Introduction to OSPF

A.Prasad Babu

www.juniper.net

Module Objectives

- After completing this module, you should be able to describe
  - OSPF protocol basics
  - Each LSA type
  - Helpful CLI commands

Juniper Your Net

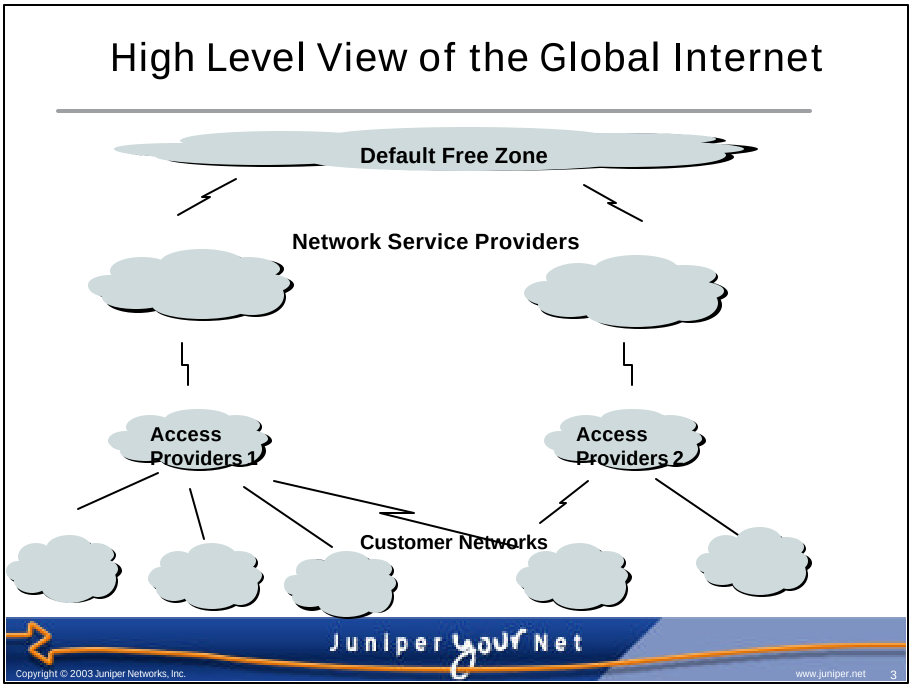
High Level View of the Global Internet
Default Free Zone
Network Service Providers
Access
Providers 1
Access
Providers 2
Customer Networks
Juniper your Net
Copyright © 2003 Juniper Networks, Inc.
www.juniper.net
3

Hierarchy of Routing Protocols within an AS

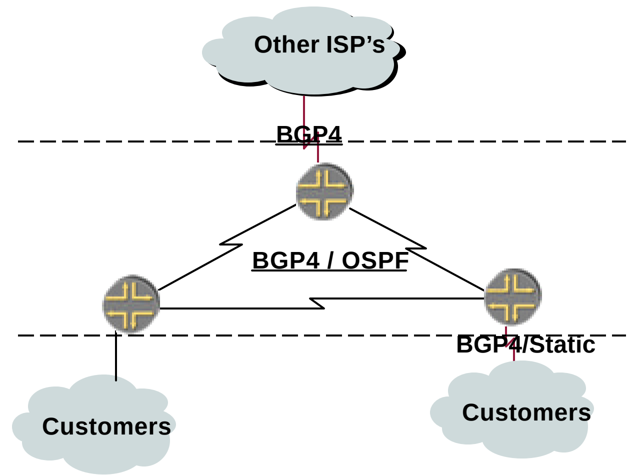

www.juniper.net 4

What Is an IGP?

- Interior Gateway Protocol
- Within an Autonomous System
- Carries information about internal prefixes
- Examples—OSPF, ISIS

Interior vs. Exterior – Routing Protocols

- Interior
  - Automatic discovery
  - Generally trust your IGP routers
  - Routes go to all IGP routers
- Exterior
  - Specifically configured peers
  - Connecting with outside networks
  - Set administrative boundaries

Juniper your Net

In the Beginning
Was Distance Vector...

- Also known as Bellman-Ford, Ford-Fulkerson
- Very simple algorithm
- Distance Vector Protocols include
  - RIP
  - BGP (but usually called Path Vector)
  - Cisco’s IGRP
  - Cisco’s EIGRP

Juniper your Net

www.juniper.net 7

Routing by Rumor

- Distributed calculation
- Each router knows only what its neighbor tells it

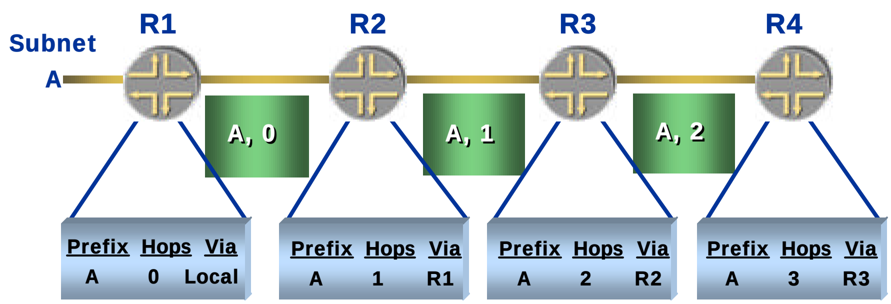

Problems with Distance Vector

- Slow convergence
  - A direct result of the distributed calculation
  - Triggered updates help
  - Kludges such as hold-down timers reduce transient errors, but increase convergence time
- Single-hop routing loops
  - Solution: split horizon
- Counting to infinity
  - Solution: make infinity finite
- Synchronized periodic updates
  - Solution: update jitter timers

Link-state Protocols

- Also known as shortest path, Dijkstra
- Algorithm based on graph theory, providing better loop avoidance
- Local computation means faster convergence
- Link-state protocols include
  - OSPF
  - IS-IS
  - ATM PNNI
  - IBM APPN
  - MPLS CSPF

Juniper Your Net

Fundamental Link-state Concepts

- Adjacency
- Information flooding
- Link-state database
- SPF calculation

Juniper your Net

SPF Calculation Example

Example Topology

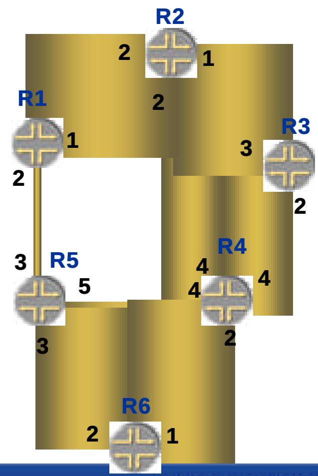

Link-state Database

R1-R2, 1
R1-R5, 2
R2-R1, 2
R2-R3, 1
R2-R4, 2
R3-R2, 3
R3-R4, 2
R4-R2, 4
R4-R3, 4
R4-R5, 4
R4-R6, 2
R5-R1, 3
R5-R4, 5
R5-R6, 3
R6-R4, 1
R6-R5, 2

SPF Calculation Example

- Loop-free, lowest-cost path to every node

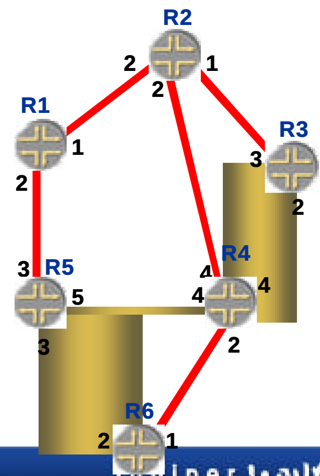

Path

R2-R2, 0
R2-R3, 1
R2-R1, 2
R2-R4, 2
R2-R5, 4
R2-R6, 4

Problems with Link State

- Information flooding load
  - Solution: sequence numbers and aging
  - Solution: areas
- Stale LS database entries
  - Solution: periodic database refresh
- .5(n2-n) adjacencies
  on multi-access networks
  - Solution: designated routers
- Memory and CPU overload
  - Solution: areas

Juniper Your Net

www.juniper.net 14

Link State Algorithms

- IS-IS
- OSPF

Open Shortest Path First (OSPF) Routing Protocol: Background

- Runs directly over IP (Protocol 89)
- Each router maintains an identical database (within areas)
- Each router constructs a tree of shortest paths by running SPF algorithm on the database
- Tree provides route to each known destination

OSPF Protocol Overview

- OSPF is a link-state IGP that routes packets within a single AS
- OSPF floods LSAs to distribute link-state information
- Each router uses these LSAs to create a complete routing table for the network
- OSPF uses the SPF algorithm to calculate the best route
- OSPF is defined in
  - RFC 2328, *OSPF Version 2*
  - RFC 1587, *The OSPF NSSA Option*

OSPF Protocol Overview

- OSPF routes IP packets based solely on the destination IP address in the IP header
- OSPF detects topological changes quickly and calculates new loop-free routes with a minimum of routing overhead traffic
- JUNOS software supports OSPF Version 2, including virtual links, stub and NSSA areas, and authentication

OSPF Terminology

- Areas
  - Single AS can be divided into smaller groups called areas
- ABRs
  - Routers that belong to more than one area are called area border routers
- Backbone area
  - Backbone area (0.0.0.0) distributes routing information between areas
- ASBRs
  - Routers that exchange routing information with routers in other ASs are called AS border routers

OSPF Terminology

- Stub areas
  - Areas through or into which AS external advertisements are not flooded
- Not-So-Stubby areas
  - Not-so-stubby areas allow external routes to be flooded within the area
- Transit areas
  - Pass traffic from one adjacent area to the backbone
  - Traffic does not originate in, nor is it destined for, the transit area

Juniper Your Net

www.juniper.net 20

Media Support

- OSPF
  - Broadcast (LANs)
  - Point-to-point
  - Point-to-multipoint
  - NBMA

Juniper your Net

OSPF Topology

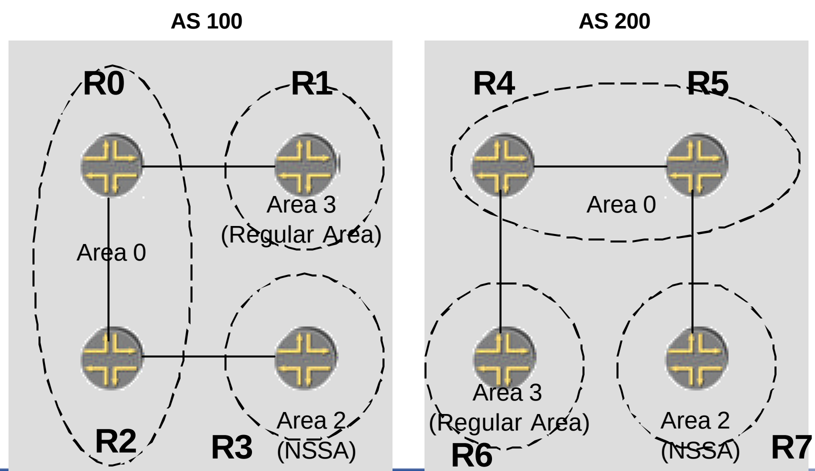

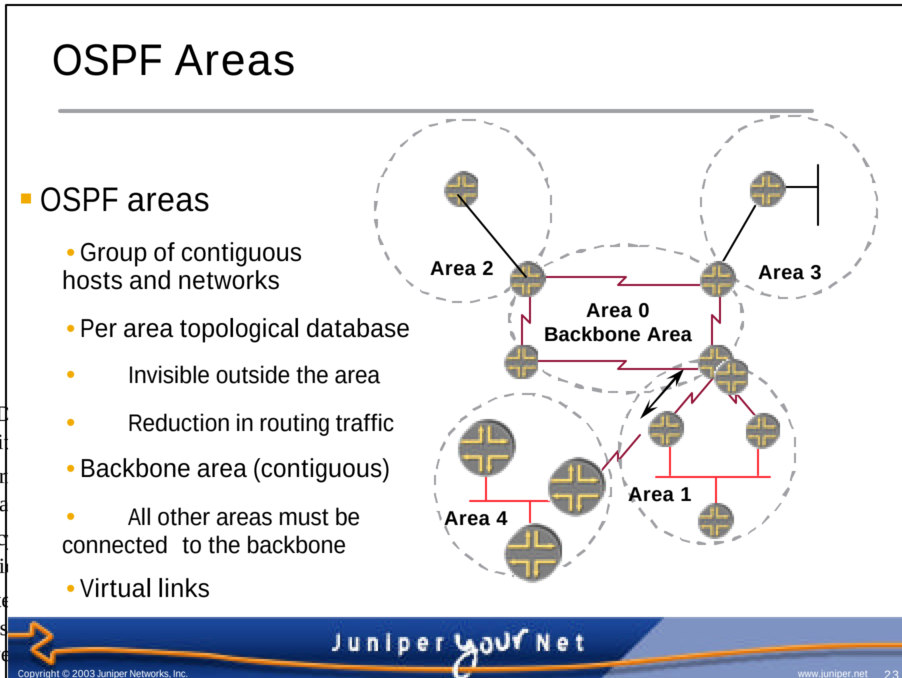

Router ID ... ies the router wi... .

Neighbor ... bors are dynamica...

Adjacenc... informati...

Link state ... r's interfaces ... ted link state adve...

Hello protocol: The part of the OSPF protocol used to establish and maintain neighboring relationships. On multi-access networks the Hello Protocol can also dynamically discover neighboring routers.

Designated router: Each multi-access network that has at least two attached routers has a Designated Router. The Designated Router generates a link state advertisement for the multi-access network and has other special responsibilities in the running of the protocol. The Designated Router is elected by the Hello Protocol.

Router and Area IDs: OSPF

- Router ID and area ID specified separately
- Each is 32-bit number
- AID associated with interface
- RID
  - Explicitly specified RID
  - Loopback address
  - Highest interface IP address

Classification of Routers

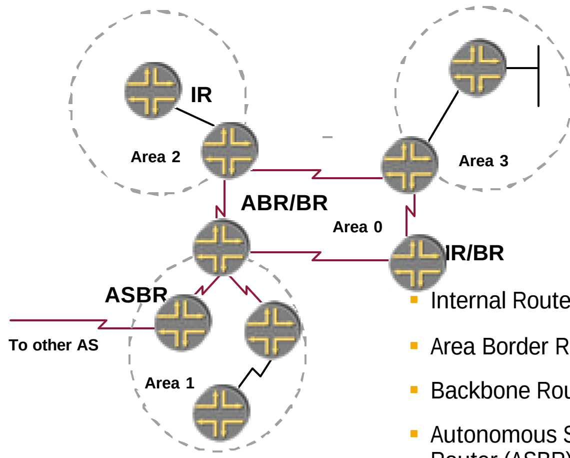

- Internal Router (IR)
- Area Border Router (ABR)
- Backbone Router (BR)
- Autonomous System Border Router (ASBR)

OSPF Route Types

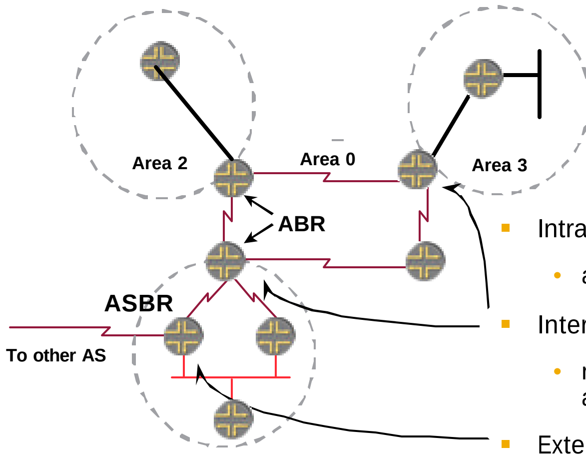

- Intra-area Route
  - all routes inside an area
- Inter-area Route
  - routes advertised from one area to another by an Area Border Router
- External Route
  - routes imported into OSPF from other protocol or static routes

OSPF Packet Types

- OSPF systems send five types of packets
  - Hello packets
  - Database description packets
  - Link-state request packets
  - Link-state update packets
  - Link-state acknowledgment packets
- Link-state request, link-state update, and link-state acknowledgment packets reliably flood link-state advertisement packets
- There are several types

Juniper Your Net

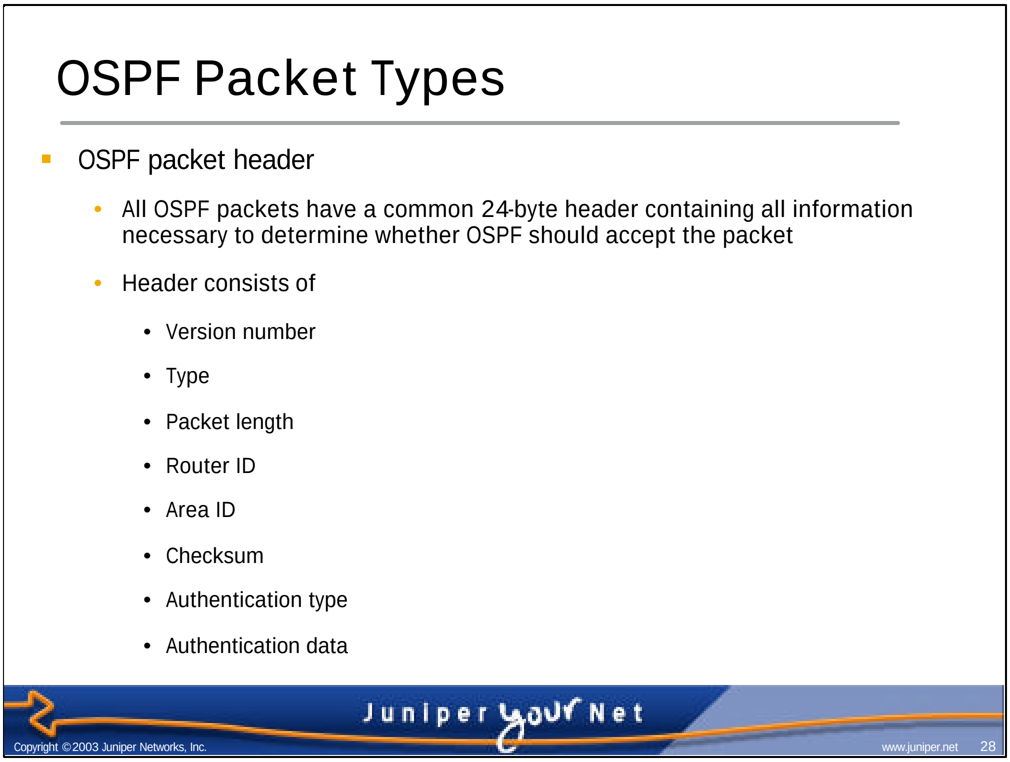

The OSPF packet header contains:

- Version number—The current OSPF version number is 2.
- Type—Type of OSPF packet.
- Packet length—Length of the packet, in bytes, including the header.
- Router ID—IP address of the router from which the packet originated.
- Area ID—Identifier of the area in which the packet is traveling. Each OSPF packet is associated with a single area. Packets traveling over a virtual link are labeled with the backbone area ID, 0.0.0.0. You configure the area ID with the area statements.
- Checksum—Fletcher checksum.
- Authentication type—Authentication scheme to use for the packet. You configure the authentication type with the `authentication-type` statement.
- Authentication data—The authentication information itself.

OSPF Packet Types

- Hello packets
  - Routers periodically send multicast Hello packets out all OSPF interfaces, including virtual links, to establish and maintain neighbor relationships
  - Hello packets consist of the OSPF header plus
    - Network mask
    - Hello interval
    - Options
    - Router priority
    - Router dead interval
    - Designated router
    - Backup designated router
    - Neighbor

The OSPF hello packet contains:

- Network mask—Network mask associated with the interface.
- Hello interval—How often the router sends hello packets. All routers on a shared network must use the same hello interval. You configure this interval with the hello-interval statement.
- Options—Optional capabilities of the router's OSPF implementation, such as stub routing or IP type of service based routing.
- Router priority—The router priority used to determine the designated router. You can configure this value with the priority statement.
- Router dead interval—How long the router waits without receiving any OSPF packets from a router before declaring that router to be nonoperational. All routers on a shared network must use the same router dead interval. You can configure this value with the dead-interval statement.
- Designated router—IP address of the designated router.
- Backup designated router—IP address of the backup designated router.
- Neighbor—IP addresses of the routers from which valid hello packets have been received within the time specified by the router dead interval.

The Hello Protocol

- Responsible for establishing and maintaining neighbor relationships
- Elects designated router on multi-access networks

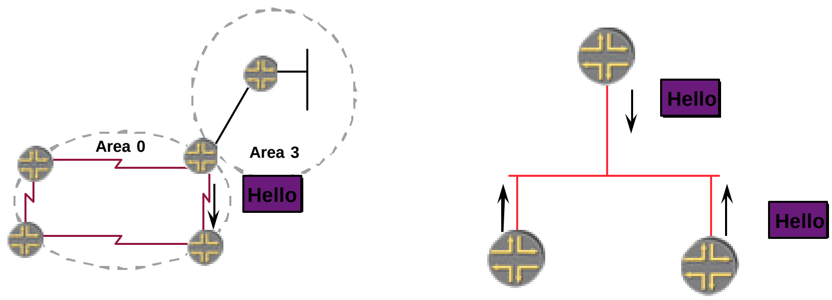

Designated Routers: OSPF

- Highest priority becomes DR
  - 0-255, default 128
  - Highest router ID tie-breaker
- Backup designated router
  - Speeds recovery from failed DR
- DR cannot be pre-empted
  - The DR is usually the first active router
- Adjacencies formed only with DR and BDR

Juniper Your Net

Neighbor Discovery and Maintenance: OSPF

- Hello packets
  - Establish two-way communication
  - Advertise optional capabilities
  - DR/BDR election/discovery
  - Serve as keepalives
  - 10s default hello interval, dead interval 4X
- Most hello fields must match for adjacency
  - Area ID, authentication, network mask, hello interval, router dead interval, options
  - Changing values causes adjacency disruption

Juniper your Net

OSPF Packet Types

- Database description packets
  - Exchanged during initialization
  - Describe contents of topological database
  - Consist of
    - OSPF header
    - Sequence number
    - LSA header

Juniper Your Net

OSPF Packet Types

- Link-state request packets
  - Sent when router detects database is stale
  - Request precise version of database
  - Consist of
    - OSPF header
    - Database information

OSPF Packet Types

- Link-state update packets
  - Carry one or more link-state advertisements
  - Multicast on physical networks that support multicast or broadcast mode
  - Acknowledged and, if retransmission is necessary, the retransmitted advertisements are sent unicast
  - Consist of
    - OSPF header
    - Number of advertisements
    - Link-state advertisements

Juniper Your Net

OSPF Packet Types

- Link-state acknowledgment packets
  - Sent in response to link-state update packets
  - Acknowledge the update packets have been received successfully
  - Can include responses to multiple update packets
  - Consist of
    - OSPF header
    - Link-state advertisement header

Juniper Your Net

OSPF Packet Types

- Link-state advertisements packet types
  - Router link advertisements
  - Network link advertisements
  - Network summary link advertisements
  - ASBR summary link advertisements
  - AS external link advertisements
  - NSSA external link advertisements
- Each LSA type describes a portion of the OSPF routing domain

OSPF LSAs

- Multiple LSA types

Type	LSA
1	Router LSA
2	Network LSA
3	Network Summary LSA
4	ASBR Summary LSA
5	AS External LSA
6	Group Membership LSA
7	Not-so-stubby Area LSA
8	External Attributes LSA
9-11	Opaque LSAs

Database Synchronization: OSPF

- Database synchronization driven by state machine
- Master/slave election
- Database synchronization
  - Database description packets
  - Link-state request packets
  - Link-state update packets
  - Link-state acknowledgement packets

Juniper Your Net

Database Refresh: OSPF

- LSA refresh every 30 minutes
- MaxAge = 1 hour
- Up-counting timer
- Design flaw: cannot change MaxAge

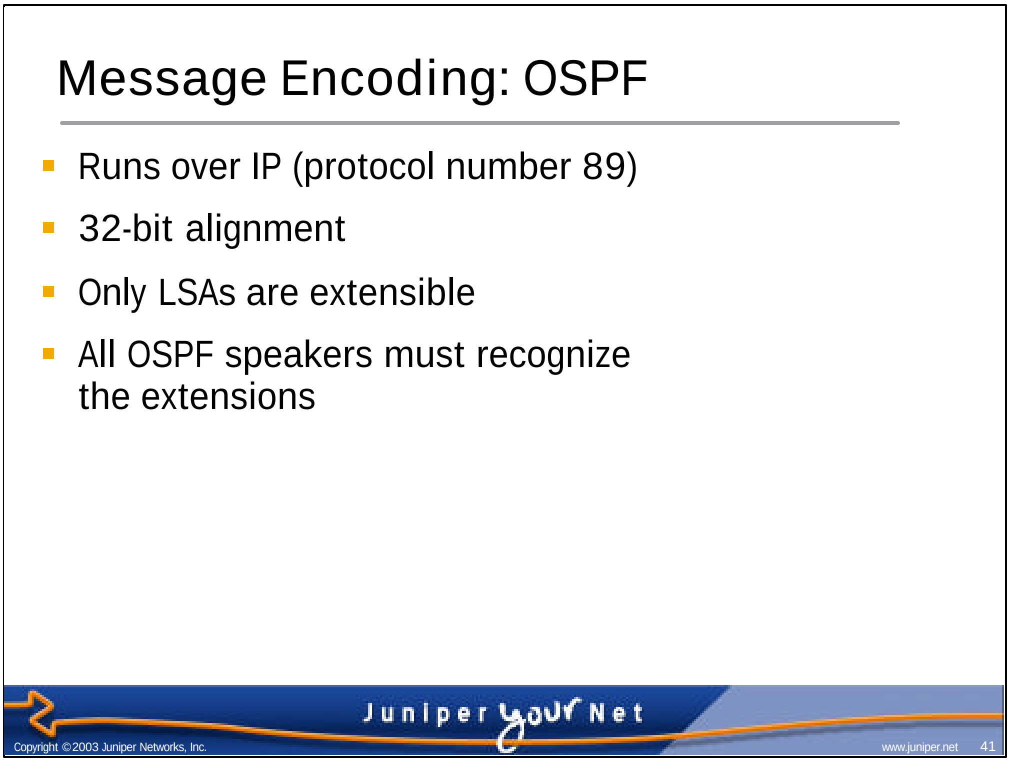
Message Encoding: OSPF
Runs over IP (protocol number 89)
32-bit alignment
Only LSAs are extensible
All OSPF speakers must recognize the extensions
Juniper your Net
Copyright ©2003 Juniper Networks, Inc.
www.juniper.net
41

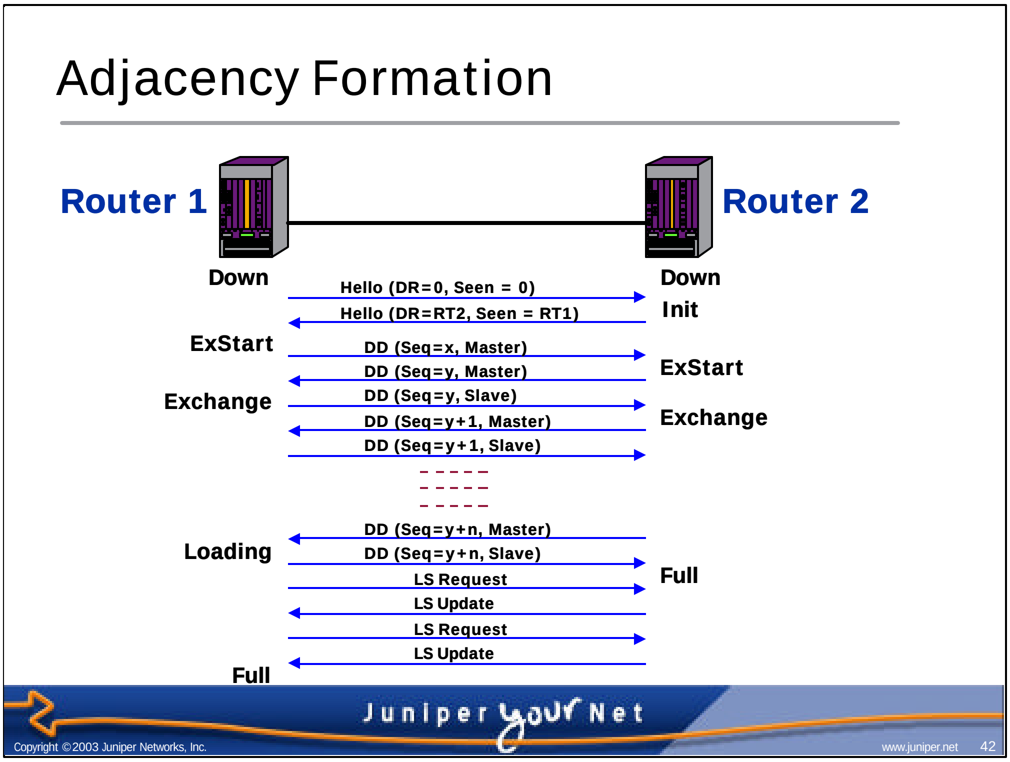

OSPF neighbors attempt to establish an *adjacency*, or a logical connection over which the neighbors agree to exchange information in their link state databases. The actual exchange process is called *database synchronization*. OSPF routers track the state of each neighbor using the following states:

- Down—This is the initial state, indicating that no recent information has been received from the neighbor.
- Init—A hello packet has been received from a neighbor, but bidirectional communication has not yet been established.
- ExStart—The router sees its own router ID in the neighbor's hello packet, indicating that bidirectional communication is established. The router sends a database description packet containing its idea of the first sequence number in the synchronization process, and claims to be the master router. The neighbor with the higher router ID becomes the master, and the other neighbor becomes the slave. The master's sequence numbers are used by both neighbors. Neighbors are considered adjacent when in this state or any of the states below.
- Exchange—The router sends database description packets to its neighbor describing its entire link-state database. If the router describes an LSA that the neighbor does not have, or that is more recent than the neighbor's copy, the neighbor can send a link-state request packet, asking for a complete copy of the LSA. Only one database description packet can be outstanding at a time, and each packet is acknowledged explicitly.
- Loading—A neighbor might enter this state. If a router has described all of its link state database, but has not yet received all the required LSAs from the neighbor, it transitions to this state.
- Full—The router's database is fully synchronized with its neighbor's database. The adjacency can now be advertised in router and network LSAs.

Not Summarized: Specific Links

- Specific link LSA advertised out
- Link state changes propagate out

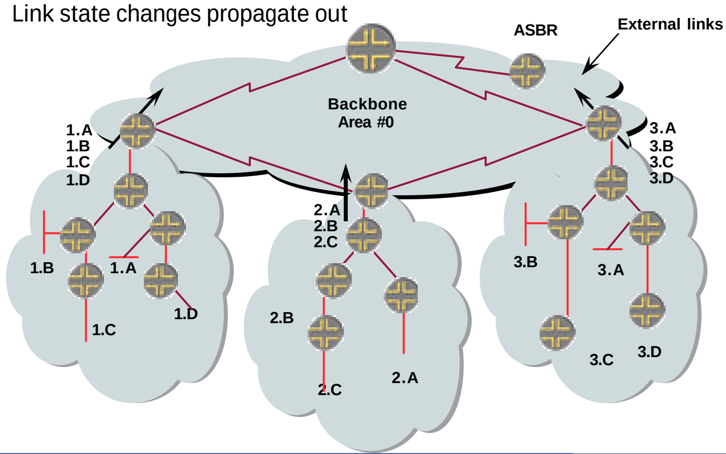

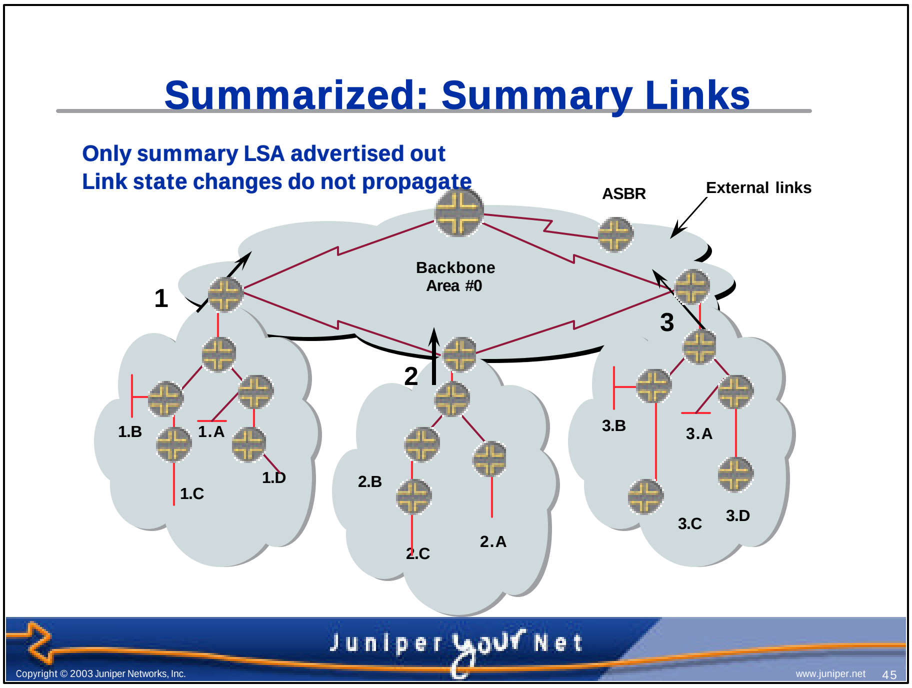
Summarized: Summary Links
Only summary LSA advertised out
Link state changes do not propagate
ASBR
External links
Backbone
Area #0
1
2
3
1.A
1.B
1.C
1.D
2.A
2.B
2.C
3.A
3.B
3.C
3.D
Juniper your Net
Copyright © 2003 Juniper Networks, Inc.
www.juniper.net
45

Not Summarized: Specific Links

- Specific link LSA advertised in
- Link state changes propagate in

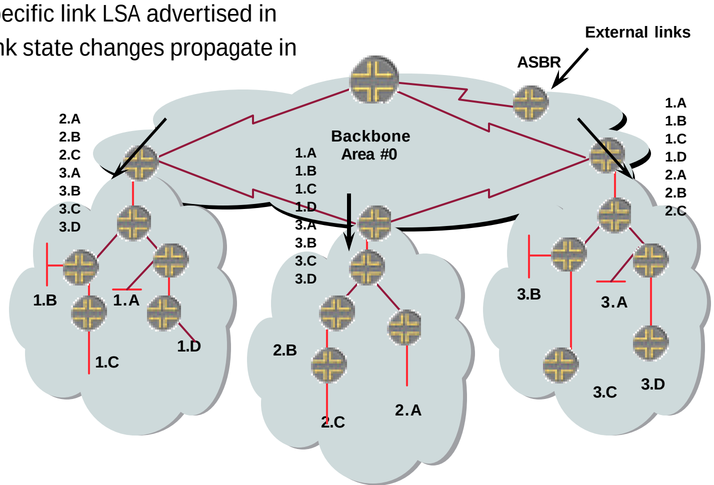

Summarized: Summary Links

- Specific link LSA advertised in
- Link state changes propagate in

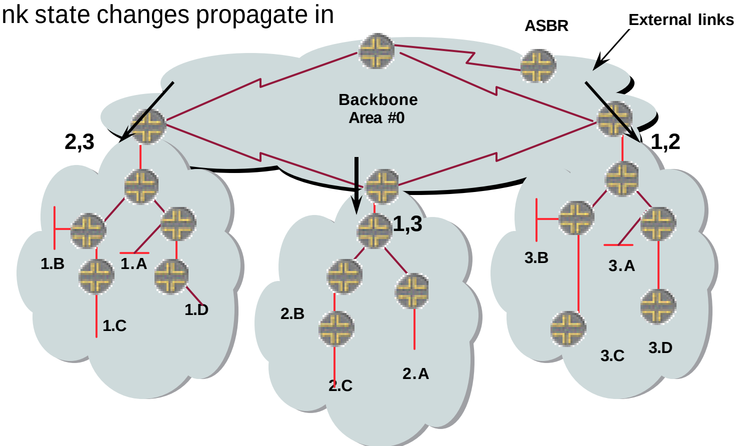

 47

Regular Area (Not a Stub)

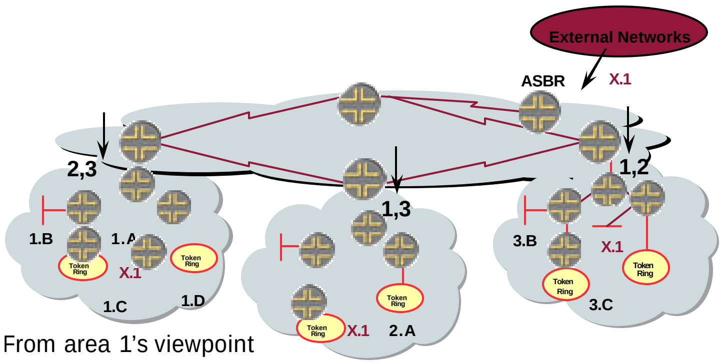

- From area 1's viewpoint
- Summary networks from other areas injected
- External networks injected, for example network X.1

Stub Areas

- Trade routing precision
  for improved scalability
- OSPF
  - Stub areas eliminate type 5 LSA load
  - Totally stubby areas extend the concept
  - All area routers must understand stubbiness

Juniper Your Net

Normal Stub Area

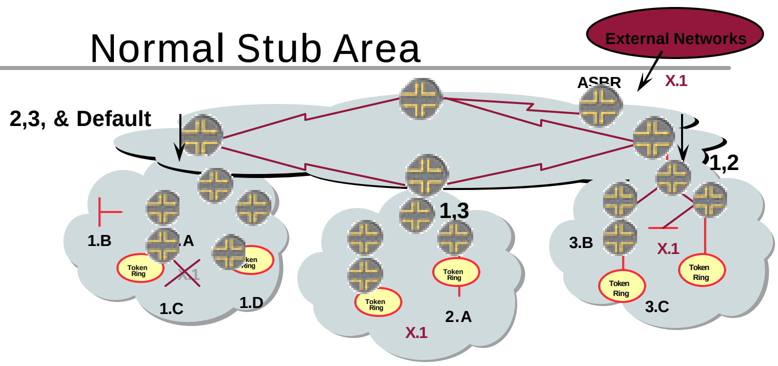

- From area 1's viewpoint
- Summary networks from other areas injected
- Default network injected into the area
  - Represents external links
- Default path to closest area border router
- Define all routers in the area as stub
  - area x stub command

Not-So-Stubby Areas

- OSPF feature
  - "Trick" to allow advertisement of external routes through stub areas (type 5 LSAs illegal)
  - All routers in area must understand type 7 LSAs

Juniper your Net

www.juniper.net 51

Not-So-Stubby Area

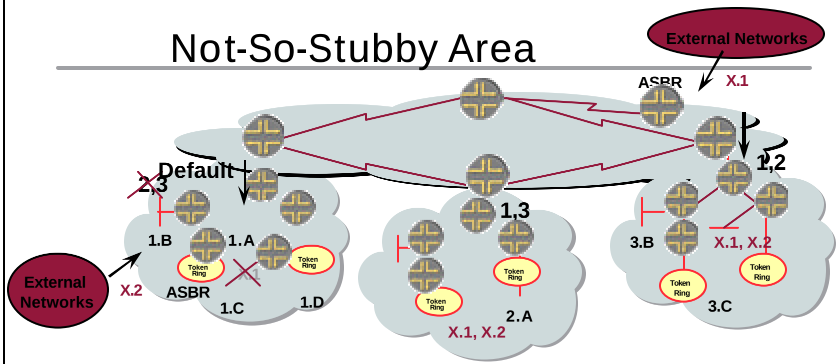

- New optional type of OSPF area
  - Supported as of 11.2 software version
- Capable of importing external routes in a limited fashion
- Type-7 LSA's carry external information within an NSSA
- NSSA Border routers translate selected type-7 LSAs into type-5 external network LSAs

Juniper your Net

www.juniper.net 52

Virtual Links

- Useful for
  - Patching partitioned backbone areas
  - Area migrations
- Should be a temporary solution!
- Full OSPF support

Major OSPF LSAs

- Router LSA
  - LSA type 1
  - Sent by all routers to describe state and cost of the router's links to the areas
  - Flooded throughout a single area only
- Network LSA
  - LSA Type 2
  - Sent by DRs to describe the network segment and all the routers attached to the segment
  - Flooded throughout a single area only

Juniper Your Net

Major OSPF LSAs

- Network summary LSA
  - LSA Type 3
  - Sent by ABRs to describe the routes that they know about in other areas
  - Flooded throughout the ABR's associated areas
- ASBR summary LSA
  - LSA Type 4
  - Sent by ABRs to describe the location of ASBRs, which are originating AS external LSAs
  - Flooded throughout the ABR's associated areas

Major OSPF LSAs

- AS external LSA
  - LSA Type 5
  - Sent by ASBRs to describe external routes they know about
  - Flooded throughout the AS (except into stub areas)
- NSSA external LSA
  - LSA Type 7
  - Used by not-so-stubby routers to import a selected number of external routes into the stub area
  - Flooded throughout the stub area only
  - Translated into AS external LSAs by an ABR

Juniper Your Net

Security

- OSPF support authentication
  - Plain-text passwords (sniffable!)
  - MD5 cryptographic hash
- Authentication especially important with OSPF
  - Runs over IP, so subject to spoofing
    and other attacks

Juniper Your Net

Checkpoint

- Can you now describe
  - OSPF protocol basics?
  - Each LSA type?

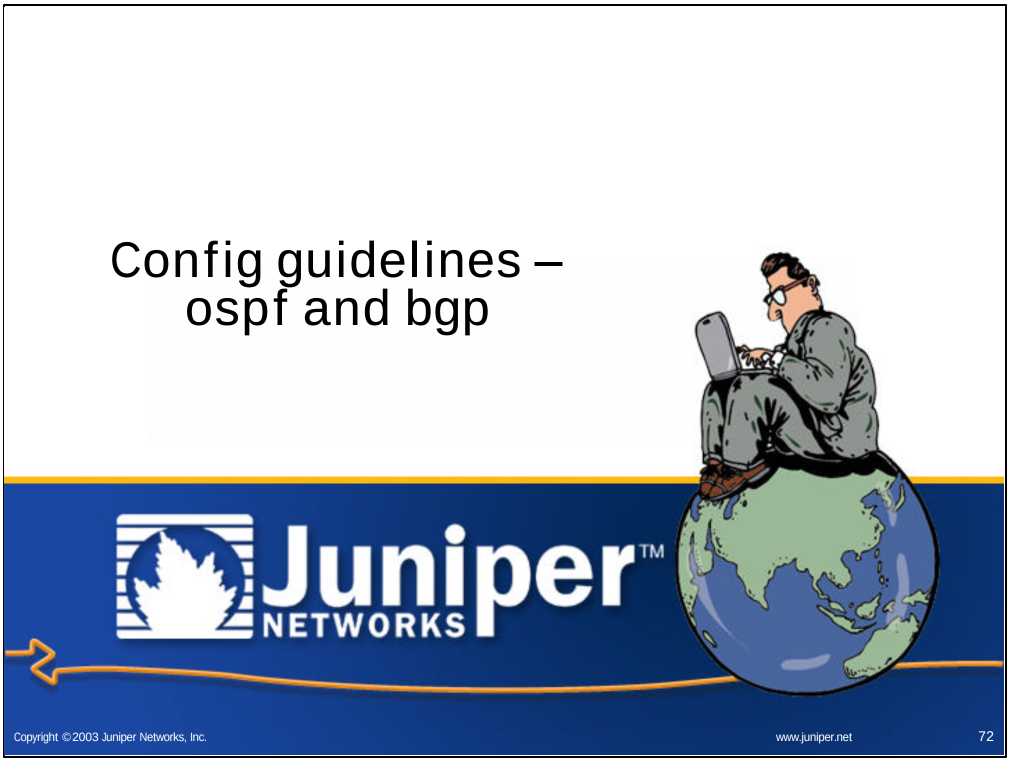
Config guidelines –
ospf and bgp
Juniper™
NETWORKS
Copyright © 2003 Juniper Networks, Inc.
www.juniper.net
72

Basic OSPF Configuration

Basic OSPF configuration, 1

Enable IP on All Router Interfaces

[edit] set interface *interface-name* unit 0 family inet

[edit]

Enable all interfaces on area 0.0.0.0

[edit]	edit protocols OSPF
[edit protocols ospf]	set area 0 interface all
[edit protocols ospf]	set area 0 interface fxp0 disable
[edit protocols ospf]	set area 0 interface *interface-name* metric *metric-value*
[edit protocols ospf]	top
[edit]	commit

Juniper your Net

www.juniper.net 74

To check the network

To see the OSPF interfaces

[edit] exit

user@host> show ospf interface brief / extensive

To see OSPF neighborss :

user@host> show ospf neighbor brief / extensive

To see OSPF routes :

user@host> show ospf route detail

To see the database

user@host> show OSPF database brief

Show OSPF Interfaces

- `show ospf interface ?`
  - View status of an interface

```
user@host> show ospf interface ?
Possible completions:
 <[Enter]>            Execute this command
 brief                Show brief status
 detail               Show detailed status
 extensive            Show extensive status
user@host> show ospf interface brief
Intf          State     Area      DR ID          BDR ID           Nbrs
ge-1/2/3.0    DRother   0.0.0.0   10.250.240.8   10.250.240.35       3
ge-2/0/0.0    DR        1.0.0.0   10.250.240.17  10.250.240.11       2
ge-2/1/9.0    DR        1.0.0.0   10.250.240.17  10.250.240.9        1
ge-4/1/3.0    DR        1.0.0.0   10.250.240.17  10.250.240.10       1
```

Juniper Your Net

The normal interface states your are likely to see are:

Backup—The interface is connected to a multiaccess network and is the backup designated router.

DR—The interface is connected to a multiaccess network and is the designated router for that network.

DRother—The interface is connected to a multiaccess network and is not the designated or the backup designated router on that network.

PtToPt—The interface is connected to a point-to-point network or virtual link.

Waiting —The router is trying to determine the backup designated router on the specified interface.

Show OSPF Interfaces

- `show ospf interface extensive`
  - View more extensive OSPF interface information

```
user@host> show ospf interface extensive
Intf              State    Area       DR ID         BDR ID        Nbrs
ge-1/2/3.0        Drother  0.0.0.0    10.250.240.8  10.250.240.35     3
  Type LAN, address 192.168.254.227, mask 255.255.255.240, MTU 1500, cost 1
  DR addr 192.168.254.230, BDR addr 192.168.254.229, adj count 2

fxp2.0            DR       1.0.0.0    10.250.240.17 10.250.240.11     2
  Type LAN, address 10.1.1.130, mask 255.255.255.128, MTU 1500, cost 1
  DR addr 10.1.1.130, BDR addr 10.1.1.131, adj count 2

fxp1.0            DR       1.0.0.0    10.250.240.17 10.250.240.9      1
  Type LAN, address 10.1.2.2, mask 255.255.255.240, MTU 1500, cost 1
  DR addr 10.1.2.2, BDR addr 10.1.2.1, adj count 1
```

Juniper Your Net

View OSPF Statistics

- `show ospf statistics`
  - View basic OSPF protocol statistics

```
user@host> show ospf statistics
Packet type          Total              Last 5 seconds
              Sent       Received       Sent       Received
   Hello     505739      990495          4             5
     DbD         20          26          0             0
   LSReq          6           5          0             0
LSUpdate      27060       15319          0             0
   LSAck      10923       52470          0             0
LSAs retransmitted: 16, last 5 seconds: 0
Receive errors:
 862 no interface found
 115923 no virtual link found
```

Show OSPF Route Example

- **user@host>** `show ospf route detail`
- **Prefix Route/Path Type Metric Next hop i/f Next hop addr**
- **1.1.1.0/24 Ext2 Network 0 ge-0/0/0.0 10.10.0.16**
- **area 0.0.0.0, options 0x0, origin 1.1.1.1**
- **1.1.1.1/32 Intra AS BR 1 ge-0/0/0.0 10.10.0.16**
- **area 0.0.0.0, options 0x0, origin 1.1.1.1**
- **1.2.3.0/24 Ext2 Network 0 ge-0/0/0.0 10.10.0.16**
- **area 0.0.0.0, options 0x0, origin 1.1.1.1**

Juniper Your Net

View OSPF Database

- `show ospf database`
  - View the LSA database

```
user@host> show ospf database ?
Possible completions:
  <[Enter]>              Execute this command
  advertising-router     Router ID of advertising router
  area                   OSPF area ID
  asbrsummary            Show OSPF summary ASBR link-state database
  brief                  Show brief view
  detail                 Show detailed view
  extensive              Show extensive view
  extern                 Show OSPF external link-state database
  lsa-id                 LSA ID
  netsummary             Show OSPF summary network link-state database
  network                Show OSPF network link-state database
  nssa                   Show OSPF NSSA link-state database
  router                 Show OSPF router link-state database
```

Show OSPF Database Example

```
user@host> show ospf database brief
OSPF link state database, area 0.0.0.0
Type     ID              Adv Rtr          Seq         Age   Cksum  Len
Router   10.250.240.8    10.250.240.8     0x800001fc  2388  0x3684  36
Router   10.250.240.17   10.250.240.17    0x80000217  1835  0x444c  36
Router   10.250.240.32   10.250.240.32    0x80000232  1876  0x0158  36
Router   10.250.240.35   10.250.240.35    0x80000291  1100  0x4aa5  36
Network  192.168.254.230 10.250.240.8     0x800001cc   117  0xab67  40
Summary  10.1.2.0        10.250.240.17    0x80000216  1535  0x1729  28
Summary  10.1.3.34       10.250.240.8     0x8000013a  2217  0x842f  28
```


Basic BGP
Configuration
Juniper™
NETWORKS
Copyright © 2003 Juniper Networks, Inc.
www.juniper.net
82

IBGP Configuration Overview

```
bgp {
    traceoptions {
        file bgp_log size 5m files 10;
        flag open;
        flag state;
        flag normal;
    }
    group core {
        type internal;
        local-address 192.168.255.7;
        peer-as 65000;
        authentication-key "$9$eDbKX-Y2aUi.oJfz3npuBIE";
        neighbor 192.168.255.1;
        neighbor 192.168.255.2;
        neighbor 192.168.255.3;
    }
}
```


This example shows an empty policy, with each optional part present. The to and from keywords are optional but it is likely you will use at least one. The then keyword is also optional, and if you do not specify it the system continues processing the next term in the policy. If there is no additional term, the final action is evaluated. If there is no final action, the default policy action is taken.

EBGP Configuration Overview

```
bgp {
group core {
        type internal;
         local-address 192.168.255.7;
        peer-as 65000;
        authentication-key "$9$eDbKX-Y2aUi.oJfz3npuBIE";
        neighbor 192.168.255.1;
        neighbor 192.168.255.2;
}
group AS_692 {
        type external;
        peer-as 692;
        import next-hop-self;
        authentication-key"$9$c6ZylMdb2JUHM8ZjkP3n/Ct";
        neighbor 172.16.5.1;
    }
}
policy-options {
    policy-statement next-hop-self {
        then {
            nexthop self;
        }
    }
}
}
```

Juniper Your Net

www.juniper.net 84

This example shows an empty policy, with each optional part present. The to and from keywords are optional but it is likely you will use at least one. The then keyword is also optional, and if you do not specify it the syst em continues processing the next term in the policy. If there is no additional term, the final action is evaluated. If there is no final action, the default policy action is taken.

BGP Information

- Several commands display a wide variety of BGP information either from the protocol itself or from BGP routes

```
user@host> show bgp ?
 Possible completions:
   group               Show the BGP group database
   neighbor            Show the BGP neighbor database
   next-hop-database   Show the BGP next hop database
   summary             Show an overview of the BGP
    information
```

Show BGP Summary

- `show bgp summary`
  - View basic information about all BGP neighbors

```
user@host> show bgp summary
Groups: 12      Peers: 26      Unestablished peers: 2


Peer           AS    InPkt    OutPkt   OutQ   Flaps Last Up/Dn
   State|#Act/Recv/Da…
131.103.0.2    45    1225     55263    50511    0    18:22:14
   47769/50591/0
192.168.1.1    33    911      0        0        0    18:22:27    Active
192.168.1.97   23    10458    2201     41043    0    18:22:03    0/0/0
192.168.1.100  432   10458    163      17643    0    17:01:18    Active
```

Juniper Your Net

Output fields:

- Groups—Number of BGP groups.
- Peers—Number of BGP peers.
- Unestablished peers—Number of unestablished BGP peers.
- Peer—Address of each BGP peer. Each peer has one line of output.
- AS—Peer’s AS number.
- InPkt—Number of packets received from the peer.
- OutPkt—Number of packets sent to the peer.
- OutQ—Count of the number of BGP packets that are queued to be transmitted to a particular neighbor. It usually is 0 because the queue is emptied quickly.
- Last Up/Down—Last time since the neighbor transitioned to or form the established state.
- State/#Act/Recv/Damped—Displays either the BGP state or, if the neighbor is connected, the number of paths received from the neighbor, the number of these paths that have been accepted as active and are being used for forwarding, and the number of routes being damped.

Show BGP Neighbor

```
user@host> show bgp neighbor
Peer: 11.1.1.2+179 AS 29        Local: 11.1.1.1+1048 AS 29
  Type: Internal  State: Established    Flags: <>
  Last State: OpenConfirm       Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference HoldTime>
        Holdtime: 90    Preference: 170
  Number of flaps: 1
  Error: "Cease" Sent: 1 Recv: 0
  Peer ID: 11.1.1.2     Local ID: 0.0.0.0       Active Holdtime: 90
  NLRI advertised by peer: unicast
  NLRI for this session: unicast
  Group Bit: 0  Send state: in sync
  Table inet.0
    Active Prefixes: 0
```


Output fields:

- Peer—Address of each BGP peer. Each peer has one line of output.
- Type—Type of peer (internal or external).
- State—BGP state for this neighbor.
- Flags—Internal peer-specific flags for this neighbor.
- Last State—BGP state that this neighbor was in prior to the current state.
- Last Event—Last BGP state transition event.
- Last Error—Last notification sent to the neighbor.
- Options—Configuration options that are in effect for this neighbor.
- Holdtime—Configured hold time for this neighbor.
- Preference—Configure preference for routes learned from the neighbor.
- Peer ID—Neighbor's router ID.
- Local ID—Local system's router ID.
- Active Holdtime—Hold-time value that was negotiated during the BGP open.
- Group Bit—Internal bit being used for the peer group.
- Send state—Whether all peers in the group have received all their updates (in sync or out of sync).
- Active Prefixes—Number of prefixes accepted as active from this neighbor.
- Last traffic (seconds)—How recently a BGP message was sent or received between the local system and this neighbor.
- Output Queue—Number of BGP update messages that are pending for transmission to the neighbor.
- Deleted routes—Prefixes that are queued for withdrawal through pending update messages.
- Queued AS Path—An AS path that is queued for transmission in an update message.

Show BGP Next Hop

- `show bgp next-hop-database`

```
user@host> show bgp next-hop-database brief

1.0.0.0/8

Source: 10.168.1.222 Nexthop: 10.168.1.222

10.168.1.222/32 MED 20 Next hops 192.168.200.2
   192.168.200.102

2.0.0.0/8

Source: 10.168.1.222 Nexthop: 10.168.1.222

10.168.1.222/32 MED 20 Next hops 192.168.200.2
   192.168.200.102

3.0.0.0/8

Source: 10.168.1.222 Nexthop: 10.168.1.222

10.168.1.222/32 MED 20 Next hops 192.168.200.2
   192.168.200.102
```


This command displays brief information about the entries in the IBGP synchronization database.

Output fields:

ip-address—Prefix.

Source—Route from which prefix was learned.

Next hop—Type of next hop for the route that was used to resolve the BGP next hop.

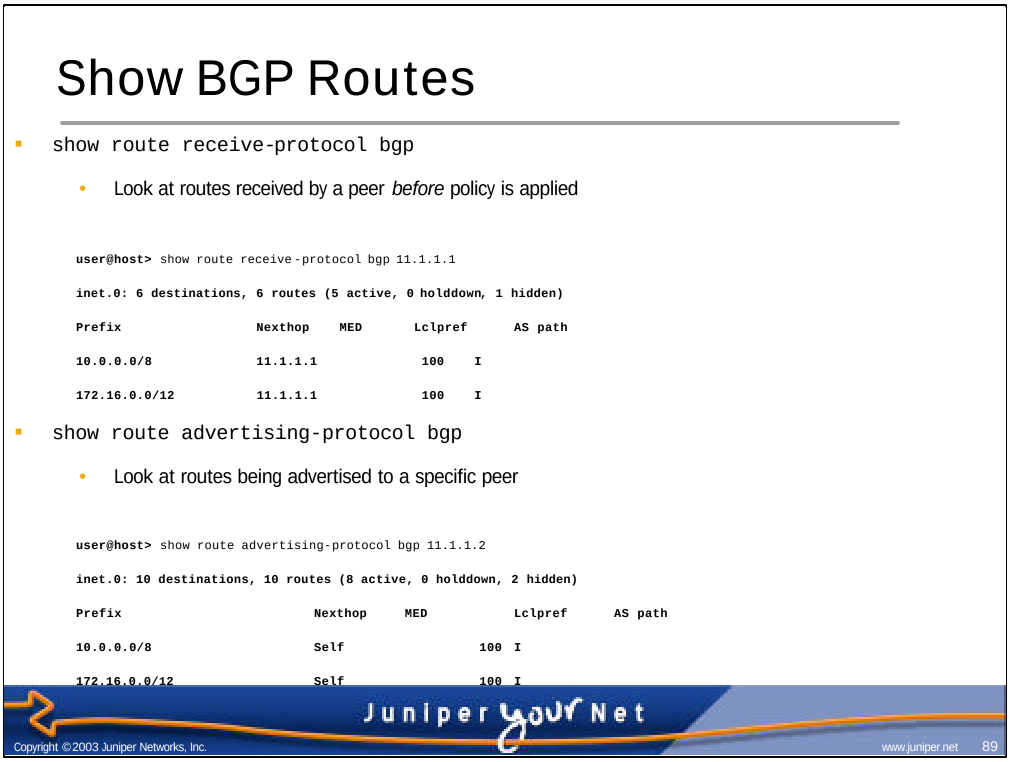


```
user@host> show route receive-protocol bgp 11.1.1.1
```

This command displays the routing information as it was received through a particular neighbor of a particular dynamic routing protocol. This information includes the routes that the local router advertised to the neighbor. The information reflects the routes before they are filtered by that protocol's import policy statements. This command works for routing protocols BGP, RIP, DVMRP, and PIM only.

```
user@host> show route advertising-protocol bgp 11.1.1.2
```

This command displays the routing information as it has been prepared for advertisement to a particular neighbor of a particular dynamic routing protocol. The information reflects the routes that the routing table exported into the routing protocol and that were filtered by that protocol's export routing policy statements. This command works for routing protocols BGP, RIP, DVMRP and PIM only.

Looking at Specific Routes

- `show route extensive`
  - Look at specific entries in the routing table

```
user@host> show route 172.16.0.0 extensive
inet.0: 6 destinations, 6 routes (5 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.0.0/12 (1 entry, 1 announced)
TSI:
BGP_Sync_Any dest 172.16.0.0/12 MED 0
        *BGP    Preference: 170/-101
                Nexthop: 11.1.1.1 via fxp0.0, selected
                State: <Active Int Ext>
                Local AS:    29 Peer AS:    29
                Age: 1d 9:46:54 Metric2: 0
                Task: BGP_29.11.1.1.1+1048
                Announcement bits (2): 0-KRT 2-BGP_Sync_Any
                AS path: I
                BGP next hop: 11.1.1.1
                Localpref: 100
                Router ID: 172.18.1.1
```


Juniper Your Net

Routing Policy Overview

www.juniper.net

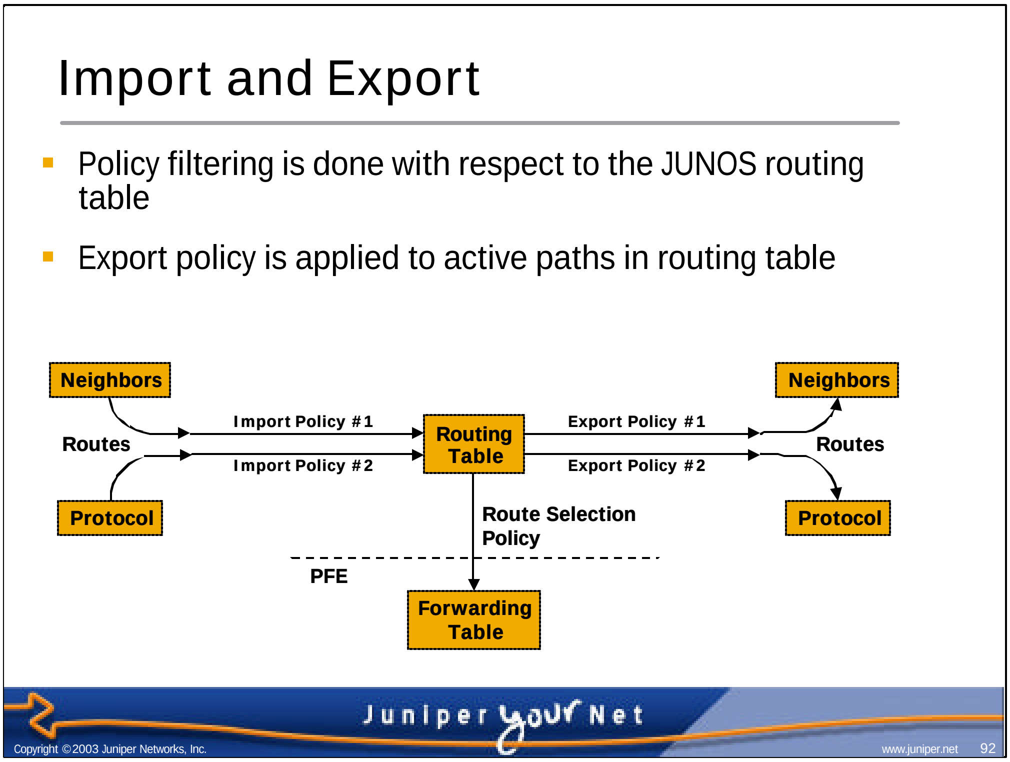

Before discussing the design of routing policy, it is necessary to understand two terms—*import* and *export*. JUNOS routing policy uses these terms to describe how routes move between the routing protocols and the routing tables:

- When a routing protocol places its routes into the routing table, this process is referred to as *importing* routes into the routing table.
- When a dynamic routing protocol uses the routes in the routing table to send protocol advertisements, the protocol takes the route from the routing table. This process is referred to as *exporting routes* from the routing table.
- The process of moving routes between a routing protocol and the routing table is always described *from the point of view of the routing table*. That is, routes are *imported into* a routing table from a routing protocol and they are *exported from* a routing table to a routing protocol. It is important to remember this distinction when working with routing policy.

When evaluating routes for export, the policy software uses only active routes from the routing table. For example, if you have multiple routes to the same destination and one route has a more attractive metric, only that route is evaluated by the policy processing software. Said another way, export policy does not evaluate all routes, but instead it evaluates only those routes a routing protocol could correctly advertise to a neighbor.

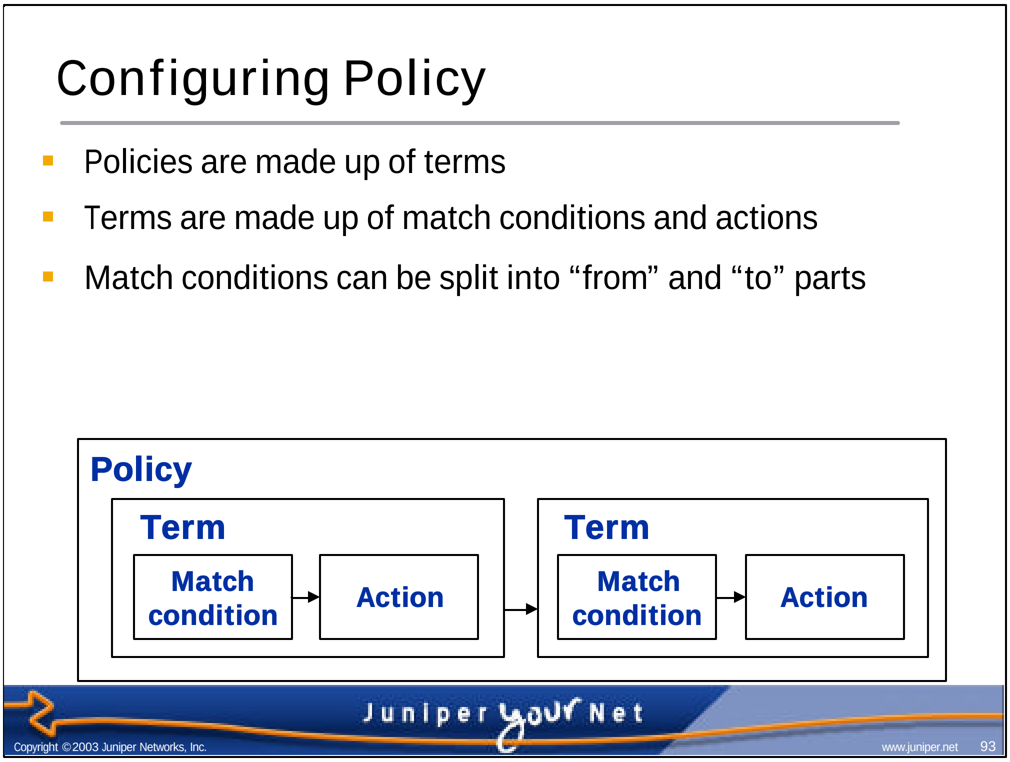

Each routing policy is introduced by the keyword policy-statement and is identified by a name, *policy-name*:

policy-statement *policy-name* { ... }

Each policy statement consists of one or more *terms*. Each term is introduced by the keyword term and identified by a name, *term-name*:

term *term-name* { ... }

As with all names in the JUNOS software, the policy and term names can contain letters, numbers, and hyphens (–), and can be up to 255 characters long. To include spaces in the name, enclose the entire name in quotation marks (double quotes).

Each policy term consists of statements that define match conditions and actions to take if the conditions are matched. There are three statements: a from statement and a to statement, which together specify the match conditions, and a then statement, which specifies the action to take if a match occurs. If you are not matching specific routes, the to and from statements are optional.

The from statement defines conditions used to match a route's source. You can specify one or more of the match conditions previously discussed. If you specify more than one, they all must match the route for a match to occur.

The from statement is optional. If you omit it, all routes are considered to match.

The to statement matches conditions that pertain to the route's destination or the protocol into which the route is being advertised. You can specify one or more of the match conditions previously discussed. If you specify more than one, they all must match the route for a match to occur.

The to statement is optional. If you omit it, all routes are considered to match.

Configuring Policy

- Basic policy syntax

```
policy-options {
    policy-statement policy-name {
        term term-name {
            from {
                match-conditions;
            }
            to {
                match-conditions;
            }
            then {
                action;
            }
        }
        final-action;
    }
}
```

A policy can have multiple terms

 94

This example shows an empty policy, with each optional part present. The to and from keywords are optional but it is likely you will use at least one. The then keyword is also optional, and if you do not specify it the syst em continues processing the next term in the policy. If there is no additional term, the final action is evaluated. If there is no final action, the default policy action is taken.

Configuring Policy

- Example

```
policy-options {
    policy-statement advertise-ospf {
        term pick-ospf {
            from protocol ospf;
            then accept;
        }
    }
}

protocols bgp {
    export advertise-ospf;
}
```


Many real-world programming languages do not require braces around orphaned statements, and the JUNOS routing policy language is no exception.

The example above has the extra braces removed from the parts of the policy containing only single items.

How would you rewrite the following policy to remove the extra braces?

```
policy-statement set-preference {
    term term1 {
        from {
            nexthop [ 10.0.0.1 10.0.0.2 ];
        }
        to {
            area 51;
        }
        then {
            preference 10;
            accept;
        }
    }
    term term2 {
        from {
            nexthop 10.0.0.3;
        }
        to {
            area 51;
        }
        then {
            preference 15;
        }
    }
}
```

Configuring Policy

- The same sample, written another way

```
policy-options {
    policy-statement advertise-ospf {
        from protocol ospf;
        then accept;
    }
}

protocols bgp {
    export advertise-ospf;
}
```

 www.juniper.net 96

If a policy contains only one term, you can omit the `term` statement to make the policy easier to read.

If the policy contains only one match or action clause, you can omit the enclosing braces to make the policy easier to read. The extra braces were removed from the `from` and `then` clauses in the example above to improve clarity.

In addition, the final policy action (used when no terms accept or reject routes) does not need to be surrounded by a `term` clause, so the example below is valid.

Export routes to neighbors that come only from OSPF area 15:

```
policy-options {
    policy-statement only-ospf-area-15 {
        term accept-ospf-15 {
            from {
                protocol ospf;
                area 15;
            }
            then accept;
        }
        then reject;
    }
}
```

Test Your Knowledge (I)

- Answer

```
policy-options {
    policy-statement isis-level2 {
        from {
            protocol isis;
            level 2;
        }
        then accept;
    }
}

protocols bgp {
    export isis-level2;
}
```

Logical AND function