     interface "to-103"
          address 10.0.0.104/24
          port 1/1/1
     exit
     autonomous-system 100
     router-id 10.10.10.104
#------------------------------------------
...
ALU-B>config>router#
```

BGP Components

Use the CLI syntax displayed below to configure the following BGP attributes:

- Configuring BGP
- Configuring Group Attributes
- Configuring Neighbor Attributes
- Configuring Route Reflection

Configuring BGP

Once the BGP protocol instance is created, the `no shutdown` command is not required since BGP is administratively enabled upon creation. Minimally, to enable BGP on a router, you must associate an autonomous system number for the router, have a preconfigured router ID or system interface, create a peer group, neighbor, and associate a peer AS number. There are no default groups or neighbors. Each group and neighbor must be explicitly configured.

All parameters configured for BGP are applied to the group and are inherited by each peer, but a group parameter can be overridden on a specific basis. BGP command hierarchy consists of three levels:

- the global level
- the group level
- the neighbor level

For example:

CLI Syntax:
```
config>router# bgp       (global level)
    group                (group level)
       neighbor          (neighbor level)
```

Note: Careful planning is essential to implement commands that can affect the behavior of global, group, and neighbor levels. Because the BGP commands are hierarchical, analyze the values that can disable features on a particular level.

The following example displays the basic BGP configuration:

```
ALU-B>config>router# info
#------------------------------------------
# BGP Configuration
#------------------------------------------
#------------------------------------------
# BGP
#------------------------------------------
     bgp
     exit
#------------------------------------------
ALU-B>config>router#
```

Configuring Group Attributes

A group is a collection of related BGP peers. The group name should be a descriptive name for the group. Follow your group, name, and ID naming conventions for consistency and to help when troubleshooting faults. All parameters configured for a peer group are applied to the group and are inherited by each peer (neighbor), but a group parameter can be overridden on a specific neighbor-level basis.

The following example displays group configuration command usage:

Example:

```
config>router# bgp
config>router>bgp# group headquarters1
config>router>bgp>group# description “HQ execs”
config>router>bgp>group# local-address 10.0.0.104
config>router>bgp>group# disable-communities
config>router>bgp>group# exit
```

The following example displays the BGP group configuration:

```
ALU-B>config>router>bgp# info
----------------------------------------------
...
     group "headquarters1"
          description "HQ execs"
          local-address 10.0.0.104
          disable-communities standard extended
          exit
     exit
...
----------------------------------------------
ALU-B>config>router>bgp#
```

Configuring Neighbor Attributes

After you create a group name and assign options, add neighbors within the same autonomous system to create IBGP connections. All parameters configured for the peer group level are applied to each neighbor, but a group parameter can be overridden on a specific neighbor basis.

The following example displays neighbor configuration command usage:

Example:

```
config>router# bgp
config>router>bgp# group headquarters1
config>router>bgp>group# neighbor 10.0.0.5
config>router>bgp>group# peer-as 100
config>router>bgp>group# passive
config>router>bgp>group# exit
config>router>bgp>group# neighbor 10.10.10.103
config>router>bgp>group# peer-as 100
config>router>bgp>group# exit
config>router>bgp>group# neighbor 17.5.0.2
config>router>bgp>group>neighbor$ hold-time 90
config>router>bgp>group>neighbor$ keepalive 30
config>router>bgp>group>neighbor$ min-as-origination 15
config>router>bgp>group>neighbor$ local-preference 170
config>router>bgp>group>neighbor$ peer-as 100
config>router>bgp>group>neighbor$ exit
config>router>bgp>group# neighbor 17.5.1.2
config>router>bgp>group>neighbor$ hold-time 90
config>router>bgp>group>neighbor$ keepalive 30
config>router>bgp>group>neighbor$ min-as-origination 15
config>router>bgp>group>neighbor$ local-preference 100
config>router>bgp>group>neighbor$ min-route-advertisement
30
config>router>bgp>group>neighbor$ preference 170
config>router>bgp>group>neighbor$ peer-as 100
config>router>bgp>group>neighbor$ exit
config>router>bgp>group# info
```

The following example displays neighbors configured in group “headquarters1”.

```
ALU-B>config>router>bgp# info
----------------------------------------------
...
     group "headquarters1"
          description "HQ execs"
          local-address 10.0.0.104
          disable-communities standard extended
          neighbor 10.0.0.5
               passive
               peer-as 100
          exit
          neighbor 10.0.0.106
               peer-as 100
          exit
          neighbor 17.5.0.2
               hold-time 90
               keepalive 30
               min-as-origination 15
               local-preference 170
               peer-as 100
          exit
          neighbor 17.5.1.2
               hold-time 90
               keepalive 30
               min-as-origination 15
               local-preference 100
               min-route-advertisement 30
               preference 170
               peer-as 100
          exit
     exit
...
----------------------------------------------
ALU-B>config>router>bgp#
```

Configuring Route Reflection

Route reflection can be implemented in autonomous systems with a large internal BGP mesh to reduce the number of IBGP sessions required. One or more routers can be selected to act as focal points for internal BGP sessions. Several BGP speaking routers can peer with a route reflector. A route reflector forms peer connections to other route reflectors. A router assumes the role as a route reflector by configuring the `cluster` *`cluster-id`* command. No other command is required unless you want to disable reflection to specific peers.

If you configure the `cluster` command at the global level, then all subordinate groups and neighbors are members of the cluster. The route reflector cluster ID is expressed in dotted-decimal notation. The ID should be a significant topology-specific value. No other command is required unless you want to disable reflection to specific peers.

If a route reflector client is fully meshed, the `disable-client-reflect` command can be enabled to stop the route reflector from reflecting redundant route updates to a client.

The following example displays route reflection configuration command usage:

Example:

```
config>router# bgp
config>router>bgp# cluster 0.0.0.100
config>router>bgp# group “Santa Clara”
config>router>bgp>group$ local-address 10.0.0.103
config>router>bgp>group# neighbor 10.0.0.91
config>router>bgp>group>neighbor$ peer-as 100
config>router>bgp>group>neighbor# exit
config>router>bgp>group# neighbor 10.0.0.92
config>router>bgp>group>neighbor$ peer-as 100
config>router>bgp>group>neighbor# exit
config>router>bgp>group# neighbor 10.0.0.93
config>router>bgp>group>neighbor$ disable-client-refl
config>router>bgp>group>neighbor# peer-as 100
config>router>bgp>group>neighbor# exit
```

The following example displays a route reflection configuration:

```
ALU-B>config>router>bgp# info
----------------------------------------------
     cluster 0.0.0.100
     group "Santa Clara"
          local-address 10.0.0.103
          neighbor 10.0.0.91
               peer-as 100
          exit
          neighbor 10.0.0.92
               peer-as 100
          exit
          neighbor 10.0.0.93
               disable-client-reflect
               peer-as 100
          exit
     exit
----------------------------------------------
ALU-B>config>router>bgp#
```

BGP Configuration Management Tasks

This section discusses the following BGP configuration management tasks:

- Modifying an AS Number
- Modifying the BGP Router ID
- Modifying the Router-Level Router ID
- Deleting a Neighbor
- Deleting Groups
- Editing BGP Parameters

Modifying an AS Number

You can modify an AS number on a 7705 SAR but the new AS number will not be used until the BGP instance is restarted either by administratively disabling or enabling the BGP instance, or by rebooting the system with the new configuration.

Since the AS number is defined in the `config>router` context, not in the BGP configuration context, the BGP instance is not aware of the change. Re-examine the plan detailing the autonomous systems, the 7705 SARs belonging to each group, group names, and peering connections.

➡ **Note:** Changing an AS number on a 7705 SAR could cause configuration inconsistencies if associated peer-as values are not also modified as required. At the group and neighbor levels, BGP will re-establish the peer relationships with all peers in the group with the new AS number.

Use the following CLI syntax to change an autonomous system number:

CLI Syntax:

```
config>router# autonomous-system as-number
```

CLI Syntax:

```
config>router# bgp
   group name
      neighbor ip-addr
         peer-as as-number
```

Example:

```
config>router# autonomous-system 400
config>router# bgp
config>router>bgp# group headquarters1
config>router>bgp>group# neighbor 10.10.10.103
config>router>bgp>group# peer-as 400
config>router>bgp>group# exit
```

Modifying the BGP Router ID

Changing the router ID number in the BGP context causes the new value to overwrite the router ID configured on the router level, system interface level, or the value inherited from the MAC address.

Note: Changing the router ID on a router could cause configuration inconsistencies if associated values are not also modified.

When configuring a new router ID, protocols are not automatically restarted with the new router ID. The next time a protocol is initialized or reinitialized, the new router ID is used. An interim period of time can occur when different protocols use different router IDs. To force the new router ID, issue the `shutdown` and `no shutdown` commands for each protocol that uses the router ID or restart the entire router.

Example:

```
config>router>bgp# router-id 10.0.0.104
config>router>bgp# shutdown
config>router>bgp# router-id 10.0.0.123
config>router>bgp# no shutdown
```

This example displays the BGP configuration with the BGP router ID specified:

```
ALU-B>config>router>bgp# info detail
----------------------------------------------
     no shutdown
     no description
     no always-compare-med
     ibgp-multipath load-balance
. . .
     router-id 10.0.0.123
----------------------------------------------
ALU-B>config>router>bgp#
```

Modifying the Router-Level Router ID

Changing the router ID number in the `config>router` context causes the new value to overwrite the router ID configured on the protocol level, system interface level, or the value inherited from the MAC address.

Note: Changing the router ID on a router could cause configuration inconsistencies if associated values are not also modified.

When configuring a new router ID, protocols are not automatically restarted with the new router ID. The next time a protocol is initialized or reinitialized, the new router ID is used. An interim period of time can occur when different protocols use different router IDs. To force the new router ID, issue the `shutdown` and `no shutdown` commands for each protocol that uses the router ID or restart the entire router.

Use the following CLI syntax to change a router ID:

CLI Syntax: `config>router# router-id` *`router-id`*

Example:

```
config>router# router-id 10.10.10.104
config>router# no shutdown
config>router>bgp# shutdown
config>router>bgp# no shutdown
```

The following example displays the router ID configuration:

```
ALU-A>config>router# info
#------------------------------------------
# IP Configuration
#------------------------------------------
     interface "system"
          address 10.10.10.104/32
     exit
     interface "to-103"
          address 10.0.0.104/24
          port 1/1/1
     exit
     autonomous-system 100
     router-id 10.10.10.104
#------------------------------------------
ALU-B>config>router#
```

Deleting a Neighbor

In order to delete a neighbor, you must shut down the neighbor before issuing the `no neighbor` *`ip-addr`* command.

Use the following CLI syntax to delete a neighbor:

CLI Syntax:

```
config>router# bgp
     group name
        no neighbor ip-address
        shutdown
           no peer-as as-number
           shutdown
```

Example:

```
config>router# bgp
config>router>bgp# group headquarters1
config>router>bgp>group# neighbor 10.0.0.103
config>router>bgp>group>neighbor# shutdown
config>router>bgp>group>neighbor# exit
config>router>bgp>group# no neighbor 10.0.0.103
```

The following example displays the “headquarters1” configuration with the neighbor 10.0.0.103 removed.

```
ALU-B>config>router>bgp# info
----------------------------------------------
     group "headquarters1"
          description “HQ execs”
               local-address 10.0.0.104
               neighbor 10.0.0.5
                    passive
                    peer-as 300
               exit
          exit
----------------------------------------------
ALU-B>config>router>bgp#
```

Deleting Groups

In order to delete a group, the neighbor configurations must be shut down first. After each neighbor is shut down, you must shut down the group before issuing the `no group` *`name`* command.

Use the following CLI syntax to shut down a peer and neighbor and then delete a group:

CLI Syntax:

```
config>router# bgp
     no group name
     shutdown
        no neighbor ip-address
     shutdown
shutdown
```

Example:

```
config>router# bgp
config>router>bgp# group headquarters1
config>router>bgp>group# neighbor 10.0.0.105
config>router>bgp>group>neighbor# shutdown
config>router>bgp>group>neighbor# exit
config>router>bgp>group# neighbor 10.0.0.103
config>router>bgp>group# shutdown
config>router>bgp>group# exit
config>router>bgp# no headquarters1
```

If you try to delete the group without shutting down the peer-group, the following message appears:

```
ALU-B>config>router>bgp# no group headquarters1
MINOR: CLI BGP Peer Group should be shutdown before deleted. BGP
Peer Group not deleted.
```

Editing BGP Parameters

You can change existing BGP parameters in the CLI. The changes are applied immediately.

CLI Syntax:
```
config>router# bgp
     group name
     . . .
        neighbor ip-address
     . . .
```

Example:
```
config>router# bgp
```

Refer to BGP Components on page 253 for a complete list of BGP parameters.

BGP Command Reference

Command Hierarchies

- Configuration Commands
  → Global BGP Commands
  → Group BGP Commands
  → Neighbor BGP Commands
  → Other BGP-Related Commands
- Show Commands
- Clear Commands
- Debug Commands

Configuration Commands

Global BGP Commands

```
config
    — router [router-name]
            — [no] bgp
                    — [no] advertise-inactive
                    — [no] aggregator-id-zero
                    — [no] as-path-ignore
                    — authentication-key [authentication-key | hash-key] [hash | hash2]
                    — no authentication-key
                    — [no] bfd-enable
                    — cluster cluster-id
                    — no cluster
                    — connect-retry seconds
                    — no connect-retry
                    — [no] damping
                    — description description-string
                    — no description
                    — [no] disable-client-reflect
                    — disable-communities [standard] [extended]
                    — [no] disable-communities
                    — [no] disable-fast-external-failover
                    — [no] enable-peer-tracking
                    — export policy-name [policy-name…(up to 5 max)]
                    — no export [policy-name]
                    — family [ipv4] [vpn-ipv4]
                    — no family
                    — [no] graceful-restart
                            — stale-routes-time time
                            — no stale-routes-time
                    — [no] group name
                    — hold-time seconds
                    — no hold-time
                    — [no] ibgp-multipath
                    — import policy-name [policy-name …(up to 5 max)]
                    — no import [policy-name]
                    — keepalive seconds
                    — no keepalive
                    — local-as as-number [private]
                    — no local-as
                    — local-preference local-preference
                    — no local-preference
                    — loop-detect {drop-peer | discard-route | ignore-loop | off}
                    — no loop-detect
                    — min-as-origination seconds
                    — no min-as-origination
                    — min-route-advertisement seconds
                    — no min-route-advertisement
                    — multipath integer
                    — no multipath
                    — [no] outbound-route-filtering
                            — [no] extended-community
                                    — [no] accept-orf
                                    — send-orf [comm-id...(up to 32 max)]
                                    — no send-orf comm-id
                    — preference preference
                    — no preference
                    — [no] rapid-withdrawal
                    — [no] remove-private
                    — route-target-list comm-id [comm-id...(up to 15 max)]
                    — no route-target-list [comm-id]
                    — router-id ip-address
                    — no router-id
                    — [no] shutdown
                    — [no] vpn-apply-export
                    — [no] vpn-apply-import
```

Group BGP Commands

```
config
    — router [router-name]
            — [no] bgp
                    — [no] group name
                                    — [no] advertise-inactive
                                    — [no] aggregator-id-zero
                                    — authentication-key [authentication-key | hash-key] [hash |
                                      hash2]
                                    — no authentication-key
                                    — [no] bfd-enable
                                    — cluster cluster-id
                                    — no cluster
                                    — connect-retry seconds
                                    — no connect-retry
                                    — [no] damping
                                    — description description-string
                                    — no description
                                    — [no] disable-client-reflect
                                    — disable-communities [standard] [extended]
                                    — no disable-communities
                                    — [no] disable-fast-external-failover
                                    — [no] enable-peer-tracking
                                    — export policy-name [policy-name…(up to 5 max)]
                                    — no export [policy-name]
                                    — family [ipv4] [vpn-ipv4]
                                    — no family
                                    — [no] graceful-restart
                                            — stale-routes-time time
                                            — no stale-routes-time
                                    — hold-time seconds
                                    — no hold-time
                                    — import policy-name [policy-name …(up to 5 max)]
                                    — no import [policy-name]
```

```
                                        — keepalive seconds
                                        — no keepalive
                                        — local-address ip-address
                                        — no local-address
                                        — local-as as-number [private]
                                        — no local-as
                                        — local-preference local-preference
                                        — no local-preference
                                        — loop-detect {drop-peer | discard-route | ignore-loop | off}
                                        — no loop-detect
                                        — min-as-origination seconds
                                        — no min-as-origination
                                        — min-route-advertisement seconds
                                        — no min-route-advertisement
                                        — [no] neighbor ip-address
                                        — [no] next-hop-self
                                        — [no] outbound-route-filtering
                                                — [no] extended-community
                                                    — [no] accept-orf
                                                    — send-orf [comm-id...(up to 32 max)]
                                                    — no send-orf comm-id
                                        — [no] passive
                                        — peer-as as-number
                                        — no peer-as
                                        — preference preference
                                        — no preference
                                        — prefix-limit limit
                                        — no prefix-limit
                                        — [no] remove-private
                                        — [no] shutdown
                                        — [no] vpn-apply-export
                                        — [no] vpn-apply-import
```

Neighbor BGP Commands

```
config
    — router [router-name]
            — [no] bgp
                    — [no] group name
                            — [no] neighbor ip-address
                                    — [no] advertise-inactive
                                    — [no] aggregator-id-zero
                                    — authentication-key [authentication-key | hash-key] [hash |
                                      hash2]
                                    — no authentication-key
                                    — [no] bfd-enable
                                    — cluster cluster-id
                                    — no cluster
                                    — connect-retry seconds
                                    — no connect-retry
                                    — [no] damping
                                    — description description-string
— no description
— [no] disable-client-reflect
— disable-communities [standard] [extended]
— no disable-communities
— [no] disable-fast-external-failover
— [no] enable-peer-tracking
— export policy-name [policy-name…(up to 5 max)]
— no export [policy-name]
— family [ipv4] [vpn-ipv4]
— no family
— [no] graceful-restart
     — stale-routes-time time
     — no stale-routes-time
— hold-time seconds
— no hold-time
— import policy-name [policy-name …(up to 5 max)]
— no import [policy-name]
— keepalive seconds
— no keepalive
— local-address ip-address
— no local-address
— local-as as-number [private]
— no local-as
— local-preference local-preference
— no local-preference
— loop-detect {drop-peer | discard-route | ignore-loop | off}
— no loop-detect
— min-as-origination seconds
— no min-as-origination
— min-route-advertisement seconds
— no min-route-advertisement
— [no] next-hop-self
— [no] outbound-route-filtering
     — [no] extended-community
       — [no] accept-orf
       — send-orf [comm-id...(up to 32 max)]
       — no send-orf comm-id
— [no] passive
— peer-as as-number
— no peer-as
— preference preference
— no preference
— prefix-limit limit
— no prefix-limit
— [no] remove-private
— [no] shutdown
— [no] vpn-apply-export
— [no] vpn-apply-import
```

Other BGP-Related Commands

```
config
     — router [router-name]
               — aggregate ip-prefix/ip-prefix-length [summary-only]
               — autonomous-system as-number
               — no autonomous-system
               — router-id ip-address
               — no router-id
```

Show Commands

```
show
     — router [router-instance]
               — bgp
                         — damping [damp-type] [detail]
                         — damping [ip-prefix | prefix-length] [detail]
                         — group [name] [detail]
                         — neighbor [ip-address [[family] filter1 [brief]]]
                         — neighbor [as-number [[family] filter2]]
                         — neighbor ip-address orf [filter3]
                         — neighbor ip-address graceful-restart
                         — next-hop [family] [ip-address] [detail]
                         — paths
                         — routes [family] [brief]
                         — routes [family] prefix [detail | longer | hunt [brief]]
                         — routes [family] community comm-id
                         — routes [family] aspath-regex reg-ex
                         — summary [all]
                         — summary [family family] [neighbor ip-address]
```

Clear Commands

```
clear
    — router
        — bgp
            — damping [{prefix/ip-prefix-length] [neighbor ip-address]} | {group name}]
            — flap-statistics [{prefix/mask [neighbor ip-address] | [group group-name] | [regex
              reg-exp | policy policy-name}]
            — neighbor {ip-address | as as-number | external | all} [soft | soft-inbound]
            — neighbor {ip-address | as as-number | external | all} statistics
            — neighbor ip-address end-of-rib
            — protocol
```

Debug Commands

```
debug
    — router
        — bgp
            — events [neighbor ip-address | group name]
            — no events
            — graceful-restart [neighbor ip-address | group name]
            — no graceful-restart
            — keepalive [neighbor ip-address | group name]
            — no keepalive
            — notification [neighbor ip-address | group name]
            — no notification
            — open [neighbor ip-address | group name]
            — no open
            — [no] outbound-route-filtering
            — packets [neighbor ip-address | group name]
            — no packets
            — route-refresh [neighbor ip-address | group name]
            — no route-refresh
            — rtm [neighbor ip-address | group name]
            — no rtm
            — socket [neighbor ip-address | group name]
            — no socket
            — timers [neighbor ip-address | group name]
            — no timers
            — update [neighbor ip-address | group name]
            — no update
```

Command Descriptions

Configuration Commands

Configuration Commands

bgp

Syntax [**no**] **bgp**

Context config>router

Description This command creates the BGP protocol instance and BGP configuration context. BGP is administratively enabled upon creation.

The **no** form of the command deletes the BGP protocol instance and removes all configuration parameters for the BGP instance. BGP must be shut down before deleting the BGP instance. An error occurs if BGP is not shut down first.

advertise-inactive

Syntax [**no**] **advertise-inactive**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command enables the advertising of inactive BGP routes to other BGP peers. By default, BGP only advertises BGP routes to other BGP peers if a given BGP route is chosen by the route table manager as the most preferred route within the system and is active in the forwarding plane. This command allows system administrators to advertise a BGP route even though it is not the most preferred route within the system for a given destination.

The **no** form of the command disables the advertising of inactive BGP routers to other BGP peers.

Default **no advertise-inactive**

aggregator-id-zero

Syntax [**no**] **aggregator-id-zero**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command is used to set the router ID in the BGP aggregator path attribute to zero when BGP aggregates routes. This prevents different routers within an AS from creating aggregate routes that contain different AS paths.

When BGP is aggregating routes, it adds the aggregator path attribute to the BGP Update messages. By default, BGP adds the AS number and router ID to the aggregator path attribute.

When this command is enabled, BGP adds the router ID to the aggregator path attribute. This command is used at the group level to revert to the value defined under the global level, and this command is used at the neighbor level to revert to the value defined under the group level.

The **no** form of the command used at the global level reverts to the default, where BGP adds the AS number and router ID to the aggregator path attribute.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **no aggregator-id-zero** — BGP adds the AS number and router ID to the aggregator path attribute

as-path-ignore

Syntax [**no**] **as-path-ignore**

Context config>router>bgp

Description This command specifies whether the AS path is used to determine the best BGP route.

If this option is present, the AS paths of incoming routes are not used in the route selection process.

The **no** form of the command removes the parameter from the configuration.

Default **no as-path-ignore**

authentication-key

Syntax **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
no authentication-key

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures the BGP authentication key.

Authentication is performed between neighboring routers before setting up the BGP session by verifying the password. Authentication is performed using the MD5 message-based digest.

The authentication key can be any combination of ASCII characters up to 255 characters long.

The **no** form of the command reverts to the default value.

Default **MD5 Authentication is disabled by default**

Parameters *authentication-key* — the authentication key. The key can be any combination of ASCII characters up to 255 characters in length (unencrypted). If spaces are used in the string, the entire string must be enclosed in quotation marks (“ ”).

hash-key — the hash key. The key can be any combination of ASCII characters up to 342 characters in length (encrypted). If spaces are used in the string, the entire string must be enclosed in quotation marks (“ ”). This is useful when a user must configure the parameter, but for security purposes, the actual unencrypted key value is not provided.

hash — specifies that the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2 — specifies that the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

bfd-enable

Syntax [**no**] **bfd-enable**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command enables the use of bidirectional forwarding (BFD) to control the state of the associated protocol interface. By enabling BFD on a given protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for BFD are set via the BFD command under the IP interface.

The **no** form of this command removes BFD from the associated IGP/BGP protocol adjacency.

Default **no bfd-enable**

cluster

Syntax **cluster** *cluster-id*
no cluster

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures the cluster ID for a route reflector server.

Route reflectors are used to reduce the number of IBGP sessions required within an AS. Normally, all BGP speakers within an AS must have a BGP peering with every other BGP speaker in the AS. A route reflector and its clients form a cluster. Peers that are not part of the cluster are considered to be non-clients.

When a route reflector receives a route, it must first select the best path from all the paths received. If the route was received from a non-client peer, then the route reflector sends the route to all clients in the cluster. If the route came from a client peer, the route reflector sends the route to all non-client peers and to all client peers except the originator.

For redundancy, a cluster can have multiple route reflectors.

The **no** form of the command deletes the cluster ID and effectively disables route reflection for the given group.

Default **no cluster**

Parameters *cluster-id* — the route reflector cluster ID, expressed in dotted-decimal notation

Values any 32-bit number in dotted-decimal notation. (0.0.0.1 to 255.255.255.255)

connect-retry

Syntax **connect-retry** *seconds*
no connect-retry

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures the BGP connect retry timer value in seconds. When this timer expires, BGP tries to reconnect to the configured peer. This configuration parameter can be set at three levels: global level (applies to all peers), peer-group level (applies to all peers in group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of the command used at the global level reverts to the default value.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **120 s**

Parameters *seconds* — the BGP connect retry timer value, in seconds, expressed as a decimal integer

Values 1 to 65535

damping

Syntax [**no**] **damping**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command enables BGP route damping for learned routes that are defined within the route policy. Use damping to reduce the number of Update messages sent between BGP peers and reduce the load on peers without affecting the route convergence time for stable routes. Damping parameters are set via route policy definitions.

The **no** form of the command used at the global level reverts route damping.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

When damping is enabled and the route policy does not specify a damping profile, the default damping profile is used. This profile is always present and consists of the following parameters:

• half-life:	15 min
• max-suppress:	60 min
• suppress-threshold:	3000
• reuse-threshold:	750

Default **no damping**

description

Syntax **description** *description-string*
no description

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command creates a text description stored in the configuration file for a configuration context.

The **no** form of the command removes the description string from the context.

Default **No description is associated with the configuration context**

Parameters *string* — the description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

disable-client-reflect

Syntax [**no**] **disable-client-reflect**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command disables the reflection of routes by the route reflector to the clients in a specific group or neighbor.

This command only disables the reflection of routes from other client peers. Routes learned from non-client peers are still reflected to all clients.

The **no** form re-enables client reflection of routes.

Default **no disable-client-reflect**

disable-communities

Syntax **disable-communities** [**standard**] [**extended**]
no disable-communities

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures BGP to disable sending communities.

Parameters **standard —** specifies standard communities that existed before VPRNs or RFC 2547

extended — specifies BGP communities that were expanded after the concept of RFC 2547 was introduced, to include handling the VRF target

disable-fast-external-failover

Syntax [**no**] **disable-fast-external-failover**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures BGP fast external failover.

For EBGP neighbors, this feature controls whether the router should drop an EBGP session immediately upon an interface-down event, or whether the BGP session should be kept up until the hold-time expires.

When fast external failover is disabled, the EBGP session stays up until the hold-time expires or the interface comes back up. If the BGP routes become unreachable as a result of the down IP interface, BGP withdraws the unavailable route immediately from other peers.

enable-peer-tracking

Syntax [**no**] **enable-peer-tracking**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command enables BGP peer tracking. BGP peer tracking allows a BGP peer to be dropped immediately if the route used to resolve the BGP peer address is removed from the IP routing table and there is no alternative available. The BGP peer will not wait for the hold timer to expire; therefore, the BGP reconvergence process is accelerated.

The **no** form of the command disables peer tracking.

Default **no enable-peer-tracking**

export

Syntax **export** *policy-name* [*policy-name…*(up to 5 max)]
no export [*policy-name*]

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command specifies the export route policy used to determine which routes are advertised to peers. Route policies are configured in the **config>router>policy-options** context. Refer to the section on “Route Policy” in the 7705 SAR OS Router Configuration Guide.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific level is used.

When multiple policy names are specified, the policies are evaluated in the order in which they are specified. A maximum of five (5) policy names can be configured. The first policy that matches is applied.

When multiple export commands are issued, the last command entered overrides the previous command.

When no export policies are specified, BGP routes are advertised and non-BGP routes are not advertised (by default).

The **no** form of the command removes the policy association with the BGP instance. To remove association of all policies, use the **no export** command without arguments.

Default **no export**

Parameters *policy-name* — the route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

family

Syntax **family** [**ipv4**] [**vpn-ipv4**]
no family

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command specifies the address family or families to be supported over BGP peerings in the base router. This command is additive, so issuing the **family** command adds the specified address family to the list.

The **no** form of the command removes the specified address family from the associated BGP peerings. If an address family is not specified, the supported address family is reset to the default.

Default **ipv4**

Parameters **ipv4 —** supports IPv4 routing information

vpn-ipv4 — exchanges IPv4 VPN routing information

graceful-restart

Syntax [**no**] **graceful-restart**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command enables graceful restart for BGP. If the control plane of a GR-capable router fails, the neighboring routers (GR helpers) temporarily preserve neighbor information, so packets continue to be forwarded through the failed GR router using the last known routes. The helper state remains until the peer completes its restart or exits if the GR timer value is exceeded.

The **no** form of the command disables graceful restart and removes all graceful restart configurations in the BGP instance.

Default **no graceful-restart**

stale-routes-time

Syntax **stale-routes-time** *time*
no stale-routes-time

Context config>router>bgp>graceful-restart
config>router>bgp>group>graceful-restart
config>router>bgp>group>neighbor>graceful-restart

Description This command configures the maximum amount of time in seconds that stale routes should be maintained after a graceful restart is initiated.

The **no** form of the command resets the stale routes time back to the default value.

Default **360 s**

Parameters *time* — the amount of time that stale routes should be maintained after a graceful restart is initiated

Values 1 to 3600 s

group

Syntax [**no**] **group** *name*

Context config>router>bgp

Description This command creates a context to configure a BGP peer group.

The **no** form of the command deletes the specified peer group and all configurations associated with the peer group. The group must be shut down before it can be deleted.

Default **no group**

Parameters *name* — the peer group name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

hold-time

Syntax **hold-time** *seconds*
no hold-time

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures the BGP hold time, expressed in seconds.

The BGP hold time specifies the maximum time BGP waits between successive messages (either Keepalive or Update) from its peer, before closing the connection. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in group) or neighbor level (only applies to specified peer). The most specific value is used.

Even though the 7705 SAR OS implementation allows setting the **keepalive** time separately, the configured **keepalive** timer is overridden by the **hold-time** value under the following circumstances.

- If the specified **hold-time** is less than the configured **keepalive** time, then the operational **keepalive** time is set to a third of the **hold-time**; the configured **keepalive** time is not changed.
- If the **hold-time** is set to 0, then the operational value of the **keepalive** time is set to 0; the configured **keepalive** time is not changed. This means that the connection with the peer is up permanently and no keepalive packets are sent to the peer.

The **no** form of the command used at the global level reverts to the default value.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **90 s**

Parameters *seconds* — the hold-time, in seconds, expressed as a decimal integer. A value of 0 indicates the connection to the peer is permanently up.

Values 0, 3 to 65535

ibgp-multipath

Syntax [**no**] **ibgp-multipath**

Context config>router>bgp

Description This command enables IBGP multipath load balancing when adding BGP routes to the route table if the route resolving the BGP next-hop offers multiple next-hops.

The **no** form of the command disables the IBGP multipath load balancing feature.

Default **no ibgp-multipath**

import

Syntax **import** *policy-name* [*policy-name…*(up to 5 max)]
no import [*policy-name*]

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command specifies the import route policy to be used to determine which routes are accepted from peers. Route policies are configured in the **config>router>policy-options** context. Refer to the section on "Route Policy" in the 7705 SAR OS Router Configuration Guide.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific level is used.

When multiple policy names are specified, the policies are evaluated in the order in which they are specified. A maximum of five (5) policy names can be specified. The first policy that matches is applied.

When multiple **import** commands are issued, the last command entered will override the previous command.

When an import policy is not specified, BGP routes are accepted by default.

The **no** form of the command removes the policy association with the BGP instance. To remove association of all policies, use **no import** without arguments.

Default **no import**

Parameters *policy-name* — the route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

keepalive

Syntax **keepalive** *seconds*
no keepalive

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures the BGP keepalive timer. A Keepalive message is sent every time this timer expires.

The **keepalive** parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used. The **keepalive** value is generally one-third of the **hold-time** interval. Even though the 7705 SAR OS implementation allows the **keepalive** value and the **hold-time** interval to be independently set, under the following circumstances, the configured **keepalive** value is overridden by the **hold-time** value.

- If the specified **keepalive** value is greater than the configured **hold-time**, then the specified value is ignored, and the **keepalive** value is set to one third of the current **hold-time** value.
- If the specified **hold-time** interval is less than the configured **keepalive** value, then the **keepal ive** value is reset to one third of the specified **hold-time** interval.
- If the **hold-time** interval is set to 0, then the configured value of the **keepalive** value is ignored. This means that the connection with the peer is up permanently and no **keepalive** packets are sent to the peer.

The **no** form of the command used at the global level reverts to the default value.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **30 s**

Parameters *seconds* — the keepalive timer, in seconds, expressed as a decimal integer

Values 0 to 21845

local-address

Syntax **local-address** *ip-address*
no local-address

Context config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures the local IP address used by the group or neighbor when communicating with BGP peers.

Outgoing connections use the **local-address** as the source of the TCP connection when initiating connections with a peer.

When a local address is not specified, the 7705 SAR OS uses the system IP address when communicating with IBGP peers and uses the interface address for directly connected EBGP peers. This command is used at the neighbor level to revert to the value defined under the group level.

The **no** form of the command removes the configured local address for BGP.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **no local-address**

Parameters *ip-address* — the local address, expressed in dotted-decimal notation. The allowed value is a valid routable IP address on the router, either an interface or system IP address.

Values ipv4-address: a.b.c.d (host bits must be 0)

local-as

Syntax **local-as** *as-number* [**private**]
no local-as

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures a BGP virtual autonomous system (AS) number.

In addition to the AS number configured for BGP in the **config>router>autonomous-system** context, a virtual (local) AS number is configured. The virtual AS number is added to the as-path message before the router's AS number makes the virtual AS the second AS in the as-path.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). By specifying this parameter at each neighbor level, it is possible to have a separate AS number per EBGP session.

When a command is entered multiple times for the same AS, the last command entered is used in the configuration. The **private** attribute can be added or removed dynamically by reissuing the command.

Changing the local AS at the global level in an active BGP instance causes the BGP instance to restart with the new local AS number.

Changing the local AS at the global level in an active BGP instance causes BGP to re-establish the peer relationships with all peers in the group with the new local AS number.

Changing the local AS at the neighbor level in an active BGP instance causes BGP to re-establish the peer relationship with the new local AS number.

This is an optional command and can be used in the following situation.

Example: Provider router P is moved from AS1 to AS2. The customer router that is connected to P, however, is configured to belong to AS1. To avoid reconfiguring the customer router, the **local-as** value on router P can be set to AS1. Thus, router P adds AS1 to the as-path message for routes it advertises to the customer router.

The **no** form of the command used at the global level will remove any virtual AS number configured.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **no local-as**

Parameters *as-number* — the virtual autonomous system number expressed as a decimal integer

Values 1 to 65535

private **—** specifies that the local-as is hidden in paths learned from the peering

local-preference

Syntax **local-preference** *local-preference*
no local-preference

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command enables setting the BGP local preference attribute in incoming routes if not specified and configures the default value for the attribute.

This value is used if the BGP route arrives from a BGP peer without the **local-preference** integer set.

The specified value can be overridden by any value set via a route policy. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of the command at the global level specifies that incoming routes with local preference set are not overridden and routes arriving without local preference set are interpreted as if the route had a local preference value of 100.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **no local-preference**

Parameters *local-preference* — the local preference value to be used as the override value, expressed as a decimal integer

Values 0 to 4294967295

loop-detect

Syntax **loop-detect** {**drop-peer** | **discard-route** | **ignore-loop** | **off**}
no loop-detect

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures how the BGP peer session handles loop detection in the AS path.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

Dynamic configuration changes of **loop-detect** are not recognized.

The **no** form of the command used at the global level reverts to the default (**ignore- loop**).

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **ignore-loop**

Parameters **drop-peer** — sends a notification to the remote peer and drops the session

discard-route — discards routes received from a peer with the same AS number as the router itself. This option prevents routes looped back to the router from being added to the routing information base and consuming memory. When this option is changed, the change will not be active for an established peer until the connection is re-established for the peer.

ignore-loop — ignores routes with loops in the AS path, but maintains peering

off — disables loop detection

min-as-origination

Syntax **min-as-origination** *seconds*
no min-as-origination

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures the minimum interval, in seconds, at which a path attribute, originated by the local router, can be advertised to a peer.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of the command used at the global level reverts to the default.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **15 s**

Parameters *seconds* — the minimum path attribute advertising interval, in seconds, expressed as a decimal integer

Values 2 to 255

min-route-advertisement

Syntax **min-route-advertisement** *seconds*
no min-route-advertisement

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures the minimum interval, in seconds, at which a prefix can be advertised to a peer.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of the command used at the global level reverts to the default.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **30 s**

Parameters *seconds* — the minimum route advertising interval, in seconds, expressed as a decimal integer

Values 2 to 255

multipath

Syntax **multipath** *integer*
no multipath

Context config>router>bgp

Description This command enables BGP multipath.

When multipath is enabled, BGP load shares traffic across multiple links. Multipath can be configured to load share traffic across a maximum of 16 routes. If the equal cost routes available are more than the configured value, then routes with the lowest next-hop IP address value are chosen.

This configuration parameter is set at the global level (applies to all peers).

Multipath is effectively disabled if the value is set to 1. When multipath is disabled and multiple equal cost routes are available, the route with the lowest next-hop IP address will be used.

The **no** form of the command reverts to the default where **multipath** is disabled.

Default **no multipath**

Parameters *integer* — the number of equal cost routes to use for multipath routing. If more equal cost routes exist than the configured value, routes with the lowest next-hop value are chosen. Setting this value to 1 disables multipath.

Values 1 to 16

neighbor

Syntax [**no**] **neighbor** *ip-address*

Context config>router>bgp>group

Description This command creates a BGP peer/neighbor instance within the context of the BGP group.

This command can be issued repeatedly to create multiple peers and their associated configurations.

The **no** form of the command is used to remove the specified neighbor and the entire configuration associated with the neighbor. The neighbor must be administratively shut down before attempting to delete it. If the neighbor is not shut down, the command will not result in any action except a warning message on the CLI indicating that the neighbor is still administratively up.

Default **no neighbors are defined**

Parameters *ip-address* — the IP address of the BGP peer router in dotted-decimal notation

Values ipv4-address: a.b.c.d (host bits must be 0)

next-hop-self

Syntax [**no**] **next-hop-self**

Context config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures the group or neighbor to always set the next-hop path attribute to its own physical interface when advertising to a peer.

This command is primarily used to avoid third-party route advertisements when connected to a multi-access network.

The **no** form of the command used at the group level allows third-party route advertisements in a multi-access network.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **no next-hop-self**

outbound-route-filtering

Syntax [**no**] **outbound-route-filtering**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command opens the configuration tree for sending or accepting BGP filter lists from peers (outbound route filtering (ORF)).

Default **no outbound-route-filtering**

extended-community

Syntax [**no**] **extended-community**

Context config>router>bgp>outbound-route-filtering
config>router>bgp>group>outbound-route-filtering
config>router>bgp>group>neighbor>outbound-route-filtering

Description This command opens the configuration tree for sending or accepting extended-community-based BGP filters. In order for the **no** version of the command to work, all sub-commands (**send-orf**, **accept-orf**) must be removed first.

Default **no extended-community**

accept-orf

Syntax [**no**] **accept-orf**

Context config>router>bgp>outbound-route-filtering>extended-community
config>router>bgp>group>outbound-route-filtering>extended-community
config>router>bgp>group>neighbor>outbound-route-filtering>extended-community

Description This command instructs the router to negotiate the receive capability in the BGP outbound route filtering (ORF) negotiation with a peer, and to accept filters that the peer wishes to send.

The **no** form of the command causes the router to remove the accept capability in the BGP ORF negotiation with a peer, and to clear any existing ORF filters that are currently in place.

Default **no accept-orf**

send-orf

Syntax **send-orf** [*comm-id*...(up to 32 max)]
no send-orf [*comm-id*]

Context config>router>bgp>outbound-route-filtering>extended-community
config>router>bgp>group>outbound-route-filtering>extended-community
config>router>bgp>group>neighbor>outbound-route-filtering>extended-community

Description This command instructs the router to negotiate the send capability in the BGP outbound route filtering (ORF) negotiation with a peer.

This command also causes the router to send a community filter, prefix filter, or AS path filter configured as an inbound filter on the BGP session to its peer as an ORF Action ADD.

The **no** form of this command causes the router to remove the send capability in the BGP ORF negotiation with a peer.

The **no** form also causes the router to send an ORF remove action for a community filter, prefix filter, or AS path filter configured as an inbound filter on the BGP session to its peer.

If the *comm-id* parameter(s) are not exclusively route target communities, the router will extract appropriate route targets and use those. If, for some reason, the *comm-id* parameter(s) specified contain no route targets, the router will not send an ORF.

Default **no send-orf**

Parameters *comm-id* — any community policy that consists exclusively of route target extended communities. If the policy is not specified, then the ORF policy is automatically generated from configured route target lists, accepted client route target ORFs, and locally configured route targets.

Values

comm-id:	target:{*ip-addr*:*comm-val* | *as-number*:*ext-comm-val*}
ip-addr:	a.b.c.d
comm-val:	0 to 65535
as-number:	0 to 65535
ext-comm-val:	0 to 4294967295

passive

Syntax [**no**] **passive**

Context config>router>bgp>group
config>router>bgp>group>neighbor

Description This command enables and disables passive mode for the BGP group or neighbor. When in passive mode, BGP will not attempt to actively connect to the configured BGP peers but responds only when it receives a connect open request from the peer.

The **no** form of the command used at the group level disables passive mode, and BGP actively attempts to connect to its peers.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **no passive**

peer-as

Syntax **peer-as** *as-number*

Context config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures the autonomous system number for the remote peer. The peer AS number must be configured for each configured peer.

For IBGP peers, the peer AS number must be the same as the autonomous system number of this router configured under the global level. This is a required command for each configured peer.

This command may be configured under the group level for all neighbors in a particular group.

Default **no AS numbers are defined**

Parameters *as-number* — the autonomous system number, expressed as a decimal integer

Values 1 to 65535

preference

Syntax **preference** *preference*
no preference

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures the route preference for routes learned from the configured peer(s).

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The value of preference behaves as follows: the lower the preference, the higher the chance of the route being the active route. The 7705 SAR OS assigns the highest default preference to BGP routes when compared to routes that are direct, static or learned via MPLS or OSPF.

The **no** form of the command used at the global level reverts to the default value.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **170**

Parameters *preference* — the route preference, expressed as a decimal integer

Values 1 to 255

prefix-limit

Syntax **prefix-limit** *limit*
no prefix-limit

Context config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures the maximum number of routes that BGP can learn from a peer.

When the number of routes reaches 90% of this limit, an SNMP trap is sent. When the limit is exceeded, the BGP peering is dropped and disabled.

The **no** form of the command removes the prefix limit.

Default **no prefix-limit**

Parameters *limit* — the number of routes that can be learned from a peer, expressed as a decimal integer

Values 1 to 4294967295

rapid-withdrawal

Syntax [**no**] **rapid-withdrawal**

Context config>router>bgp

Description This command disables the delay (Minimum Route Advertisement) on sending BGP withdrawals. Normal route withdrawals may be delayed up to the minimum route advertisement to allow for efficient packing of BGP Update messages.

The **no** form of the command removes this command from the configuration and returns withdrawal processing to the normal behavior.

Default **no rapid-withdrawal**

remove-private

Syntax [**no**] **remove-private**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command allows private AS numbers to be removed from the AS path before advertising them to BGP peers. The **no** form of the command includes private AS numbers in the AS path attribute.

When the **remove-private** parameter is set at the global level, it applies to all peers regardless of group or neighbor configuration. When the parameter is set at the group level, it applies to all peers in the group regardless of the neighbor configuration.

The 7705 SAR OS recognizes the set of AS numbers that are defined by IANA as private. These are AS numbers in the range 64512 through 65535, inclusive.

The **no** form of the command used at the global level reverts to the default value.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **no remove-private**

route-target-list

Syntax **route-target-list** *comm-id* [*comm-id* **..**(up to 15 max)]
no route-target-list [*comm-id*]

Context config>router>bgp

Description This command specifies the route target(s) to be accepted and advertised from/to route reflector clients. If the **route-target-list** is a non-null list, only routes with one or more of the given route targets are accepted or advertised to route reflector clients.

This command is only applicable if the router is a route-reflector server. This parameter has no effect on non-route-reflector clients.

If the **route-target-list** is assigned at the global level, then the list applies to all route-reflector clients connected to the system.

The **no** form of the command with a specified route target community removes the specified community from the **route-target-list**.

The **no** form of the command entered without a route target community removes all communities from the list.

Default **no route-target-list**

Parameters *comm-id* — the route target community

Values		
	comm-id:	target:{*ip-addr*:*comm-val* | *as-number*:*ext-comm-val*}
	ip-addr:	a.b.c.d
	comm-val:	0 to 65535
	as-number:	1..65535
	ext-comm-val:	0 to 4294967295

router-id

Syntax **router-id** *ip-address*
no router-id

Context config>router>bgp

Description This command specifies the router ID to be used with this BGP instance. If no router ID is specified, the system interface IP address is used.

Changing the BGP router ID on an active BGP instance causes the BGP instance to restart with the new router ID. The router ID must be set to a valid host address.

Default **no router-id**

.

Parameters *ip-address* — the router ID, expressed in dotted-decimal notation. The allowed value is a valid routable IP address on the router, either an interface or system IP address. It is highly recommended that this address be the system IP address.

Values a.b.c.d

shutdown

Syntax [**no**] **shutdown**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

The **no** form of this command administratively enables an entity.

Unlike other commands and parameters where the default state is not indicated in the configuration file, the **shutdown** and **no shutdown** states are always indicated in system-generated configuration files.

Default administrative states for services and service entities are described in **Special Cases**.

The **no** form of the command places an entity in an administratively enabled state.

Special Cases

BGP Global — the BGP protocol is created in the **no shutdown** state

BGP Group — BGP groups are created in the **no shutdown** state

BGP Neighbor — BGP neighbors/peers are created in the **no shutdown** state

vpn-apply-export

Syntax [**no**] **vpn-apply-export**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command causes the base instance BGP export route policies to be applied to VPN-IPv4 routes.

The **no** form of the command disables the application of the base instance BGP export route policies to VPN-IPv4 routes.

Default **no vpn-apply-export**

vpn-apply-import

Syntax [**no**] **vpn-apply-import**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command causes the base instance BGP import route policies to be applied to VPN-IPv4 routes.

The **no** form of the command disables the application of the base instance BGP import route policies to VPN-IPv4 routes.

Default **no vpn-apply-import**

Other BGP-Related Commands

aggregate

Syntax **aggregate** *ip-prefix/ip-prefix-length* [**summary-only**]
no aggregate *ip-prefix/ip-prefix-length*

Context config>router

Description This command creates an aggregate route.

Use this command to group a number of routes with common prefixes into a single entry in the routing table. This reduces the number of routes that need to be advertised by this router and reduces the number of routes in the routing tables of downstream routers.

Both the original components and the aggregated route (source protocol aggregate) are offered to the Routing Table Manager (RTM). Subsequent policies can be configured to assign protocol-specific characteristics, such as the OSPF tag, to aggregate routes.

Multiple entries with the same prefix but a different mask can be configured; routes are aggregated to the longest mask. If one aggregate is configured as 10.0/16 and another as 10.0.0/24, then route 10.0.128/17 would be aggregated into 10.0/16 and route 10.0.0.128/25 would be aggregated into 10.0.0/24. If multiple entries are made with the same prefix and the same mask, the previous entry is overwritten.

The **no** form of the command removes the aggregate.

Default **no aggregate**

Parameters *ip-prefix/ip-prefix-length* — the destination address of the aggregate route in dotted-decimal notation

Values	ip-prefix:	a.b.c.d (host bits must be 0)
	ip-prefix-length:	0 to 32

summary-only — suppresses advertisement of more specific component routes for the aggregate. To remove the **summary-only** option, enter the same aggregate command without the **summary-only** parameter.

autonomous-system

Syntax **autonomous-system** *as-number*
no autonomous-system

Context config>router

Description This command configures the autonomous system (AS) number for the router. A router can only belong to one AS. An AS number is a globally unique number within an AS. This number is used to exchange exterior routing information with neighboring ASs and as an identifier of the AS itself.

If the AS number is changed on a router with an active BGP instance, the new AS number is not used until the BGP instance is restarted either by administratively disabling/enabling (**shutdown**/ **no shutdown**) the BGP instance or rebooting the system with the new configuration.

Default **no autonomous-system**

Parameters *as-number* — the autonomous system number expressed as a decimal integer

Values 1 to 65535

router-id

Syntax **router-id** *ip-address*
[**no**] **router-id**

Context config>router

Description This command configures the router ID for the router instance.

The router ID is used by both OSPF and BGP routing protocols in this instance of the routing table manager. IS-IS uses the router ID value as its system ID.

When configuring a new router ID, protocols are not automatically restarted with the new router ID. The next time a protocol is initialized, the new router ID is used. This can result in an interim period of time when different protocols use different router IDs.

To force the new router ID to be used, issue the **shutdown** and **no shutdown** commands for each protocol that uses the router ID, or restart the entire router.

The **no** form of the command reverts to the default value.

Default The system uses the system interface address (which is also the loopback address). If a system interface address is not configured, use the last 32 bits of the chassis MAC address.

Parameters *ip-address* — 32-bit router ID, expressed in dotted-decimal notation or as a decimal value

Show Commands

router

Syntax **router** [*router-instance*]

Context show

Description The command displays router instance information.

Parameters *router-instance* — specifies either the router name or service ID

Values	router-name:	Base, management
	service-id:	1 to 2147483647
Default	Base	

bgp

Syntax **bgp**

Context show>router

Description This command enables the context to display BGP related information.

damping

Syntax **damping** [*damp-type*] [**detail**]
damping [*ip-prefix* **|** *prefix-length*] [**detail**]

Context show>router>bgp

Description This command displays BGP routes that have been dampened due to route flapping. This command can be entered with or without a route parameter.

When the **detail** keyword is included, more detailed information displays.

When only the command is entered (without any parameters included except detail), then all dampened routes are listed.

When a parameter is specified, then the matching route or routes are listed.

When a **decayed**, **history**, or **suppressed** keyword is specified, only those types of dampened routes are listed.

Parameters *ip-prefix* — the specified IP prefix and length

Values *ipv4-prefix* a.b.c.d (host bits must be 0)

ipv4-prefix-length 0 to 32

damp-type — the type of damping to display

Values
decayed displays damping entries that are decayed but are not suppressed
history displays damping entries that are withdrawn but have history
suppressed displays damping entries suppressed because of route damping

detail — displays detailed information

Output The following output is an example of BGP damping information, and Table 1 describes the fields.

Sample Output

```
*A:7705_ALU-2>show>router>bgp# damping
===============================================================================
 BGP Router ID : 55.55.55.55       AS : 1       Local AS : 1
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best
===============================================================================
BGP Damped Routes
===============================================================================
Flag  Network                                                   Reuse
      From
      AS-Path
-------------------------------------------------------------------------------
ud*i  12.149.7.0/24                                             00h00m00s
      10.0.28.1
      60203 65001 19855 3356 1239 22406
si    24.155.6.0/23                                             00h43m41s
      10.0.28.1
      60203 65001 19855 3356 2914 7459
si    24.155.8.0/22                                             00h38m31s
      10.0.28.1
      60203 65001 19855 3356 2914 7459
===============================================================================
```

```
*A:7705_ALU-2>show>router>bgp# damping 10.10.10.0/32 detail
===============================================================================
BGP Router ID : 10.0.0.14         AS : 65206   Local AS : 65206
===============================================================================
Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best
===============================================================================
BGP Damped Routes 15.203.192.0/18
===============================================================================
-------------------------------------------------------------------------------
Network : 15.203.192.0/18
-------------------------------------------------------------------------------
```

```
Network         : 15.203.192.0/18        Peer: 10.0.28.1
NextHop         : 10.0.28.1              Reuse time: 00h00m00s
Peer AS         : 60203                  Peer Router-Id: 32.32.27.203
Local Pref      : none
Age             : 00h00m42s              Last update: 02d01h20m
FOM Present     : 2003                   FOM Last upd.: 2025
Number of Flaps : 2                      Flags: ud*i
Path            : 60203 65001 19855 3356 702 1889
Applied Policy  : default-damping-profile
-------------------------------------------------------------------------------
Paths : 1
===============================================================================
*A:7705_ALU-2>show>router>bgp#


*A:7705_ALU-2>show>router>bgp# damping suppressed detail
===============================================================================
 BGP Router ID : 55.55.55.55       AS : 1       Local AS : 1
===============================================================================
Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best
===============================================================================
BGP Damped Routes (Suppressed)
===============================================================================
-------------------------------------------------------------------------------
Network : 44.44.48.0/20
-------------------------------------------------------------------------------
Network         : 15.142.48.0/20         Peer: 10.0.28.1
NextHop         : 10.0.28.1              Reuse time: 00h29m22s
Peer AS         : 60203                  Peer Router-Id: 32.32.27.203
Local Pref      : none
Age             : 00h01m28s              Last update: 02d01h20m
FOM Present     : 2936                   FOM Last upd.: 3001
Number of Flaps : 3                      Flags: si
Path            : 60203 65001 19855 3356 702 1889
Applied Policy  : default-damping-profile
-------------------------------------------------------------------------------
Network : 44.44.128.0/19
-------------------------------------------------------------------------------
Network         : 15.200.128.0/19        Peer : 10.0.28.1
NextHop         : 10.0.28.1              Reuse time : 00h29m22s
Peer AS         : 60203                  Peer Router-Id: 32.32.27.203
Local Pref      : none
Age             : 00h01m28s              Last update: 02d01h20m
FOM Present     : 2936                   FOM Last upd.: 3001
Number of Flaps : 3                      Flags: si
Path            : 60203 65001 19855 3356 702 1889
Applied Policy  : default-damping-profile
===============================================================================
*A:7705_ALU-2>show>router>bgp#
```

```
A:ALA-12# show router bgp damping detail
===============================================================================
BGP Router ID : 10.0.0.14 AS : 65206 Local AS : 65206
===============================================================================
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, - best
===============================================================================
BGP Damped Routes
===============================================================================
-------------------------------------------------------------------------------
Network : 12.149.7.0/24
-------------------------------------------------------------------------------
Network         : 12.149.7.0/24          Peer : 10.0.28.1
NextHop         : 10.0.28.1              Reuse time : 00h00m00s
Peer AS         : 60203                  Peer Router-Id : 32.32.27.203
Local Pref      : none
Age             : 00h22m09s              Last update : 02d00h58m
FOM Present     : 738                    FOM Last upd. : 2039
Number of Flaps : 2                      Flags : ud*i
Path : 60203 65001 19855 3356 1239 22406
Applied Policy : default-damping-profile
-------------------------------------------------------------------------------
A:ALA-12#
```

Table 1: Show BGP Damping Output Fields

Label	Description
`BGP Router ID`	The local BGP router ID
`AS`	The configured autonomous system number
`Local AS`	The configured or inherited local AS for the specified peer group. If not configured, then it is the same value as the AS.
`Flag(s)`	Legend: Status codes: u - used s - suppressed h - history d - decayed * - valid If an * is not present, then the status is invalid Origin codes: i - IGP e - EGP ? - incomplete > - best
`Network`	IP prefix and mask length for the route
`From`	The originator ID path attribute value

Table 1: Show BGP Damping Output Fields (Continued)

Label	Description
`Reuse`	The time when a suppressed route can be used again
`AS-Path`	The BGP AS path for the route
`Peer`	The router ID of the advertising router
`NextHop`	The BGP next hop for the route
`Reuse time`	The time when the route can be reused
`Peer AS`	The autonomous system number of the advertising router
`Peer Router-Id`	The router ID of the advertising router
`Local Pref`	The BGP local preference path attribute for the route
`Age`	The length of time in hour/minute/second (HH:MM:SS) format
`Last update`	The time that BGP was updated last in day/hour/minute (DD:HH:MM) format
`FOM Present`	The current Figure of Merit (FOM) value
`FOM Last upd.`	The last update Figure of Merit (FOM) value
`Number of Flaps`	The number of route flaps in the neighbor connection
`Path`	The BGP AS path for the route
`Applied Policy`	The applied route policy name

group

Syntax **group** [*name*] [**detail**]

Context show>router>bgp

Description This command displays group information for a BGP peer group. This command can be entered with or without parameters.

When this command is entered without a group name, information about all peer groups displays.

When the command is issued with a specific group name, information only pertaining to that specific peer group displays.

The “State” field displays the BGP group’s operational state. Valid states are:

Up — BGP global process is configured and running

Down — BGP global process is administratively shut down and not running

Disabled — BGP global process is operationally disabled. The process must be restarted by the operator.

Parameters *name* — displays information for the BGP group specified

detail — displays detailed information

Output The following output is an example of BGP group information, and Table 2 describes the fields.

Sample Output

```
*A:7705_ALU-2>show>router>bgp# group
===============================================================================
BGP Group
===============================================================================
-------------------------------------------------------------------------------
Group            : bgp_group
-------------------------------------------------------------------------------
Group Type       : No Type               State            : Up
Peer AS          : n/a                   Local AS         : 1
Local Address    : n/a                   Loop Detect      : Ignore
Import Policy    : None Specified / Inherited
Export Policy    : None Specified / Inherited
Hold Time        : 90                    Keep Alive       : 30
Cluster Id       : None                  Client Reflect   : Enabled
NLRI             : Unicast               Preference       : 170
TTL Security     : Disabled              Min TTL Value    : n/a
Graceful Restart : Enabled               Stale Routes Time: 360
Auth key chain   : n/a
Bfd Enabled      : Disabled

List of Peers
- 44.44.44.44 :

Total Peers      : 1                     Established      : 0
-------------------------------------------------------------------------------
Peer Groups : 1
```

```
*A:7705_ALU-2>show>router>bgp# group detail

===============================================================================
BGP Group  (detail)
===============================================================================
-------------------------------------------------------------------------------
Group            : bgp_group
-------------------------------------------------------------------------------
Group Type       : No Type              State            : Up
Peer AS          : n/a                  Local AS         : 1
Local Address    : n/a                  Loop Detect      : Ignore
Connect Retry    : 120                  Authentication   : None
Local Pref       : 100                  MED Out          : 0
Multihop         : 0 (Default)          AS Override      : Disabled
Min Route Advt.  : 30                   Min AS Originate : 15
Prefix Limit     : No Limit             Passive          : Disabled
Next Hop Self    : Disabled             Aggregator ID 0  : Disabled
Remove Private   : Disabled             Damping          : Enabled
Import Policy    : None Specified / Inherited
Export Policy    : None Specified / Inherited
Hold Time        : 90                   Keep Alive       : 30
Cluster Id       : None                 Client Reflect   : Enabled
NLRI             : Unicast              Preference       : 170
TTL Security     : Disabled             Min TTL Value    : n/a
Graceful Restart : Enabled              Stale Routes Time: 360
Auth key chain   : n/a
Bfd Enabled      : Disabled

List of Peers
- 44.44.44.44 :

Total Peers      : 1                    Established      : 0
-------------------------------------------------------------------------------
Peer Groups : 1
===============================================================================
*A:7705_ALU-2>show>router>bgp#
```

Table 2: Show BGP Group Output Fields

Label	Description
Group	The BGP group name
Group Type	No Type — peer type not configured External — peer type configured as external BGP peers Internal — peer type configured as internal BGP peers
State	Disabled — the BGP peer group has been operationally disabled Down — the BGP peer group is operationally inactive Up — the BGP peer group is operationally active
Peer AS	The configured or inherited peer AS for the specified peer group
Local AS	The configured or inherited local AS for the specified peer group
Local Address	The configured or inherited local address for originating peering for the specified peer group
Loop Detect	The configured or inherited loop detect setting for the specified peer group
Connect Retry	The configured or inherited connect retry timer value
Authentication	None — no authentication is configured MD5 — MD5 authentication is configured
Local Pref	The configured or inherited local preference value
MED Out	The configured or inherited MED value assigned to advertised routes without a MED attribute
Import Policy	The configured import policies for the peer group
Export Policy	The configured export policies for the peer group
Hold Time	The configured hold time setting
Keep Alive	The configured keepalive setting
Cluster Id	The configured route reflector cluster ID None — No cluster ID has been configured
Client Reflect	Disabled — the BGP route reflector will not reflect routes to this neighbor Enabled — the BGP route reflector is configured to reflect routes to this neighbor

Table 2: Show BGP Group Output Fields (Continued)

Label	Description
NLRI	The type of network layer reachability information that the specified peer group can accept Unicast — IPv4 unicast routing information can be carried
Preference	The configured route preference value for the peer group
TTL Security	The state of the TTL security
Min TTL Value	The minimum value configured for TTL
Graceful Restart	The state of graceful restart
Stale Routes Time	The length of time that stale routes are kept in the route table
Auth key chain	The value for the authentication key chain
Bfd Enabled	Enabled — BFD is enabled Disabled — BFD is disabled
List of Peers	A list of BGP peers configured under the peer group
Total Peers	The total number of peers configured under the peer group
Established	The total number of peers that are in an established state
Peer Groups	The number of peer groups
Multihop	The maximum number of router hops a BGP connection can traverse
AS Override	The setting of the AS override
Min Route Advt.	The minimum amount of time that must pass between route updates for the same IP prefix
Min AS Originate	The minimum amount of time that must pass between updates for a route originated by the local router
Prefix Limit	No Limit — no route limit assigned to the BGP peer group 1 — 4294967295 — the maximum number of routes BGP can learn from a peer
Passive	Disabled — BGP attempts to establish a BGP connection with a neighbor in the specified peer group Enabled — BGP will not actively attempt to establish a BGP connection with a neighbor in the specified peer group

Table 2: Show BGP Group Output Fields (Continued)

Label	Description
Next Hop Self	Disabled — BGP is not configured to send only its own IP address as the BGP next hop in route updates to neighbors in the peer group Enabled — BGP sends only its own IP address as the BGP next hop in route updates to neighbors in the specified peer group
Aggregator ID 0	Disabled — BGP is not configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates sent to the neighbor in the peer group Enabled — BGP is configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates sent to the neighbor in the peer group
Remove Private	Disabled — BGP will not remove all private AS numbers from the AS path attribute in updates sent to the neighbor in the peer group Enabled — BGP removes all private AS numbers from the AS path attribute in updates sent to the neighbor in the peer group
Damping	Disabled — the peer group is configured not to dampen route flaps Enabled — the peer group is configured to dampen route flaps

neighbor

Syntax **neighbor** [*ip-address* [[*family*] *filter1* [**brief**]]]
neighbor [*as-number* [[*family*] *filter2*]]
neighbor [*ip-address*] **orf** [*filter3*]
neighbor [*ip-address*] **graceful-restart**

Context show>router>bgp

Description This command displays BGP neighbor information. This command can be entered with or without any parameters.

When this command is issued without any parameters, information about all BGP peers displays.

When the command is issued with a specific IP address or ASN, information regarding only that specific peer or peers with the same AS displays.

When either **received-routes** or **advertised-routes** is specified, the routes received from or sent to the specified peer are listed (see second output example).

Note: This information is not available by SNMP.

When either **history** or **suppressed** is specified, the routes learned from those peers that either have a history or are suppressed (respectively) are listed.

The “State” field displays the BGP peer’s protocol state. In addition to the standard protocol states, this field can also display the “Disabled” operational state, which indicates that the peer is operationally disabled and must be restarted by the operator.

Parameters *ip-address* — the specified IP address for which to display information

Values *ipv4-address*: a.b.c.d (host bits must be 0)

as-number — the specified AS number for which to display information

Values 1 to 65535

family — the type of routing information to be distributed by this peer group

Values **ipv4** — displays only those BGP peers that have the IPv4 family enabled and not those capable of exchanging IP-VPN routes
vpn-ipv4 —displays the contents of the multicast routing table

filter1 — displays information for the specified IP address

Values **received-routes** — displays the number of routes received from this peer
advertised-routes — displays the number of routes advertised by this peer
history — displays statistics for dampened routes
suppressed — displays the number of paths from this peer that have been suppressed by damping
detail — displays detailed information pertaining to *filter1*

filter2 — displays information for the specified AS number

Values **history** — displays statistics for dampened routes
suppressed — displays the number of paths from this peer that have been suppressed by damping
detail — displays detailed information pertaining to *filter2*

brief — displays information in a brief format. This parameter is only supported with received-routes and advertised-routes.

orf — displays outbound route filtering for the BGP instance. ORF (Outbound Route Filtering) is used to inform a neighbor of targets (using target-list) that it is willing to receive. This mechanism helps lessen the update exchanges between neighbors and saves CPU cycles to process routes that could have been received from the neighbor only to be dropped/ignored.

filter3 — displays path information for the specified IP address

Values **send** — displays the number of paths sent to this peer
receive — displays the number of paths received from this peer

graceful-restart — displays neighbors configured for graceful restart

Output The following outputs are examples of BGP neighbor information:

- BGP neighbor (standard and detailed) (Sample Output - BGP Neighbor (standard and detailed), Table 3)
- BGP neighbor (advertised and received) (Sample Output - BGP Neighbor (advertised- and received-routes), Table 4)
- BGP neighbor (graceful restart) (Sample Output - BGP Neighbor (graceful restart), Table 5)

Sample Output - BGP Neighbor (standard and detailed)

```
*A:7705_ALU-2>show>router>bgp# neighbor
===============================================================================
BGP Neighbor
===============================================================================
-------------------------------------------------------------------------------
Peer  : 10.10.10.12
Group : ibgp_group
-------------------------------------------------------------------------------
Peer AS              : 65000            Peer Port            : 49550
Peer Address         : 10.10.10.12
Local AS             : 65000            Local Port           : 179
Local Address        : 10.10.10.1
Peer Type            : Internal
State                : Established      Last State           : Established
Last Event           : recvKeepAlive
Last Error           : Cease
Local Family         : IPv4 VPN-IPv4
Remote Family        : IPv4 VPN-IPv4
Hold Time            : 90               Keep Alive           : 30
Active Hold Time     : 90               Active Keep Alive    : 30
Cluster Id           : None
Preference           : 170              Num of Flaps         : 0
Recd. Paths          : 19
IPv4 Recd. Prefixes  : 600              IPv4 Active Prefixes : 563
IPv4 Suppressed Pfxs : 0                VPN-IPv4 Suppr. Pfxs : 0
VPN-IPv4 Recd. Pfxs  : 8656             VPN-IPv4 Active Pfxs : 8656
Mc IPv4 Recd. Pfxs.  : 0                Mc IPv4 Active Pfxs. : 0
Mc IPv4 Suppr. Pfxs  : 0
Input Queue          : 0                Output Queue         : 0
i/p Messages         : 1141             o/p Messages         : 1041
i/p Octets           : 449029           o/p Octets           : 163814
i/p Updates          : 151              o/p Updates          : 50
TTL Security         : Disabled         Min TTL Value        : n/a
Graceful Restart     : Disabled         Stale Routes Time    : n/a
Advertise Inactive   : Disabled         Peer Tracking        : Disabled
Auth key chain       : n/a
Bfd Enabled          : Enabled
Local Capability     : RouteRefresh MP-BGP
Remote Capability    : RouteRefresh MP-BGP
Import Policy        : None Specified / Inherited
Export Policy        : stmt1

-------------------------------------------------------------------------------
Neighbors : 1
===============================================================================
*A:7705_ALU-2>show>router>bgp#
```

```
*A:7705_ALU-2>show>router>bgp# neighbor 10.10.10.12 detail

===============================================================================
BGP Neighbor
===============================================================================
-------------------------------------------------------------------------------
Peer  : 10.10.10.12
Group : iBGP
-------------------------------------------------------------------------------
Peer AS              : 65000            Peer Port            : 49550
Peer Address         : 10.10.10.12
Local AS             : 65000            Local Port           : 179
Local Address        : 10.10.10.1
Peer Type            : Internal
State                : Established      Last State           : Established
Last Event           : recvKeepAlive
Last Error           : Cease
Local Family         : IPv4 VPN-IPv4
Remote Family        : IPv4 VPN-IPv4
Connect Retry        : 120              Local Pref.          : 70
Min Route Advt.      : 30               Min AS Orig.         : 15
Multihop             : 0 (Default)      AS Override          : Disabled
Damping              : Disabled         Loop Detect          : Ignore
MED Out              : No MED Out       Authentication       : None
Next Hop Self        : Disabled         AggregatorID Zero    : Disabled
Remove Private       : Disabled         Passive              : Disabled
Peer Identifier      : 10.10.10.12      Fsm Est. Trans       : 1
Fsm Est. Time        : 22h42m46s        InUpd Elap. Time     : 22h54m31s
Prefix Limit         : No Limit
Hold Time            : 90               Keep Alive           : 30
Active Hold Time     : 90               Active Keep Alive    : 30
Cluster Id           : None             Client Reflect       : Disabled
Preference           : 170              Num of Flaps         : 0
Recd. Paths          : 19
IPv4 Recd. Prefixes  : 600              IPv4 Active Prefixes : 563
IPv4 Suppressed Pfxs : 0                VPN-IPv4 Suppr. Pfxs : 0
VPN-IPv4 Recd. Pfxs  : 8656             VPN-IPv4 Active Pfxs : 8656
Mc IPv4 Recd. Pfxs.  : 0                Mc IPv4 Active Pfxs. : 0
Mc IPv4 Suppr. Pfxs  : 0
Input Queue          : 0                Output Queue         : 0
i/p Messages         : 2881             o/p Messages         : 2777
i/p Octets           : 482089           o/p Octets           : 196798
i/p Updates          : 151              o/p Updates          : 50
TTL Security         : Disabled         Min TTL Value        : n/a
Graceful Restart     : Disabled         Stale Routes Time    : n/a
Advertise Inactive   : Disabled         Peer Tracking        : Disabled
Auth key chain       : n/a
Bfd Enabled          : Enabled
Local Capability     : RouteRefresh MP-BGP
Remote Capability    : RouteRefresh MP-BGP
Import Policy        : None Specified / Inherited
Export Policy        : stmt1

-------------------------------------------------------------------------------
Neighbors : 1
===============================================================================
*A:7705_ALU-2>show>router>bgp#
```

```
*A:7705_ALU-2>show>router>bgp# neighbor 10.0.0.11 orf
===============================================================================
BGP Neighbor 10.0.0.11 ORF
===============================================================================
Send List (Automatic)
-------------------------------------------------------------------------------
target:65535:10
target:65535:20
===============================================================================
*A:7705_ALU-2>show>router>bgp#
```

```
*A:7705_ALU-2>show>router>bgp# neighbor 10.0.0.1 orf
===============================================================================
BGP Neighbor 10.0.0.1 ORF
===============================================================================
Receive List
-------------------------------------------------------------------------------
target:65535:10
target:65535:20
===============================================================================
*A:7705_ALU-2>show>router>bgp#
```

Table 3: Show BGP Neighbor (Standard and Detailed) Output Fields

Label	Description
`Peer`	The IP address of the configured BGP peer
`Group`	The BGP peer group to which this peer is assigned
`Peer AS`	The configured or inherited peer AS for the peer group
`Peer Address`	The configured address for the BGP peer
`Peer Port`	The TCP port number used on the far-end system
`Local AS`	The configured or inherited local AS for the peer group
`Local Address`	The configured or inherited local address for originating peering for the peer group
`Local Port`	The TCP port number used on the local system
`Peer Type`	`External` — peer type configured as external BGP peers
	`Internal` — peer type configured as internal BGP peers

Table 3: Show BGP Neighbor (Standard and Detailed) Output Fields (Continued)

Label	Description
State	Idle — The BGP peer is not accepting connections
	Active — BGP is listening for and accepting TCP connections from this peer
	Connect — BGP is attempting to establish a TCP connection with this peer
	Open Sent — BGP has sent an OPEN message to the peer and is waiting for an OPEN message from the peer
	Open Confirm — BGP has received a valid OPEN message from the peer and is awaiting a KEEPALIVE or NOTIFICATION
	Established — BGP has successfully established a peering session and is exchanging routing information
Last State	Idle — The BGP peer is not accepting connections
	Active — BGP is listening for and accepting TCP connections from this peer
	Connect — BGP is attempting to establish a TCP connections with this peer
	Open Sent — BGP has sent an OPEN message to the peer and is waiting for an OPEN message from the peer
	Open Confirm — BGP has received a valid OPEN message from the peer and is awaiting a KEEPALIVE or NOTIFICATION

Table 3: Show BGP Neighbor (Standard and Detailed) Output Fields (Continued)

Label	Description
Last Event	start — BGP has initialized the BGP neighbor
	stop — BGP has disabled the BGP neighbor
	open — BGP transport connection is opened
	close — BGP transport connection is closed
	openFail — BGP transport connection failed to open
	error — BGP transport connection error
	connectRetry — the connect retry timer expired
	holdTime — the hold time timer expired
	keepAlive — the keepalive timer expired
	recvOpen — BGP has received an OPEN message
	revKeepalive — BGP has received a KEEPALIVE message
	recvUpdate — BGP has received an UPDATE message
	recvNotify — BGP has received a NOTIFICATION message
	None — no events have occurred
Last Error	The last BGP error and subcode to occur on the BGP neighbor
Local Family	The configured local family value
Remote Family	The configured remote family value
Connect Retry	The configured or inherited connect retry timer value
Local Pref.	The configured or inherited local preference value
Min Route Advt.	The minimum amount of time that must pass between route updates for the same IP prefix
Min AS Originate	The minimum amount of time that must pass between updates for a route originated by the local router
Multihop	The maximum number of router hops a BGP connection can traverse
Damping	Disabled — the BGP neighbor is configured not to dampen route flaps
	Enabled — the BGP neighbor is configured to dampen route flaps

Table 3: Show BGP Neighbor (Standard and Detailed) Output Fields (Continued)

Label	Description
`Loop Detect`	`Ignore` — The BGP neighbor is configured to ignore routes with an AS loop
	`Drop` — The BGP neighbor is configured to drop the BGP peering if an AS loop is detected
	`Off` — AS loop detection is disabled for the neighbor
`MED Out`	The configured or inherited MED value assigned to advertised routes without a MED attribute
`Authentication`	`None` — no authentication is configured
	`MD5` — MD5 authentication is configured
`Next Hop Self`	`Disabled` — BGP is not configured to send only its own IP address as the BGP next hop in route updates to the specified neighbor
	`Enabled` — BGP will send only its own IP address as the BGP next hop in route updates to the neighbor
`AggregatorID Zero`	`Disabled` — the BGP neighbor is not configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates
	`Enabled` — the BGP neighbor is configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates
`Remove Private`	`Disabled` — BGP will not remove all private AS numbers from the AS path attribute in updates sent to the specified neighbor
	`Enabled` — BGP will remove all private AS numbers from the AS path attribute in updates sent to the specified neighbor
`Passive`	`Disabled` — BGP will actively attempt to establish a BGP connection with the specified neighbor
	`Enabled` — BGP will not actively attempt to establish a BGP connection with the specified neighbor
`Peer Identifier`	The IP identifier for the peer router
`Prefix Limit`	`No Limit` — no route limit assigned to the BGP peer group
	`1 — 4294967295` — the maximum number of routes BGP can learn from a peer
`Hold Time`	The configured hold time setting
`Keep Alive`	The configured keepalive setting

Table 3: Show BGP Neighbor (Standard and Detailed) Output Fields (Continued)

Label	Description
`Active Hold Time`	The negotiated hold time, if the BGP neighbor is in an established state
`Active Keep Alive`	The negotiated keepalive time, if the BGP neighbor is in an established state
`Cluster Id`	The configured route reflector cluster ID
	`None` — no cluster ID has been configured
`Client Reflect`	`Disabled` — The BGP route reflector is configured not to reflect routes to this neighbor
	`Enabled` — The BGP route reflector is configured to reflect routes to this neighbor
`Preference`	The configured route preference value for the peer group
`Num of Flaps`	The number of route flaps in the neighbor connection
`Recd. Prefixes`	The number of routes received from the BGP neighbor
`Recd. Paths`	The number of unique sets of path attributes received from the BGP neighbor
`IPv4 Recd. Prefixes`	The number of unique sets of IPv4 path attributes received from the BGP neighbor
`IPv4 Active Prefixes`	The number of IPv4 routes received from the BGP neighbor and active in the forwarding table
`IPv4 Suppressed Pfxs`	The number of unique sets of IPv4 path attributes received from the BGP neighbor and suppressed due to route damping
`VPN-IPv4 Suppr. Pfxs`	The number of unique sets of VPN-IPv4 path attributes received from the BGP neighbor and suppressed due to route damping
`VPN-IPv4 Recd. Pfxs`	The number of unique sets of VPN-IPv4 path attributes received from the BGP neighbor
`VPN-IPv4 Active Pfxs`	The number of VPN-IPv4 routes received from the BGP neighbor and active in the forwarding table
`Mc IPv4 Recd. Pfxs`	The number of unique sets of multiclass IPv4 path attributes received from the BGP neighbor
Mc IPv4 Active Pfxs	The number of multiclass IPv4 routes received from the BGP neighbor and active in the forwarding table

Table 3: Show BGP Neighbor (Standard and Detailed) Output Fields (Continued)

Label	Description
Mc IPv4 Suppr. Pfxs	The number of unique sets of multiclass IPv4 path attributes received from the BGP neighbor and suppressed due to route damping
`Input Queue`	The number of BGP messages to be processed
`Output Queue`	The number of BGP messages to be transmitted
`i/p Messages`	The total number of packets received from the BGP neighbor
`o/p Messages`	The total number of packets sent to the BGP neighbor
`i/p Octets`	The total number of octets received from the BGP neighbor
`o/p Octets`	The total number of octets sent to the BGP neighbor
`i/p Updates`	The total number of updates received from the BGP neighbor
`o/p Updates`	The total number of updates sent to the BGP neighbor
`TTL Security`	The state of the TTL security
`Min TTL Value`	The minimum value configured for TTL
`Graceful Restart`	The state of graceful restart
`Stale Routes Time`	The length of time that stale routes are kept in the route table
`Auth key chain`	The value for the authentication key chain
`Bfd Enabled`	`Enabled` — BFD is enabled `Disabled` — BFD is disabled
`Local Capability`	The capability of the local BGP speaker; for example, route refresh, MP-BGP, ORF
`Remote Capability`	The capability of the remote BGP peer; for example, route refresh, MP-BGP, ORF
`Export Policy`	The configured export policies for the peer group
`Import Policy`	The configured import policies for the peer group

Sample Output - BGP Neighbor (advertised- and received-routes)

```
*A:7705_ALU-2>show>router>bgp# neighbor 44.44.44.44 advertised-routes
===============================================================================
 BGP Router ID : 55.55.55.55       AS : 1       Local AS : 1
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best
===============================================================================
BGP IPv4 Routes
===============================================================================
Flag  Network                                           LocalPref   MED
      Nexthop                                                       VPN Label
      As-Path
-------------------------------------------------------------------------------
?    10.0.0.02/32 100 none
     10.0.0.16  -
     No As-Path
?    10.0.6.04/24 100 none
     10.0.0.16 -
     No As-Path
-------------------------------------------------------------------------------
Routes : 2
===============================================================================
*A:7705_ALU-2>show>router>bgp#


*A:PE1>show>router>bgp# neighbor 10.10.10.12 advertised-routes  brief
===============================================================================
 BGP Router ID : 10.10.10.1        AS : 65000   Local AS : 65000
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best

===============================================================================
BGP IPv4 Routes
===============================================================================
Flag  Network
-------------------------------------------------------------------------------
?     10.10.10.1/32
?     12.12.1.0/24
?     20.0.0.0/24
?     21.0.0.0/24
?     88.88.1.0/24
-------------------------------------------------------------------------------
Routes : 5
===============================================================================
*A:7705_ALU-2>show>router>bgp#
```

```
*A:7705_ALU-2>show>router>bgp# neighbor 44.44.44.44 received-routes
===============================================================================
 BGP Router ID : 55.55.55.55       AS : 1       Local AS : 1
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best
===============================================================================
BGP IPv4 Routes
===============================================================================
Flag  Network                                           LocalPref   MED
      Nexthop                                                       VPN Label
      As-Path
-------------------------------------------------------------------------------
?    10.0.0.16/32 100 none
     10.0.0.16     -
     No As-Path
?    10.0.6.0/24  100 none
     10.0.0.16     -
     No As-Path
?    10.0.8.0/24  100 none
     10.0.0.16     -
     No As-Path
?    10.0.12.0/24 100 none
     10.0.0.16     -
     No As-Path
-------------------------------------------------------------------------------
Routes : 4
===============================================================================
*A:7705_ALU-2>show>router>bgp#
```

Table 4: Show BGP Neighbor (Advertised- and Received-Routes) Output Fields

Label	Description
`BGP Router ID`	The local BGP router ID
`AS`	The configured autonomous system number
`Local AS`	The configured local AS setting. If not configured, then it is the same value as the AS.
`Flag(s)`	Legend: Status codes: u - used s - suppressed h - history d - decayed * - valid If an * is not present, then the status is invalid Origin codes: i - IGP e - EGP ? - incomplete > - best
`Network`	The route IP prefix and mask length for the route
`Next Hop`	The BGP next hop for the route
`LocalPref`	The BGP local preference path attribute for the route
`MED`	The BGP Multi-Exit Discriminator (MED) path attribute for the route
`AS-Path`	The BGP AS path for the route

Sample Output - BGP Neighbor (graceful restart)

```
*A:7705_ALU-2>show>router>bgp# neighbor 20.10.120.44 graceful-restart
===============================================================================
BGP Neighbor 20.10.120.44 Graceful Restart
===============================================================================
Graceful Restart locally configured for peer  : Enabled
Peer's Graceful Restart feature               : Enabled
NLRI(s) that peer supports restart for        : IPv4-Unicast IPv4-MPLS IPv4-VPN
NLRI(s) that peer saved forwarding for        : IPv4-Unicast IPv4-MPLS IPv4-VPN
NLRI(s) that restart is negotiated for        : None
NLRI(s) of received end-of-rib markers        : IPv4-Unicast
NLRI(s) of all end-of-rib markers sent        : IPv4-Unicast
Restart time locally configured for peer      : 120 seconds
Restart time requested by the peer            : 390 seconds
Time stale routes from peer are kept for      : 360 seconds
Graceful restart status on the peer           : Not currently being helped
Number of Restarts                            : 328
Last Restart at                               : 08/20/2006 12:22:06
===============================================================================
*A:7705_ALU-2>show>router>bgp#
```

Table 5: BGP Neighbor (Graceful Restart) Output Fields

Label	Description
`BGP Neighbor`	The IP address of the BGP neighbor
`Graceful Restart locally configured for peer`	The configured state of graceful restart for the local router
`Peer's Graceful Restart feature`	The configured state of graceful restart for the peer router
`NLRI(s) that peer supports restart for`	The families supported by the peer router for graceful restart
`NLRI(s) that peer saved forwarding for`	The families for which the peer router continued to forward packets after graceful restart.
`NLRI(s) that restart is negotiated for`	The families that negotiate restart during graceful restart
`NLRI(s) of received end-of-rib markers`	The families for which end-of-RIB markers have been received
`NLRI(s) of all end-of-rib markers sent`	The families for which end-of-RIB markers have been sent
`Restart time locally configured for peer`	The length of time configured on the local router for the peer router’s graceful restart

Table 5: BGP Neighbor (Graceful Restart) Output Fields (Continued)

Label	Description
`Restart time requested by the peer`	The length of time requested by the peer router for graceful restart
`Time stale routes from peer are kept for`	The length of time that the local router continues to support stale routes
`Graceful restart status on the peer`	The status of graceful restart on the peer router
`Number of Restarts`	The number of restarts since graceful restart is enabled between peers
`Last Restart at`	The local time of the last graceful restart

next-hop

Syntax **next-hop** [*family*] [*ip-address*] [**detail**]

Context show>router>bgp

Description This command displays BGP next-hop information.

Parameters *family* — the type of routing information to be distributed by the BGP instance

Values **ipv4** — displays only those BGP peers that have the IPv4 family enabled and not those capable of exchanging IP-VPN routes

ip-address — displays the next hop information for the specified IP address.

Values *ipv4-address*: a.b.c.d (host bits must be 0)

detail — displays the more detailed version of the output

Output The following output is an example of BGP next-hop information, and Table 6 describes the fields.

Sample Output

```
A:7705_ALU-2>show>router>bgp# next-hop
===============================================================================
 BGP Router ID : 10.10.10.1        AS : 65000   Local AS : 65000
===============================================================================

===============================================================================
BGP Next Hop
===============================================================================
Next Hop                                                        Pref Owner
   Resolving Prefix                                                  Metric
   Resolved Next Hop                                                 Ref. Count
-------------------------------------------------------------------------------
10.10.10.12                                                     15   ISIS
   0.0.0.0/0                                                         10
   88.88.1.2                                                         592
10.10.10.12                                                     15   ISIS
   0.0.0.0/0                                                         10
   88.88.2.2                                                         592
27.0.0.1                                                        15   ISIS
   27.0.0.0/24                                                       20
   88.88.1.2                                                         8
27.0.0.1                                                        15   ISIS
   27.0.0.0/24                                                       20
   88.88.2.2                                                         8
-------------------------------------------------------------------------------
Next Hops : 2
===============================================================================
A:7705_ALU-2>show>router>bgp#
```

```
A:7705_ALU-2>show>router>bgp# next-hop 27.0.0.1
===============================================================================
 BGP Router ID : 10.10.10.1        AS : 65000   Local AS : 65000
===============================================================================

===============================================================================
BGP Next Hop
===============================================================================
Next Hop                                                        Pref Owner
   Resolving Prefix                                                  Metric
   Resolved Next Hop                                                 Ref. Count
-------------------------------------------------------------------------------
27.0.0.1                                                        15   ISIS
   27.0.0.0/24                                                       20
   88.88.1.2                                                         8
27.0.0.1                                                        15   ISIS
   27.0.0.0/24                                                       20
   88.88.2.2                                                         8
-------------------------------------------------------------------------------
Next Hops : 1
===============================================================================
A:7705_ALU-2>show>router>bgp#
```

```
A:7705_ALU-2>show>router>bgp# next-hop 27.0.0.1 detail
===============================================================================
 BGP Router ID : 10.10.10.1        AS : 65000   Local AS : 65000
===============================================================================

===============================================================================
BGP Next Hop
===============================================================================
Next Hop: 27.0.0.1
-------------------------------------------------------------------------------
Resolving Prefix : 27.0.0.0/24
Preference       : 15                    Metric           : 20
Reference Count  : 8                     Owner            : ISIS
Resolved Next Hop: 88.88.1.2
Egress Label     : N/A
Resolved Next Hop: 88.88.2.2
Egress Label     : N/A
Resolved Next Hop: 88.88.3.2
Egress Label     : N/A
-------------------------------------------------------------------------------
Next Hops : 1
===============================================================================
A:7705_ALU-2>show>router>bgp#
```

Table 6: Show BGP Next-Hop Output Fields

Label	Description
`BGP Router ID`	The local BGP router ID
`AS`	The configured autonomous system number
`Local AS`	The configured local AS setting. If not configured, then the value is the same as the AS.
`Next Hop`	The next-hop address
`Resolving Prefix`	The prefix of the best next hop
`Pref` `Preference`	The BGP preference attribute for the routes
`Metric`	The metric derived from the IGP for a particular next hop
`Reference Count`	The number of routes using the resolving prefix
`Owner`	The routing protocol used to derive the best next hop
`Resolved Next Hop`	The IP address of the next hop
`Egress Label`	The VPN label used for VPN-IPv4 data
`Next Hops`	The number of next hops

paths

Syntax **paths**

Context show>router>bgp

Description This command displays a summary of BGP path attributes.

Output The following output is an example of BGP path information, and Table 7 describes the fields.

Sample Output

```
*A:7705_ALU-2>show>router>bgp# paths
===============================================================================
 BGP Router ID : 55.55.55.55 AS : 65000   Local AS : 65000
===============================================================================

===============================================================================
BGP Paths
===============================================================================
Path: No As-Path
-------------------------------------------------------------------------------
Next Hop         : 44.44.10.12
Origin           : Incomplete           Segments         : 0
MED              : None                 Local Preference : 4294967295
Refs             : 1080                 ASes             : 0
Flags            : IBGP-learned
-------------------------------------------------------------------------------
Path: No As-Path
-------------------------------------------------------------------------------
Next Hop         : 88.88.1.2
Origin           : IGP                  Segments         : 0
MED              : 10                   Local Preference : None
Refs             : 4                    ASes             : 0
Flags            : Imported
-------------------------------------------------------------------------------
Path: No As-Path
-------------------------------------------------------------------------------
Next Hop         : 44.44.10.21
Origin           : IGP                  Segments         : 0
MED              : None                 Local Preference : 100
Refs             : 1082                 ASes             : 0
Flags            : IBGP-learned
Cluster          : 10.10.10.12
Originator Id    : 10.10.10.21
-------------------------------------------------------------------------------
Paths : 3
===============================================================================
*A:7705_ALU-2>show>router>bgp#
```

Table 7: Show BGP Path Output Fields

Label	Description
BGP Router ID	The local BGP router ID
AS	The configured autonomous system number
Local AS	The configured local AS setting. If not configured, then the value is the same as the AS.
Path	The AS path attribute
Next Hop	The advertised BGP next hop
Origin	EGP — the NLRI is learned by an EGP protocol
	IGP — the NLRI is interior to the originating AS
	Incomplete — NLRI was learned another way
Segments	The number of segments in the AS path attribute
MED	The Multi-Exit Discriminator value
Local Preference	The local preference value. This value is used if the BGP route arrives from a BGP peer without the Local Pref attribute set. It is overridden by any value set via a route policy.
Refs	The number of routes using a specified set of path attributes
ASes	The number of autonomous system numbers in the AS path attribute
Flags	IBGP-learned — path attributes learned by an IBGP peering
Community	The BGP community attribute list
Cluster List	The route reflector cluster list
Originator ID	The originator ID path attribute value

routes

Syntax **routes** [*family*] [**brief**]
routes [*family*] *prefix* [**detail | longer | hunt** [**brief**]]
routes [*family*] **community** *comm-id*
routes [*family*] **aspath-regex** *reg-ex*

Context show>router>bgp

Description This command displays BGP route information.

When this command is issued without any parameters, the entire BGP routing table displays.

When this command is issued with an IP prefix/mask or IP address, the best match for the parameter displays.

Parameters *family* — the type of routing information to be distributed by the BGP instance

Values **ipv4** — displays only those BGP peers that have the IPv4 family enabled and not those capable of exchanging IP-VPN routes

vpn-ipv4 — displays the BGP peers that are IP-VPN capable

prefix — the type of routing information to display

Values [*rd*]:[*ip-address*[/*mask*]]

rd	*ip-address*:*number1* *as-number1*:*number2* *as-number2*:*number3*
number1	1 to 65535
as-number1	1 to 65535
number2	0 to 4294967295
as-number2	1 to 4294967295
number3	0 to 65535
ip-address	a.b.c.d
mask	0 to 32

brief — provides a summarized display of the set of peers to which a BGP route is advertised

hunt — displays entries for the specified route in the RIB-In, RIB-Out, and RTM

longer — displays the specified route and subsets of the route

detail — displays the more detailed version of the output

community *comm-id* **—** displays all routes with the specified BGP community

Values

comm-id	[*as-number1*:*comm-val1* | *ext-comm* | *well-known-comm*]
ext-comm	*type*:*ip-address*:*comm-val1* | *as-number1*:*comm-val2* | *as- number2*:*comm-val1*}
as-number1	0 to 65535
comm-val1	0 to 65535
type	**target**, **origin** (keywords)
ip-address	a.b.c.d

comm-val2	0 to 4294967295
as-number2	0 to 4294967295
well-known-comm	**no-export**, **no-export-subconfed**, **no-advertise** (keywords)

aspath-regex *reg-exp* — displays all routes with an AS path matching the specified regular expression *reg-exp*

Output The following output is an example of BGP route information, and Table 8 describes the fields.

Sample Output

```
*A:7705_ALU-2>show>router>bgp# routes
===============================================================================
 BGP Router ID : 10.10.10.1        AS : 65000   Local AS : 65000
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best

===============================================================================
BGP IPv4 Routes
===============================================================================
Flag  Network                                           LocalPref   MED
      Nexthop                                                       VPN Label
      As-Path
-------------------------------------------------------------------------------
u*>?  10.10.10.12/32                                    100         None
      10.10.10.12                                                   -
      No As-Path

u*>?  11.11.1.0/24                                      100         None
      10.10.10.12                                                   -
      No As-Path

*?    23.0.0.0/24                                       100         None
      10.10.10.12                                                   -
      No As-Path

*?    24.0.0.0/24                                       100         None
      10.10.10.12                                                   -
      No As-Path
-------------------------------------------------------------------------------
Routes : 4
===============================================================================
*A:7705_ALU-2>show>router>bgp#


*A:7705_ALU-2>show>router>bgp# routes brief
===============================================================================
 BGP Router ID : 10.10.10.1        AS : 65000   Local AS : 65000
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best

===============================================================================
BGP IPv4 Routes
===============================================================================
Flag  Network
-------------------------------------------------------------------------------
u*>?  10.10.10.12/32
u*>?  11.11.1.0/24
*?    23.0.0.0/24
*?    24.0.0.0/24
-------------------------------------------------------------------------------
Routes : 4
===============================================================================
*A:7705_ALU-2>show>router>bgp#
```

```
*A:7705_ALU-2>show>router>bgp# routes 13.1.0.0/24 detail
===============================================================================
BGP Router ID : 10.128.0.161 AS : 65535 Local AS : 65535
===============================================================================
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best
===============================================================================
BGP Routes
===============================================================================
Original Attributes
Network        : 13.1.0.0/24         Nexthop : 10.20.1.20
Route Dist.    : 10070:100           VPN Label : 152784
From           : 10.20.1.20          Res. Nexthop: 10.130.0.2
Local Pref.    : 100
Aggregator AS  : none                Aggregator: none
Atomic Aggr.   : Not Atomic          MED : none
Community      : target:10070:1
Cluster        : No Cluster Members
Originator Id  : None                Peer Router Id: 10.20.1.20
Flags          : Used Valid Best IGP
AS-Path        : 10070 {14730}

Modified Attributes

Network        : 13.1.0.0/24         Nexthop : 10.20.1.20
Route Dist.    : 10001:100           VPN Label : 152560
From           : 10.20.1.20          Res. Nexthop : 10.130.0.2
Local Pref.    : 100
Aggregator AS  : none                Aggregator: none
Atomic Aggr.   : Not Atomic          MED : none
Community      : target:10001:1
Cluster        : No Cluster Members
Originator Id  : None                Peer Router Id: 10.20.1.20
Flags          : Used Valid Best IGP
AS-Path        : No As-Path
-------------------------------------------------------------------------------
...
===============================================================================
A*A:7705_ALU-2>show>router>bgp#
```

```
*A:7705_ALU-2> show router bgp routes 100.0.0.0/30 hunt
===============================================================================
BGP Router ID : 10.20.1.1 AS : 100Local AS : 100
===============================================================================
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best
===============================================================================
BGP Routes
===============================================================================
RIB In Entries
-------------------------------------------------------------------------------
Network        : 100.0.0.0/30
Nexthop        : 10.20.1.2
Route Dist.    : 10.20.1.2:1         VPN Label: 131070
From           : 10.20.1.2
Res. Nexthop   : 10.10.1.2
Local Pref.    : 100                 Interface Name: to-sr7
Aggregator AS  : none                Aggregator: none
Atomic Aggr.   : Not Atomic          MED: none
Community      : target:10.20.1.2:1
Cluster        : No Cluster Members
Originator Id  : None                Peer Router Id: 10.20.1.2
Flags          : Used Valid Best IGP
AS-Path        : No As-Path
VPRN Imported  : 1 2 10 12
-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
Routes : 1
===============================================================================
*A:7705_ALU-2
```

Table 8: Show BGP Route Output Fields

Label	Description
`BGP Router ID`	The local BGP router ID
`AS`	The configured autonomous system number
`Local AS`	The configured local AS setting. If not configured, then the value is the same as the AS.
`Flag(s)`	Legend: Status codes: u - used s - suppressed h - history d - decayed * - valid If an * is not present, then the status is invalid Origin codes: i - IGP e - EGP ? - incomplete > - best
`Network`	The IP prefix and mask length
`Nexthop`	The BGP next hop
`AS-Path`	The BGP AS path attribute
`Local Pref.`	The local preference value. This value is used if the BGP route arrives from a BGP peer without the Local Pref attribute set. It is overridden by any value set via a route policy.
`MED`	The MED metric value
	`none` — MED metrics are present
`VPN Label`	The label generated by the PE's label manager
`Original Attributes`	The received BGP attributes of a route from a peer without any modification from any policy
`Modified Attributes`	The final BGP attributes of a route after the policies evaluation
`Route Dist.`	The route distinguisher identifier attached to routes that distinguishes the VPN it belongs to
`From`	The advertising BGP neighbor's IP address
`Res. Nexthop`	The resolved next hop

Table 8: Show BGP Route Output Fields (Continued)

Label	Description
Aggregator AS	The aggregator AS value
	none — aggregator AS attributes are not present
Aggregator	The aggregator attribute value
	none — aggregator attributes are not present
Atomic Aggr.	Atomic — the atomic aggregator flag is set
	Not Atomic — the atomic aggregator flag is not set
Community	The BGP community attribute list
Cluster	The route reflector cluster list
Originator Id	The originator ID path attribute value
	none — the originator ID attribute is not present
Peer Router Id	The router ID of the advertising router
VPRN Imported	The VPRNs where a particular BGP-VPN received route has been imported and installed

summary

Syntax **summary** [**all**]
summary [**family** *family*] [**neighbor** *ip-address*]

Context show>router>bgp

Description This command displays a summary of BGP neighbor information.

If confederations are not configured, that portion of the output will not display.

The “State” field displays the global BGP operational state. The valid values are:

Up — BGP global process is configured and running

Down — BGP global process is administratively shut down and not running

Disabled — BGP global process is operationally disabled. The process must be restarted by the operator.

For example, if a BGP peer is operationally disabled, then the state in the summary table shows the state ‘Disabled’.

Parameters *family* — the type of routing information to be distributed by the BGP instance

Values **ipv4** — displays only those BGP peers that have the IPv4 family enabled

vpn-ipv4 — displays the BGP peers that are IP-VPN capable

ip-address **—** clears damping information for entries received from the BGP neighbor

Values *ipv4-address*: a.b.c.d

Output The following output is an example of BGP summary information, and Table 9 describes the fields.

Sample Output

```
*A:7705_ALU-2>show>router>bgp# summary
===============================================================================
 BGP Router ID : 55.55.55.1        AS : 65000   Local AS : 65000
===============================================================================
BGP Admin State         : Up          BGP Oper State              : Up
Total Peer Groups       : 1           Total Peers                 : 1
Total BGP Paths         : 74          Total Path Memory           : 9128
Total IPv4 Remote Rts   : 600         Total IPv4 Rem. Active Rts  : 563
Total Supressed Rts     : 0           Total Hist. Rts             : 0
Total Decay Rts         : 0

Total VPN Peer Groups   : 0           Total VPN Peers             : 0
Total VPN Local Rts     : 8672
Total VPN-IPv4 Rem. Rts : 8656        Total VPN-IPv4 Rem. Act. Rts: 8656
Total VPN Supp. Rts     : 0           Total VPN Hist. Rts         : 0
Total VPN Decay Rts     : 0

===============================================================================
BGP Summary
===============================================================================
Neighbor
                   AS PktRcvd InQ  Up/Down   State|Rcv/Act/Sent (Addr Family)
                      PktSent OutQ
-------------------------------------------------------------------------------
44.44.10.12
                65000     654    0 04h11m01s 600/563/569 (IPv4)
                          557    0           8656/8656/8672 (VpnIPv4)
===============================================================================
*A:7705_ALU-2>show>router>bgp#


*A:7705_ALU-2>show>router>bgp# summary all
===============================================================================
BGP Summary
===============================================================================
Neighbor
ServiceId          AS PktRcvd InQ  Up/Down   State|Rcv/Act/Sent (Addr Family)
                      PktSent OutQ
-------------------------------------------------------------------------------
44.44.10.12
Def. Instance   65000     662    0 04h14m52s 600/563/569 (IPv4)
                          564    0           8656/8656/8672 (VpnIPv4)

===============================================================================
*A:7705_ALU-2>show>router>bgp#
```

```
*A:7705_ALU-2>show>router>bgp# summary neighbor 44.44.10.12
===============================================================================
 BGP Router ID : 44.44.10.1        AS : 65000   Local AS : 65000
===============================================================================
BGP Admin State         : Up          BGP Oper State              : Up
Total Peer Groups       : 1           Total Peers                 : 1
Total BGP Paths         : 74          Total Path Memory           : 9128
Total IPv4 Remote Rts   : 600         Total IPv4 Rem. Active Rts  : 563
Total Supressed Rts     : 0           Total Hist. Rts             : 0
Total Decay Rts         : 0

Total VPN Peer Groups   : 0           Total VPN Peers             : 0
Total VPN Local Rts     : 8672
Total VPN-IPv4 Rem. Rts : 8656        Total VPN-IPv4 Rem. Act. Rts: 8656
Total VPN Supp. Rts     : 0           Total VPN Hist. Rts         : 0
Total VPN Decay Rts     : 0

===============================================================================
BGP Summary
===============================================================================
Neighbor
                   AS PktRcvd InQ  Up/Down   State|Rcv/Act/Sent (Addr Family)
                      PktSent OutQ
-------------------------------------------------------------------------------
44.44.10.12
                65000     673    0 04h20m24s 600/563/569 (IPv4)
                          575    0           8656/8656/8672 (VpnIPv4)
===============================================================================
*A:7705_ALU-2>show>router>bgp#
```

```
*A:7705_ALU-2>show>router>bgp# summary family ipv4
===============================================================================
 BGP Router ID : 44.44.10.1        AS : 65000   Local AS : 65000
===============================================================================
BGP Admin State         : Up          BGP Oper State              : Up
Total Peer Groups       : 1           Total Peers                 : 1
Total BGP Paths         : 74          Total Path Memory           : 9128
Total IPv4 Remote Rts   : 600         Total IPv4 Rem. Active Rts  : 563
Total Supressed Rts     : 0           Total Hist. Rts             : 0
Total Decay Rts         : 0

Total VPN Peer Groups   : 0           Total VPN Peers             : 0
Total VPN Local Rts     : 8672
Total VPN-IPv4 Rem. Rts : 8656        Total VPN-IPv4 Rem. Act. Rts: 8656
Total VPN Supp. Rts     : 0           Total VPN Hist. Rts         : 0
Total VPN Decay Rts     : 0

===============================================================================
BGP IPv4 Summary
===============================================================================
Neighbor
                   AS PktRcvd PktSent  InQ OutQ Up/Down   State|Recv/Actv/Sent
-------------------------------------------------------------------------------
44.44.10.12
                65000     679     581    0    0 04h23m36s 600/563/569
===============================================================================
*A:7705_ALU-2>show>router>bgp#
```

Table 9: Show BGP Summary Output Fields

Label	Description
`BGP Router ID`	The local BGP router ID
`AS`	The configured autonomous system number
`Local AS`	The configured local AS setting. If not configured, then the value is the same as the AS.
`BGP Admin State`	`Down` — BGP is administratively disabled
	`Up` — BGP is administratively enabled
`BGP Oper State`	`Down` — BGP is operationally disabled
	`Up` — BGP is operationally enabled
`Total Peer Groups`	The total number of configured BGP peer groups
`Total Peers`	The total number of configured BGP peers
`Total BGP Paths`	The total number of unique sets of BGP path attributes learned from BGP peers
`Total Path Memory`	The total amount of memory used to store the path attributes
`Total IPv4 Remote Rts`	The total number of IPv4 routes learned from BGP peers
`Total IPv4 Remote Act. Rts`	The total number of IPv4 routes used in the forwarding table
`Total Suppressed Rts`	The total number of suppressed routes due to route damping
`Total Hist. Rts`	The total number of routes with history due to route damping
`Total Decay Rts`	The total number of decayed routes due to route damping
`Total VPN Peer Groups`	The total number of configured VPN peer groups
`Total VPN Peers`	The total number of configured VPN peers
`Total VPN Local Rts`	The total number of configured local VPN routes
`Total VPN-IPv4 Rem. Rts`	The total number of configured remote VPN-IPv4 routes
`Total VPN-IPv4 Rem. Act. Rts`	The total number of active remote VPN-IPv4 routes used in the forwarding table
`Total VPN Supp. Rts`	The total number of suppressed VPN routes due to route damping

Table 9: Show BGP Summary Output Fields (Continued)

Label	Description
`Total VPN Hist. Rts`	The total number of VPN routes with history due to route damping
`Total VPN Decay Rts`	The total number of decayed routes due to route damping
`Neighbor`	The BGP neighbor address
`AS (Neighbor)`	The BGP neighbor autonomous system number
`PktRcvd`	The total number of packets received from the BGP neighbor
`PktSent`	The total number of packets sent to the BGP neighbor
`InQ`	The number of BGP messages to be processed
`OutQ`	The number of BGP messages to be transmitted
`Up/Down`	The amount of time that the BGP neighbor has either been established or not established depending on its current state
`State|Recv/Actv/Sent (Addr Family)`	The BGP neighbor’s current state (if not established) or the number of received routes, active routes and sent routes (if established), along with the address family

Clear Commands

damping

Syntax **damping** [{*ip-prefix/ip-prefix-length*] [**neighbor** *ip-address*]} | {[**group** *name*}]

Context clear>router>bgp

Description This command clears or resets the route damping information for received routes.

Parameters *ip-prefix/ip-prefix-length* — clears damping information for entries that match the IP prefix and prefix length

Values	*ipv4-prefix*:	a.b.c.d (host bits must be 0)
	ipv4-prefix-length:	0 to 32

ip-address — clears damping information for entries received from the BGP neighbor

Values	*ipv4-address*:	a.b.c.d

name — clears damping information for entries received from any BGP neighors in the peer group

Values 32 characters maximum

flap-statistics

Syntax **flap-statistics** [{*ip-prefix/mask*] [**neighbor** *ip-address*] | [**group** *group-name*] | [**regex** *reg-exp*] | [**policy** *policy-name*}]

Context clear>router>bgp

Description This command clears route flap statistics.

Parameters *ip-prefix/mask* — clears route flap statistics for entries that match the specified IP prefix and mask length

Values	*ip-prefix*:	a.b.c.d (host bits must be 0)
	mask:	0 to 32

ip-address — clears route flap statistics for entries received from the specified BGP neighbor

Values	*ipv4-address*:	a.b.c.d

group-name — clears route flap statistics for entries received from any BGP neighbors in the specified peer group

reg-exp — clears route flap statistics for all entries that have the regular expression and the AS path that matches the regular expression

policy-name — clears route flap statistics for entries that match the specified route policy

neighbor

Syntax **neighbor** {*ip-address* | **as** *as-number* | **external** | **all**} [**soft** | **soft-inbound**]
neighbor {*ip-address* | **as** *as-number* | **external** | **all**} **statistics**
neighbor *ip-address* **end-of-rib**

Context clear>router>bgp

Description This command resets the specified BGP peer or peers. This can cause existing BGP connections to be shut down and restarted.

Parameters *ip-address* — resets the BGP neighbor with the specified IP address

Values *ipv4-address*: a.b.c.d

as-number — resets all BGP neighbors with the specified peer AS

Values 1 to 65535

external — resets all EBGP neighbors

all — resets all BGP neighbors

soft — the specified BGP neighbor(s) re-evaluates all routes in the Local-RIB against the configured export policies

soft-inbound — the specified BGP neighbor(s) re-evaluates all routes in the RIB-In against the configured import policies

statistics — the BGP neighbor statistics

end-of-rib — clears the routing information base (RIB)

protocol

Syntax **protocol**

Context clear>router>bgp

Description This command resets the entire BGP protocol.

Debug Commands

events

Syntax **events** [**neighbor** *ip-address* | **group** *name*]
no events

Context debug>router>bgp

Description This command logs all events changing the state of a BGP peer.

Parameters *ip-address* — debugs only events affecting the specified BGP neighbor

Values *ipv4-address*: a.b.c.d (host bits must be 0)

name — debugs only events affecting the specified peer group and associated neighbors

graceful-restart

Syntax **graceful-restart** [**neighbor** *ip-address* | **group** *name*]
no graceful-restart

Context debug>router>bgp

Description This command enables debugging for BGP graceful restart.

The **no** form of the command disables the debugging.

Parameters *ip-address* — debugs only events affecting the specified BGP neighbor

Values *ipv4-address*: a.b.c.d (host bits must be 0)

name — debugs only events affecting the specified peer group and associated neighbors

keepalive

Syntax **keepalive** [**neighbor** *ip-address* | **group** *name*]
no keepalive

Context debug>router>bgp

Description This command decodes and logs all sent and received Keepalive messages in the debug log.

Parameters *ip-address* — debugs only events affecting the specified BGP neighbor

Values *ipv4-address*: a.b.c.d (host bits must be 0)

name — debugs only events affecting the specified peer group and associated neighbors

notification

Syntax **notification** [**neighbor** *ip-address* | **group** *name*]
no notification

Context debug>router>bgp

Description This command decodes and logs all sent and received Notification messages in the debug log.

Parameters *ip-address* — debugs only events affecting the specified BGP neighbor

Values *ipv4-address*: a.b.c.d (host bits must be 0)

name — debugs only events affecting the specified peer group and associated neighbors

open

Syntax **open** [**neighbor** *ip-address* | **group** *name*]
no open

Context debug>router>bgp

Description This command decodes and logs all sent and received Open messages in the debug log.

Parameters *ip-address* — debugs only events affecting the specified BGP neighbor

name — debugs only events affecting the specified peer group and associated neighbors

outbound-route-filtering

Syntax [**no**] **outbound-route-filtering**

Context debug>router>bgp

Description This command enables debugging for all BGP outbound route filtering (ORF) packets. ORF is used to inform a neighbor of targets (using target-list) that it is willing to receive.

packets

Syntax **packets** [**neighbor** *ip-address* **|** **group** *name*]
packets

Context debug>router>bgp

Description This command decodes and logs all sent and received BGP packets in the debug log.

Parameters *ip-address* — debugs only events affecting the specified BGP neighbor

Values *ipv4-address*: a.b.c.d (host bits must be 0)

name — debugs only events affecting the specified peer group and associated neighbors

route-refresh

Syntax **route-refresh** [**neighbor** *ip-address* **|** **group** *name*]
no route-refresh

Context debug>router>bgp

Description This command enables and disables debugging for BGP route refresh.

Parameters *ip-address* — debugs only events affecting the specified BGP neighbor

Values *ipv4-address*: a.b.c.d (host bits must be 0)

name — debugs only events affecting the specified peer group and associated neighbors

rtm

Syntax **rtm** [**neighbor** *ip-address* **|** **group** *name*]
no rtm

Context debug>router>bgp

Description This command logs RTM changes in the debug log.

Parameters *ip-address* — debugs only events affecting the specified BGP neighbor

Values *ipv4-address*: a.b.c.d (host bits must be 0)

name — debugs only events affecting the specified peer group and associated neighbors

socket

Syntax **socket** [**neighbor** *ip-address* | **group** *name*]
no socket

Context debug>router>bgp

Description This command logs all TCP socket events to the debug log.

Parameters *ip-address* — debugs only events affecting the specified BGP neighbor

Values *ipv4-address*: a.b.c.d (host bits must be 0)

name — debugs only events affecting the specified peer group and associated neighbors

timers

Syntax **timers** [**neighbor** *ip-address* | **group** *name*]
no timers

Context debug>router>bgp

Description This command logs all BGP timer events to the debug log.

Parameters *ip-address* — debugs only events affecting the specified BGP neighbor

Values *ipv4-address*: a.b.c.d (host bits must be 0)

name — debugs only events affecting the specified peer group and associated neighbors

update

Syntax **update** [**neighbor** *ip-address* | **group** *name*]
no update

Context debug>router>bgp

Description This command decodes and logs all sent and received Update messages in the debug log.

Parameters *ip-address* — debugs only events affecting the specified BGP neighbor

Values *ipv4-address*: a.b.c.d (host bits must be 0)

name — debugs only events affecting the specified peer group and associated neighbors

Standards and Protocol Support

Standards Compliance

IEEE 802.1ag	Service Layer OAM
IEEE 802.1p/q	VLAN Tagging
IEEE 802.3	10BaseT
IEEE 802.3ah	Ethernet OAM
IEEE 802.3u	100BaseTX
IEEE 802.3x	Flow Control
IEEE 802.3z	1000BaseSX/LX
IEEE 802.3-2008	Revised base standard
ITU-T Y.1731	OAM functions and mechanisms for Ethernet-based networks

Telecom Compliance

IC CS-03 Issue 9	Spectrum Management and Telecommunications
ACTA TIA-968-A	
AS/ACIF S016 (Australia/New Zealand)	Requirements for Customer Equipment for connection to hierarchical digital interfaces
ITU-T G.703	Physical/electrical characteristics of hierarchical digital interfaces
ITU-T G.707	Network node interface for the Synchronous Digital Hierarchy (SDH)
ITU-T G.712-2001	Transmission performance characteristics of pulse code modulation channels
ITU-T G.957	Optical interfaces for equipments and systems relating to the synchronous digital hierarchy
ITU-T V.24	List of definitions for interchange circuits between data terminal equipment (DTE) and data circuit-terminating equipment (DCE)
ITU-T V.36	Modems for synchronous data transmission using 60-108 kHz group band circuits

Protocol Support

ATM

RFC 2514 Definitions of Textual Conventions and OBJECT_IDENTITIES for ATM Management, February 1999

RFC 2515 Definition of Managed Objects for ATM Management, February 1999

RFC 2684 Multiprotocol Encapsulation over ATM Adaptation Layer 5

af-tm-0121.000 Traffic Management Specification Version 4.1, March 1999

ITU-T Recommendation I.610 - B-ISDN Operation and Maintenance Principles and Functions version 11/95

ITU-T Recommendation I.432.1 - B-ISDN user-network interface - Physical layer specification: General characteristics

GR-1248-CORE - Generic Requirements for Operations of ATM Network Elements (NEs). Issue 3 June 1996

GR-1113-CORE - Bellcore, Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1, July 1994

AF-PHY-0086.001 Inverse Multiplexing for ATM (IMA)

BFD

draft-ietf-bfd-mib-00.txt Bidirectional Forwarding Detection Management Information Base

draft-ietf-bfd-base-o5.txt Bidirectional Forwarding Detection

draft-ietf-bfd-v4v6-1hop-06.txt BFD IPv4 and IPv6 (Single Hop)

draft-ietf-bfd-multihop-06.txt BFD for Multi-hop Paths

BGP

RFC 1397 BGP Default Route Advertisement
RFC 1997 BGP Communities Attribute
RFC 2385 Protection of BGP Sessions via MDS
RFC 2439 BGP Route Flap Dampening
RFC 2547bis BGP/MPLS VPNs
RFC 2918 Route Refresh Capability for BGP-4
RFC 3107 Carrying Label Information in BGP-4
RFC 3392 Capabilities Advertisement with BGP-4
RFC 4271 BGP-4 (previously RFC 1771)
RFC 4360 BGP Extended Communities Attribute
RFC 4364 BGP/MPLS IP Virtual Private Networks (VPNs) (previously RFC 2574bis BGP/MPLS VPNs)
RFC 4456 BGP Route Reflection: Alternative to Full-mesh IBGP (previously RFC 1966 and RFC 2796)
RFC 472 Graceful Restart Mechanism for BGP - GR Helper
RFC 4760 Multi-protocol Extensions for BGP (previously RFC 2858)

DHCP

RFC 1534 Interoperation between DHCP and BOOTP
RFC 2131 Dynamic Host Configuration Protocol (REV)
RFC 3046 DHCP Relay Agent Information Option (Option 82)

DIFFERENTIATED SERVICES

RFC 2474 Definition of the DS Field in the IPv4 and IPv6 Headers
RFC 2597 Assured Forwarding PHB Group
RFC 2598 An Expedited Forwarding PHB
RFC 3140 Per-Hop Behavior Identification Codes

DIGITAL DATA NETWORK MANAGEMENT

V.35
RS-232 (also known as EIA/TIA-232)

GRE

RFC 2784 Generic Routing Encapsulation (GRE)

LDP

RFC 5036 LDP Specification

IS-IS

RFC 1142 OSI IS-IS Intra-domain Routing Protocol (ISO 10589)
RFC 1195 Use of OSI IS-IS for routing in TCP/IP & dual environments
RFC 2763 Dynamic Hostname Exchange for IS-IS
RFC 2966 Domain-wide Prefix Distribution with Two-Level IS-IS
RFC 2973 IS-IS Mesh Groups
RFC 3373 Three-Way Handshake for Intermediate System to Intermediate System (IS-IS) Point-to-Point Adjacencies
RFC 3567 Intermediate System to Intermediate System (IS-IS) Cryptographic Authentication
RFC 3719 Recommendations for Interoperable Networks using IS-IS
RFC 3784 Intermediate System to Intermediate System (IS-IS) Extensions for Traffic Engineering (TE)
RFC 3787 Recommendations for Interoperable IP Networks
RFC 4205 for Shared Risk Link Group (SRLG) TLV draft-ietf-isis-igp-p2p-over-lan-05.txt
RFC 5309 Point-to-Point Operation over LAN in Link State Routing Protocols

MPLS

RFC 3031 MPLS Architecture
RFC 3032 MPLS Label Stack Encoding
RFC 3815 Definitions of Managed Objects for the Multiprotocol Label Switching (MPLS), Label Distribution Protocol (LDP)
RFC 4379 Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures

NETWORK MANAGEMENT

ITU-T X.721: Information technology- OSI-Structure of Management Information
ITU-T X.734: Information technology- OSI-Systems Management: Event Report Management Function
M.3100/3120 Equipment and Connection Models
TMF 509/613 Network Connectivity Model
RFC 1157 SNMPv1
RFC 1305 Network Time Protocol (Version 3) Specification, Implementation and Analysis
RFC 1850 OSPF-MIB
RFC 1907 SNMPv2-MIB
RFC 2011 IP-MIB
RFC 2012 TCP-MIB

RFC 2013 UDP-MIB

RFC 2030 Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI

RFC 2096 IP-FORWARD-MIB

RFC 2138 RADIUS

RFC 2206 RSVP-MIB

RFC 2571 SNMP-FRAMEWORKMIB

RFC 2572 SNMP-MPD-MIB

RFC 2573 SNMP-TARGET-&-NOTIFICATION-MIB

RFC 2574 SNMP-USER-BASED-SMMIB

RFC 2575 SNMP-VIEW-BASED ACM-MIB

RFC 2576 SNMP-COMMUNITY-MIB

RFC 2588 SONET-MIB

RFC 2665 EtherLike-MIB

RFC 2819 RMON-MIB

RFC 2863 IF-MIB

RFC 2864 INVERTED-STACK-MIB

RFC 3014 NOTIFICATION-LOG MIB

RFC 3164 The BSD Syslog Protocol

RFC 3273 HCRMON-MIB

RFC 3411 An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks

RFC 3412 Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)

RFC 3413 Simple Network Management Protocol (SNMP) Applications

RFC 3414 User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)

RFC 3418 SNMP MIB

draft-ietf-disman-alarm-mib-04.txt

draft-ietf-mpls-ldp-mib-07.txt

draft-ietf-ospf-mib-update-04.txt

draft-ietf-mpls-lsr-mib-06.txt

draft-ietf-mpls-te-mib-04.txt

IANA-IFType-MIB

OSPF

RFC 1765 OSPF Database Overflow

RFC 2328 OSPF Version 2

RFC 2370 Opaque LSA Support

RFC 3101 OSPF NSSA Option

RFC 3137 OSPF Stub Router Advertisement

RFC 3630 Traffic Engineering (TE) Extensions to OSPF

PPP

RFC 1332 PPP Internet Protocol Control Protocol (IPCP)

RFC 1570 PPP LCP Extensions

RFC 1619 PPP over SONET/SDH

RFC 1661 The Point-to-Point Protocol (PPP)

RFC 1662 PPP in HDLC-like Framing

RFC 1989 PPP Link Quality Monitoring

RFC 1990 The PPP Multilink Protocol (MP)

RFC 2686 The Multi-Class Extension to Multi-Link PPP

PSEUDOWIRES

RFC 3550 RTP: A Transport Protocol for Real-Time Applications

RFC 3985 Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture

RFC 4385 Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN

RFC 4446 IANA Allocation for PWE3

RFC 4447 Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)

RFC 4448 Encapsulation Methods for Transport of Ethernet over MPLS Networks

RFC 4553 Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)

RFC 4717 Encapsulation Methods for Transport of Asynchronous Transfer Mode (ATM) over MPLS Networks

RFC 5085 Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires

RFC 5086 Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)

draft-ietf-pwe3-redundancy-02 Pseudowire (PW) Redundancy

RADIUS

RFC 2865 Remote Authentication Dial In User Service

RFC 2866 RADIUS Accounting

RSVP-TE and FRR

RFC 2430 A Provider Architecture for DiffServ & TE

RFC 2961 RSVP Refresh Overhead Reduction Extensions

RFC 2702 Requirements for Traffic Engineering over MPLS

RFC 2747 RSVP Cryptographic Authentication

RFC 3097 RSVP Cryptographic Authentication - Updated Message Type Value

RFC 3209 Extensions to RSVP for LSP Tunnels

RFC 3210 Applicability Statement for Extensions to RSVP for LSP Tunnels

RFC 4090 Fast Reroute Extensions to RSVP-TE for LSP Tunnels

SONET/SDH

GR-253-CORE SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000

ITU-T Recommendation G.841 Telecommunication Standardization Section of ITU, Types and Characteristics of SDH Networks Protection Architecture, issued in October 1998 and as augmented by Corrigendum1 issued in July 2002

SSH

draft-ietf-secsh-architecture.txt SSH Protocol Architecture

draft-ietf-secsh-userauth.txt SSH Authentication Protocol

draft-ietf-secsh-transport.txt SSH Transport Layer Protocol

draft-ietf-secsh-connection.txt SSH Connection Protocol

draft-ietf-secsh- newmodes.txt SSH Transport Layer Encryption Modes

SYNCHRONIZATION

G.813 Timing characteristics of SDH equipment slave clocks (SEC)

G.8261 Timing and synchronization aspects in packet networks

G.8262 Timing characteristics of synchronous Ethernet equipment slave clock

GR 1244 CORE Clocks for the Synchronized Network: Common Generic Criteria

IEEE 1588v2 1588 PTP 2008

TACACS+

draft-grant-tacacs-02.txt The TACACS+ Protocol

TCP/IP

RFC 768 User Datagram Protocol

RFC 791 Internet Protocol

RFC 792 Internet Control Message Protocol

RFC 793 Transmission Control Protocol

RFC 826 Ethernet Address Resolution Protocol

RFC 854 Telnet Protocol Specification

RFC 1350 The TFTP Protocol (Rev. 2)

RFC 1812 Requirements for IPv4 Routers

Proprietary MIBs

TIMETRA-ATM-MIB.mib
TIMETRA-CAPABILITY-7705-V1.mib
TIMETRA-CFLOWD-MIB.mib
TIMETRA-CHASSIS-MIB.mib
TIMETRA-CLEAR-MIB.mib
TIMETRA-FILTER-MIB.mib
TIMETRA-GLOBAL-MIB.mib
TIMETRA-LDP-MIB.mib
TIMETRA-LOG-MIB.mib
TIMETRA-MPLS-MIB.mib
TIMETRA-OAM-TEST-MIB.mib
TIMETRA-PORT-MIB.mib
TIMETRA-PPP-MIB.mib
TIMETRA-QOS-MIB.mib
TIMETRA-ROUTE-POLICY-MIB.mib
TIMETRA-RSVP-MIB.mib
TIMETRA-SAP-MIB.mib
TIMETRA-SDP-MIB.mib
TIMETRA-SECURITY-MIB.mib
TIMETRA-SERV-MIB.mib
TIMETRA-SYSTEM-MIB.mib
TIMETRA-TC-MIB.mib

Customer documentation and product support

Customer documentation

http://www.alcatel-lucent.com/myaccess

Product manuals and documentation updates are available at alcatel-lucent.com. If you are a new user and require access to this service, please contact your Alcatel-Lucent sales representative.

Technical Support

http://www.alcatel-lucent.com/support

Documentation feedback

documentation.feedback@alcatel-lucent.com

Alcatel·Lucent