



Maintaining and Troubleshooting Routing Solutions



TSHOOT Module 7, 8

Chapter 5 Objectives

- Diagnose network layer connectivity problems using the IOS command line interface.
- Diagnose and resolve problems related to the exchange of routing information by the Enhanced Interior Gateway Routing Protocol (EIGRP).
- Diagnose and resolve problems related to the exchange of routing information by use of the Open Shortest Path First (OSPF) routing protocol.
- Diagnose and resolve problems when redistributing routes.
- Diagnose and resolve problems related to the exchange of routing information by use of the Border Gateway Protocol (BGP).

Troubleshooting OSPF



OSPF Review: Process and Operation ①

- **Reception of routing information from neighbors**
 - Routing information is exchanged in the form of link-state advertisements (LSAs)
 - LSAs contain information about elements of the network topology (routers, neighbor relationships, connected subnets, areas and redistribution)
- **Routing protocol data structures**
 - OSPF stores the LSAs that it receives in a link-state database
 - The SPF algorithm computes the shortest path to each network in terms of cost, (the OSPF metric), based on the information in the link-state database
 - Several other data structures, such as an interface table, a neighbor table, and a routing information base (RIB) are maintained
- **Route injection or redistribution**
 - Directly connected networks that are enabled for OSPF are advertised in the router's LSA
 - Routes from other sources, such as other routing protocols or static routes can also be imported into the link-state database and advertised by use of special LSAs

OSPF Review: Process and Operation ②

■ Route selection and installation

- OSPF will attempt to install the best routes, computed using the SPF algorithm, in the routing table
- OSPF discerns three different types of routes: intra-area routes, inter-area routes, and external routes
- If two routes of different types for the same prefix are available for installation in the routing table, OSPF will prefer intra-area routes over inter-area routes and both these types will be preferred over external routes, regardless of the cost of the paths
- If two equal cost routes of the same type are available, they will both be selected for installation in the routing table

■ Transmission of routing information to neighbors

- Routing information is flooded to all routers in an area by passing LSAs from neighbor to neighbor using a reliable transport mechanism
- Area Border Routers (ABRs) inject routing information from an area into the backbone area or, reversely, from the backbone area into the other areas that it is connected to.

OSPF Review: Data Structures

■ Interface table

- Lists all interfaces that have been enabled for OSPF.
- The directly connected subnets, that are associated with these interfaces, are included in the type 1 router LSA that the router injects into the OSPF link-state database for its area.
- When an interface is configured as a passive interface, it is still listed in the OSPF interface table, but no neighbor relationships are established on this interface.

■ Neighbor table

- Keeps track of all active OSPF neighbors.
- Neighbors are added to this table based on the reception of Hello packets
- Neighbors are removed when the OSPF dead time for a neighbor expires or when the associated interface goes down.
- OSPF goes through a number of states while establishing a neighbor relationship (also known as adjacency).
- The neighbor table lists the current state for each individual neighbor.

OSPF Review: Data Structures – Cont.

■ Link-state database


- This is the main data structure in which OSPF stores network topology information.
- This database contains full topology information for the areas that a router connects to, and information about the paths that are available to reach networks and subnets in other areas or other autonomous systems.
- This database is one of the most important data structures from which to gather information when troubleshooting OSPF problems.

■ Routing information base


- After executing the SPF algorithm, the results of this calculation are stored in the RIB or routing table.
- This information includes the best routes to each individual prefix in the OSPF network with their associated path costs.
- When the information in the link-state database changes, only a partial recalculation might be necessary (depending on the nature of the change).
- Routes might be added to or deleted from the RIB without the need for a full SPF recalculation.

OSPF Review: Network Areas and LSAs

A multi-area OSPF network with five routers performing no route redistribution




OSPF Review: Network Areas and LSAs – Cont.



Area DB	LSA Type-1	LSA Type-2	LSA Type-3
Area 1	2	1	5
Area 0	3	1	4
Area 2	2	0	5

OSPF Review: Network Areas and LSAs – Cont.



Router	LSA Type-1	LSA Type-2	LSA Type-3
Router A	2	1	5
Router B	5	2	9
Router C	3	1	4
Router D	5	1	9
Router E	2	0	5


OSPF Review: Information Flow Within an Area

- Two routers will become OSPF neighbors if the following parameters match in the Hello packets:
 - 1) **Hello and dead timers**
 - Neighbors must use the same Hello and dead time.
 - Broadcast and point-to-point type networks default to 10-second Hello and 40-second dead time.
 - If timers are changed on an interface, change timers for neighboring routers on that interface.
 - 2) **OSPF area number**
 - Routers will become neighbors on a link only if they both consider that link to be in the same area.
 - 3) **OSPF area type**
 - Routers will become neighbors only if they both consider the area to be the same type of area (normal, stub, or not-so-stubby area [NSSA]).
 - 4) **IP subnet and subnet mask**
 - Two routers will not become neighbors if they are not on the same subnet.
 - The exception to this rule is on a point-to-point link, where the subnet mask is not verified.
 - 5) **Authentication type and authentication data**
 - Routers will become neighbors only if they both use the same authentication type (null, clear text, or MD5).
 - If they use authentication, the authentication data (password or hash value) must match.

OSPF Review: Information Flow Within an Area

- Building a neighbor relationship or adjacency with the neighboring router consists of several states:
- **Init**
 - State of the neighbor when a Hello has been received from this neighbor, but the neighbor is not listing this router in its neighbor list yet.
- **2-Way**
 - State when the router sees its own router ID listed in the active neighbor list in the Hello packets received from that neighbor.
- **Exstart**
 - Indicates that the routers are starting the database exchange state by establishing a master and slave relationship.
- **Exchange**
 - Neighboring routers exchange link-state database description packets to determine which entries each neighbor is missing.
- **Loading**
 - Each of the two routers can request missing LSAs from the other router.
- **Full**
 - The final stage where neighbors have successfully synchronized their link-state databases.

OSPF Review: Information Flow Between Areas



- ABRs (routers B and D) play a key role in exchanging routing information between OSPF areas.
- When two neighbors in the same area exchange databases type-1 and type-2 LSAs that belong to different areas are not exchanged.
- To distribute information about subnets that are available in a particular area to other areas, the ABR generates type-3 LSAs.
- The type-3 LSAs are injected by the ABR into the area 0 database.
- Other ABRs use these type-3 LSAs to compute the best path to these subnets and then inject the information into their connected areas by use of new type-3 LSAs.

Monitoring OSPF with show Commands

- To gather information from the OSPF data structures use the following **show** commands

show ip ospf

- Displays general OSPF information, Router ID, areas, their types, SPF run count and other

show ip ospf interface [brief]

- Displays the interfaces that have been activated for OSPF.
- Listing contains interfaces that have an IP address covered by a network statement or configured using the interface configuration mode command: ip ospf process-number area area-number.

show ip ospf neighbor

- Lists all neighbors of this router on its active OSPF interfaces and shows their current state.

show ip ospf database

- Displays a summary of the OSPF link-state database content (LSA headers).
- Using additional command options, specific LSAs can be selected and the actual LSA content can be inspected.

show ip ospf statistics

- Displays how often and when the SPF algorithm was last executed.
- This command can be helpful when diagnosing routing instability.

Monitoring OSPF with debug Commands

- To observe the real-time OSPF information exchange use the following **debug** commands:

debug ip routing

- Command is not specific to the OSPF protocol.
- Displays any changes that are made to the routing table, such as installation or removal of routes.
- Can be useful in diagnosing routing protocol instabilities.

debug ip ospf packet

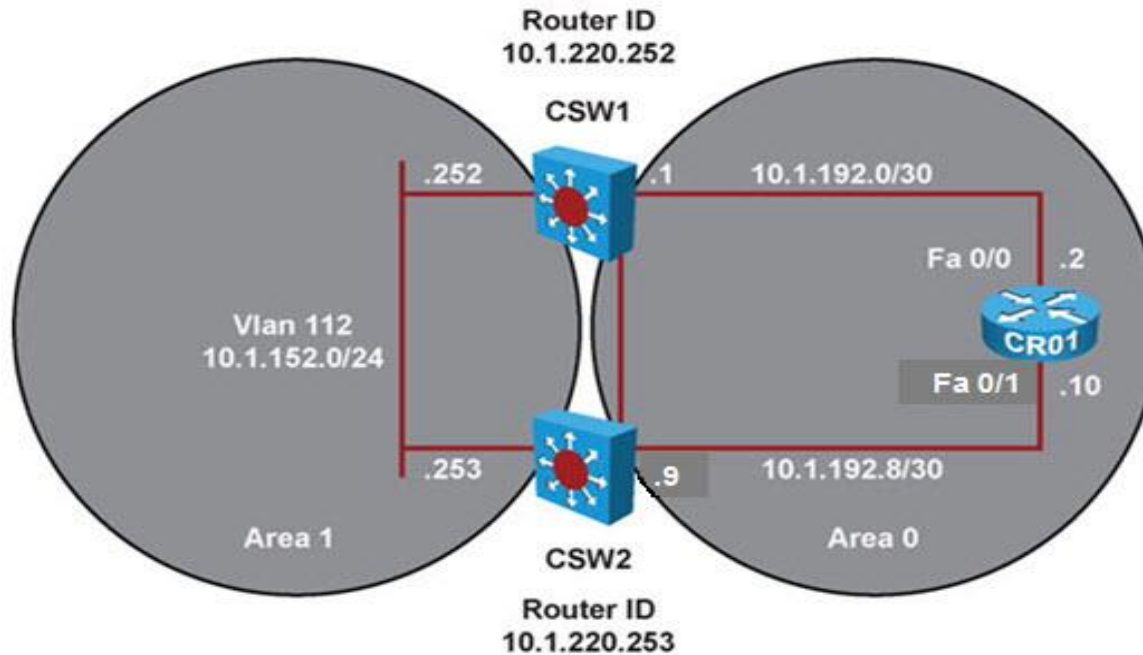
- Displays the transmission and reception of OSPF packets.
- Only the packet headers are displayed, not the content of the packets.
- Can be useful to verify if Hellos are sent and received as expected.

debug ip ospf events

- This command displays OSPF events such as reception and transmission of Hellos.
- Output also includes the establishment of neighbor relationships and the reception or transmission of LSAs.
- Can provide clues as to why neighbor Hellos might be ignored (mismatched parameters such as timers, area number, etc.).

OSPF Troubleshooting Example

Only one equal-cost OSPF path used by CR01



```
CR01# show ip route 10.1.152.0 255.255.255.0
```

```
Routing entry for 10.1.152.0/24
```


```
Known via "ospf 100", distance 110, metric 2, type inter area
```

```
Last update from 10.1.192.1 on FastEthernet0/0, 00:00:11 ago
```

```
Routing Descriptor Blocks:
```

```
* 10.1.192.1, from 10.1.220.252, 00:00:11 ago, via FastEthernet0/0  
Route metric is 2, traffic share count is 1
```


OSPF Troubleshooting Example - Cont.



```
CRO1# ping 10.1.192.9
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 10.1.192.9, timeout is 2 seconds:
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```


OSPF Troubleshooting Example - Cont.

The CR01 link-state database shows two Area 0 Type-3 summary LSAs for network 10.1.152.0, one from CSW1 and one From CSW2.

```
CR01# show ip ospf database summary 10.1.152.0
      OSPF Router with ID (10.1.220.1) (Process ID 100)
      Summary Net Link States (Area 0)
Routing Bit Set on this LSA
  LS age: 201
  Options: (No TOS-capability, DC, Upward)
  LS Type: Summary Links(Network)
  Link State ID: 10.1.152.0 (summary Network Number)
  Advertising Router: 10.1.220.252
  LS Seq Number: 80000001
  Checksum: 0x1C97
  Length: 28
  Network Mask: /24
  TOS: 0      Metric: 1

  LS age: 136
  Options: (No TOS-capability, DC, Upward)
  LS Type: Summary Links(Network)
  Link State ID: 10.1.152.0 (summary Network Number)
  Advertising Router: 10.1.220.253
  LS Seq Number: 80000001
  Checksum: 0x169C
  Length: 28
  Network Mask: /24
```

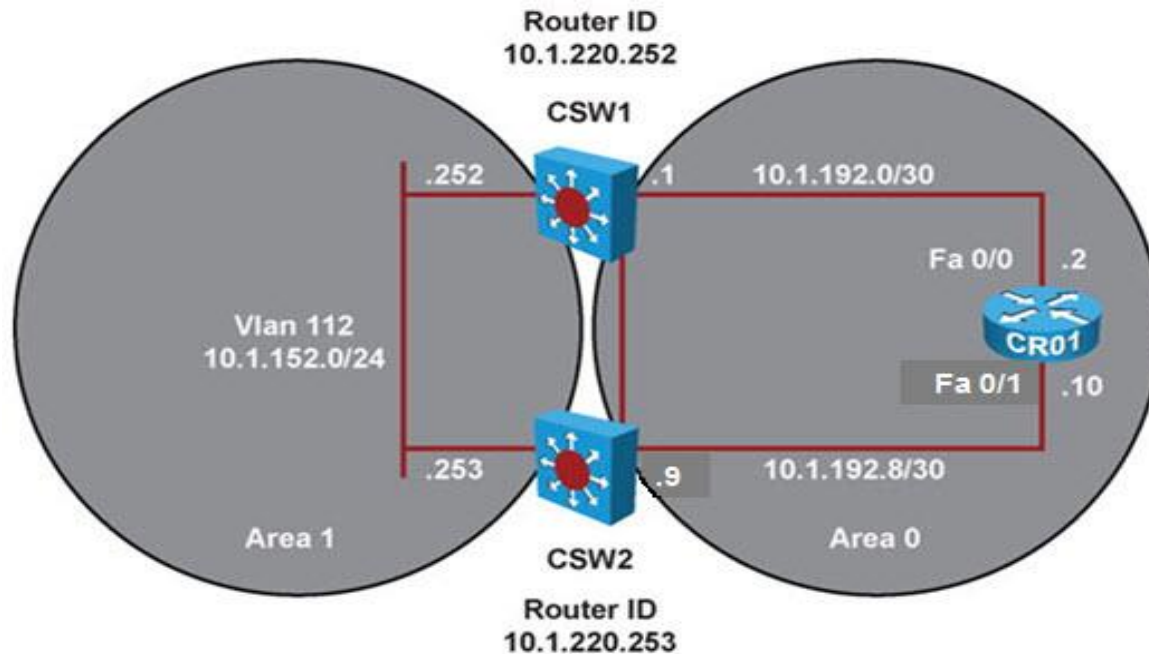
OSPF Troubleshooting Example - Cont.



```
CRO1# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.1.220.252	1	FULL/DR	00:00:33	10.1.192.1	FastEthernet0/0

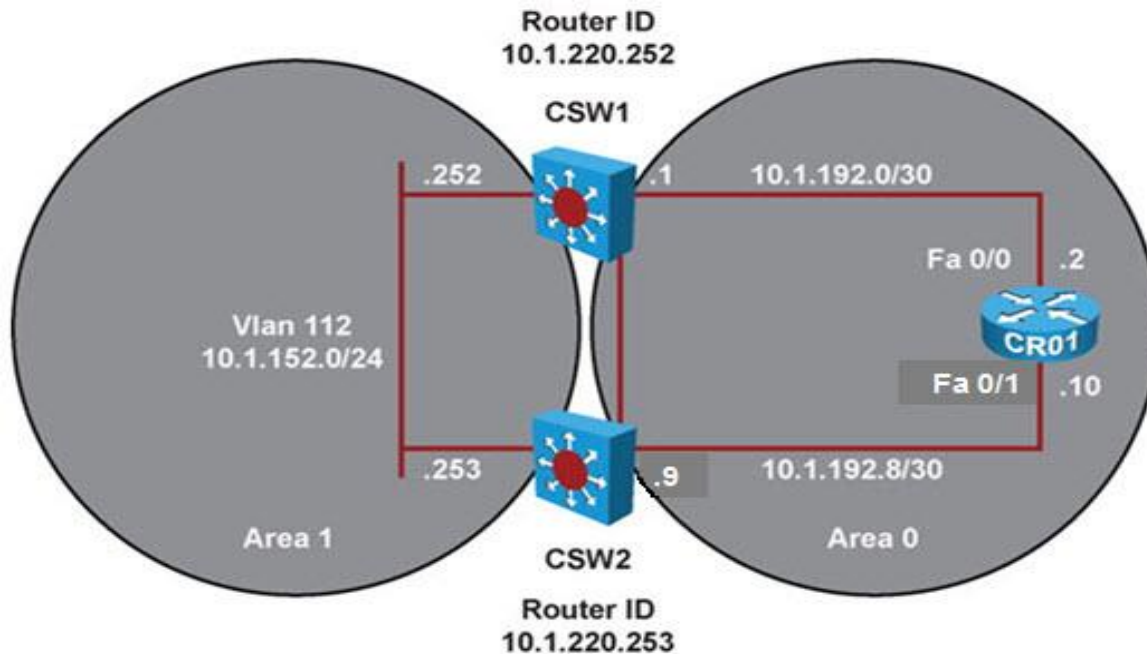
OSPF Troubleshooting Example - Cont.



```
CRO1# show ip ospf interface brief
```

Interface	PID	Area	IP Address/Mask	Cost	State	Nbrs
Lo0	100	0	10.1.220.1/32	1	LOOP	0/0
Fa0/0	100	0	10.1.192.2/30	1	BDR	1/1


OSPF Troubleshooting Example - Cont.



```
CRO1# show running-config | section router ospf
router ospf 100
  log-adjacency-changes
  network 10.1.192.2 0.0.0.0 area 0
  network 10.1.192.9 0.0.0.0 area 0
  network 10.1.220.1 0.0.0.0 area 0
```

OSPF Troubleshooting Example:

Correcting the network statement



```
CRO1(config)# router ospf 100
CRO1(config-router)# no network 10.1.192.9 0.0.0.0 area 0
CRO1(config-router)# network 10.1.192.10 0.0.0.0 area 0
```

OSPF Troubleshooting Example - Cont.

Results of **show** commands after correcting the OSPF network statement:


```
CRO1# show ip ospf interface brief
```

Interface	PID	Area	IP Address/Mask	Cost	State	Nbrs	F/C
Lo0	100	0	10.1.220.1/32	1	LOOP	0/0	
Fa0/1	100	0	10.1.192.10/30	1	BDR	1/1	
Fa0/0	100	0	10.1.192.2/30	1	BDR	1/1	

```
CRO1# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
10.1.220.253	1	FULL/DR	00:00:39	10.1.192.9	FastEthernet0/1
10.1.220.252	1	FULL/DR	00:00:31	10.1.192.1	FastEthernet0/0

OSPF Troubleshooting Example - Cont.



```
CRO1# show ip route 10.1.152.0 255.255.255.0
```

```
Routing entry for 10.1.152.0/24
```

```
Known via "ospf 100", distance 110, metric 2, type inter area
```

```
Last update from 10.1.192.9 on FastEthernet0/1, 00:00:29 ago
```

```
Routing Descriptor Blocks:
```

```
10.1.192.9, from 10.1.220.253, 00:00:29 ago, via FastEthernet0/1
```

```
Route metric is 2, traffic share count is 1
```

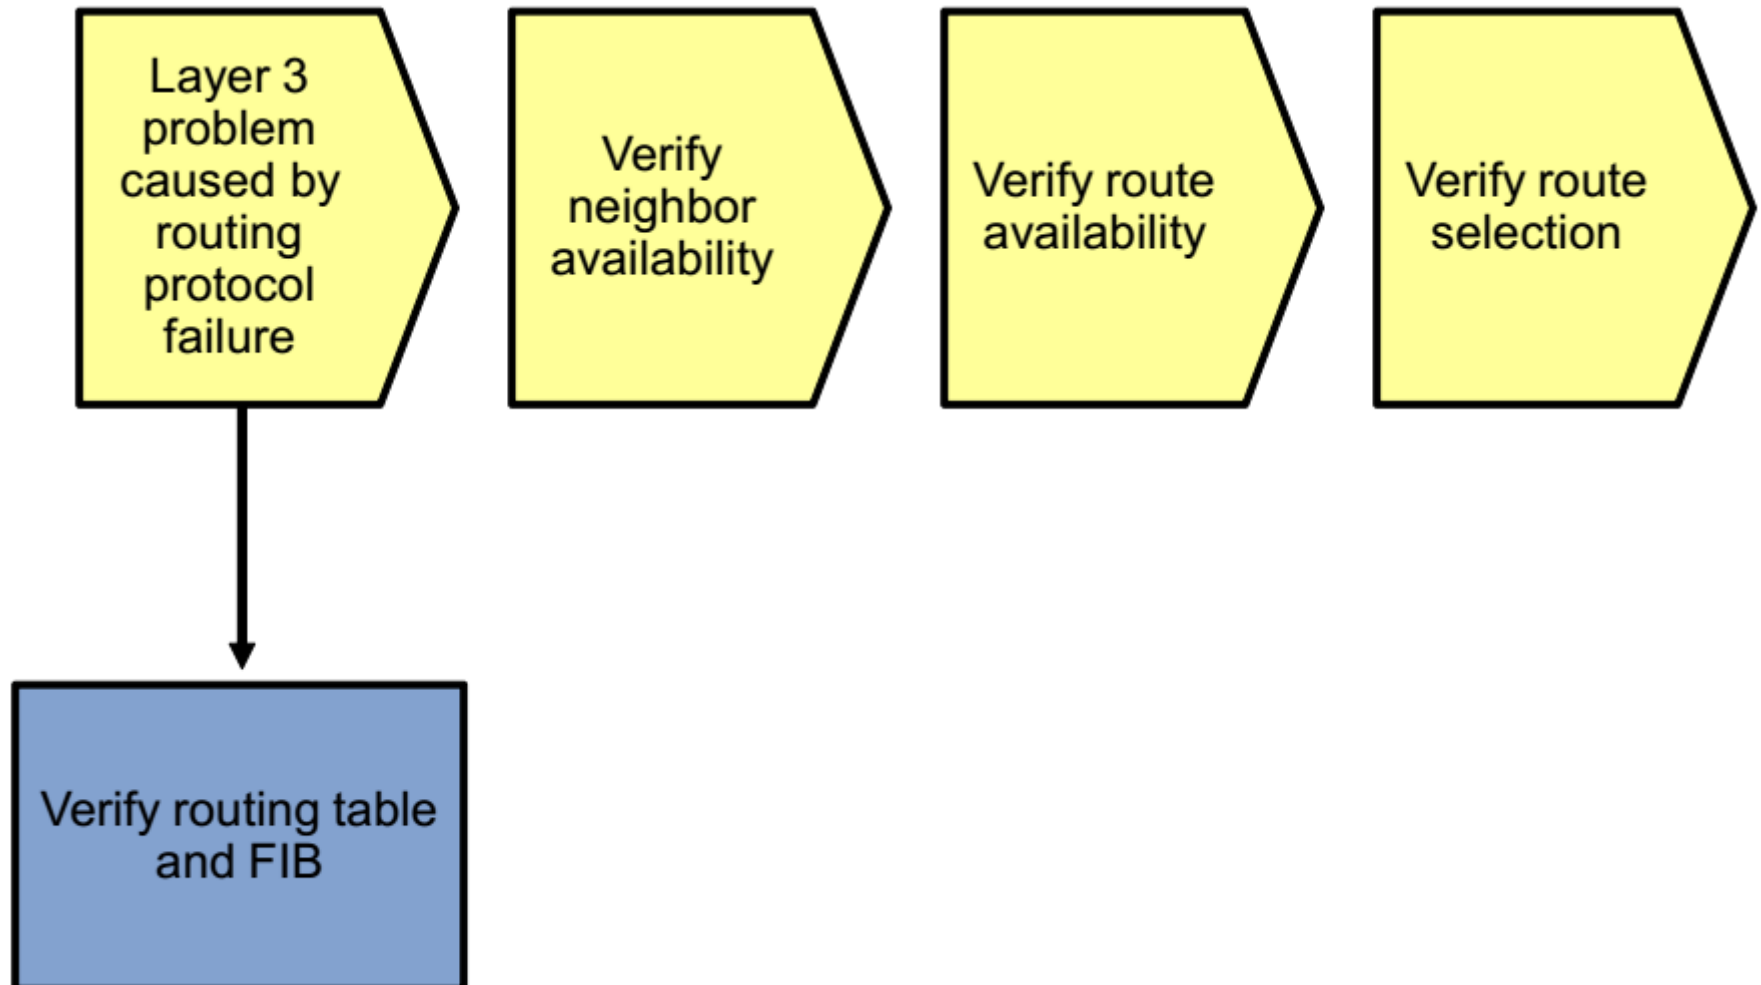
```
* 10.1.192.1, from 10.1.220.252, 00:00:29 ago, via FastEthernet0/0
```

```
Route metric is 2, traffic share count is 1
```

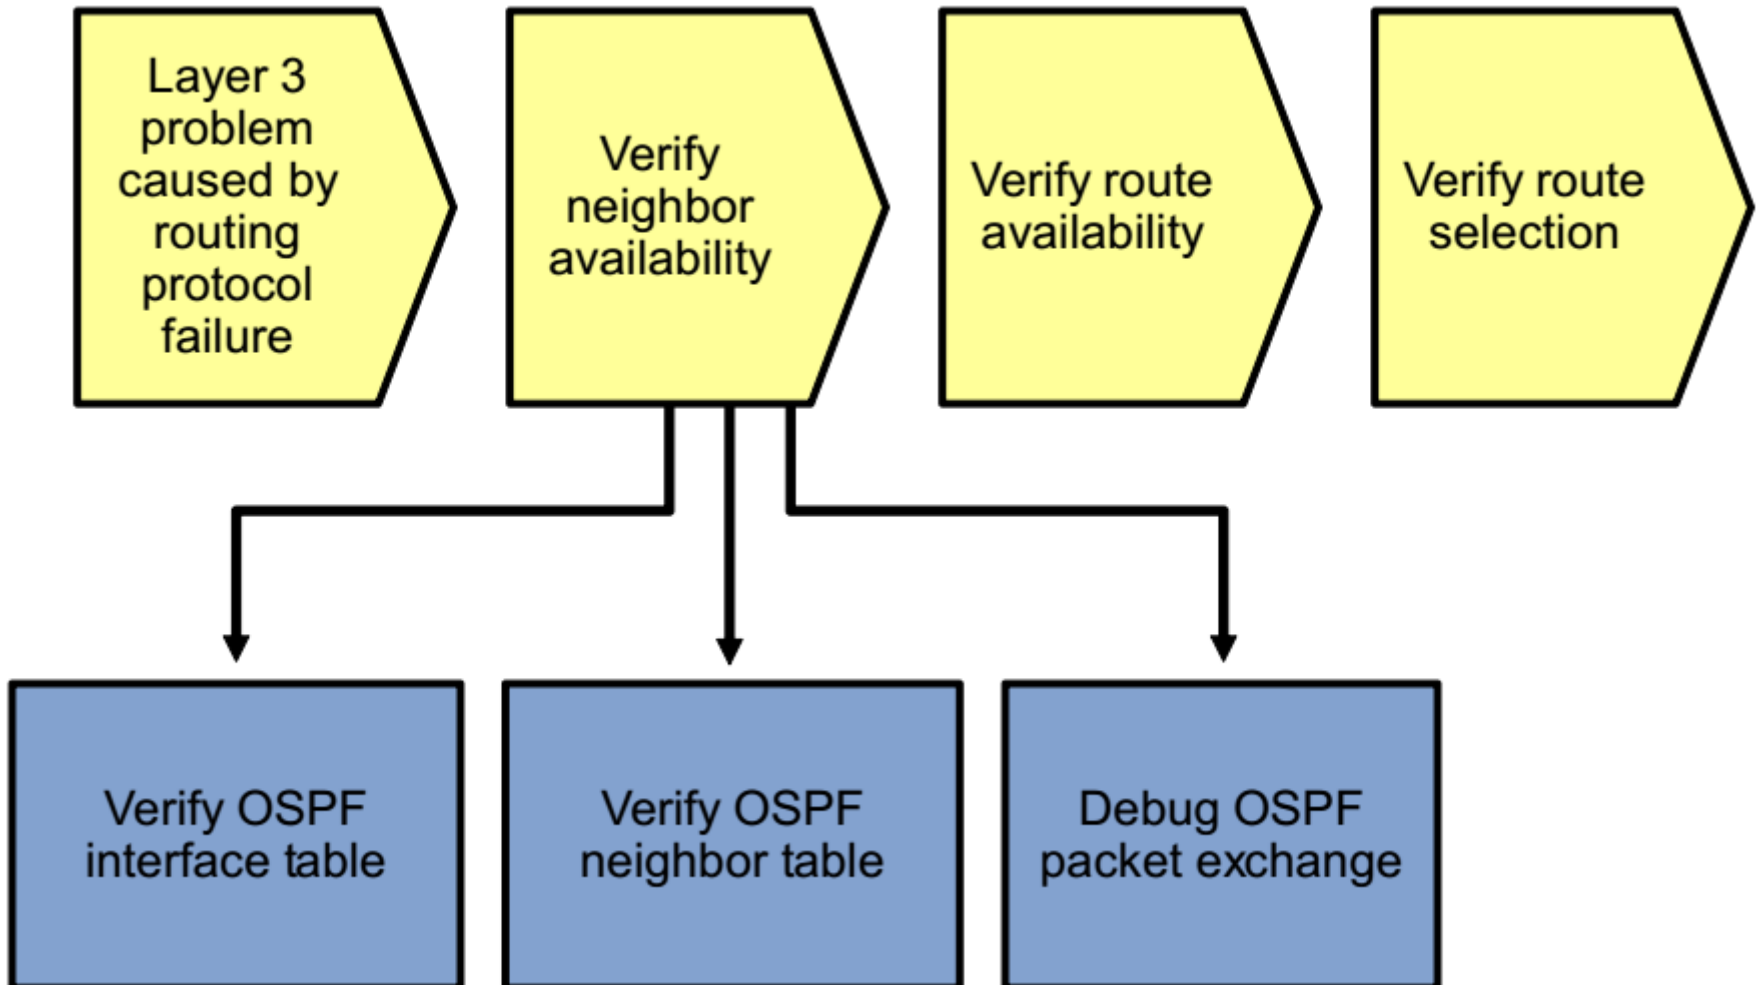

OSPF Troubleshooting Flow



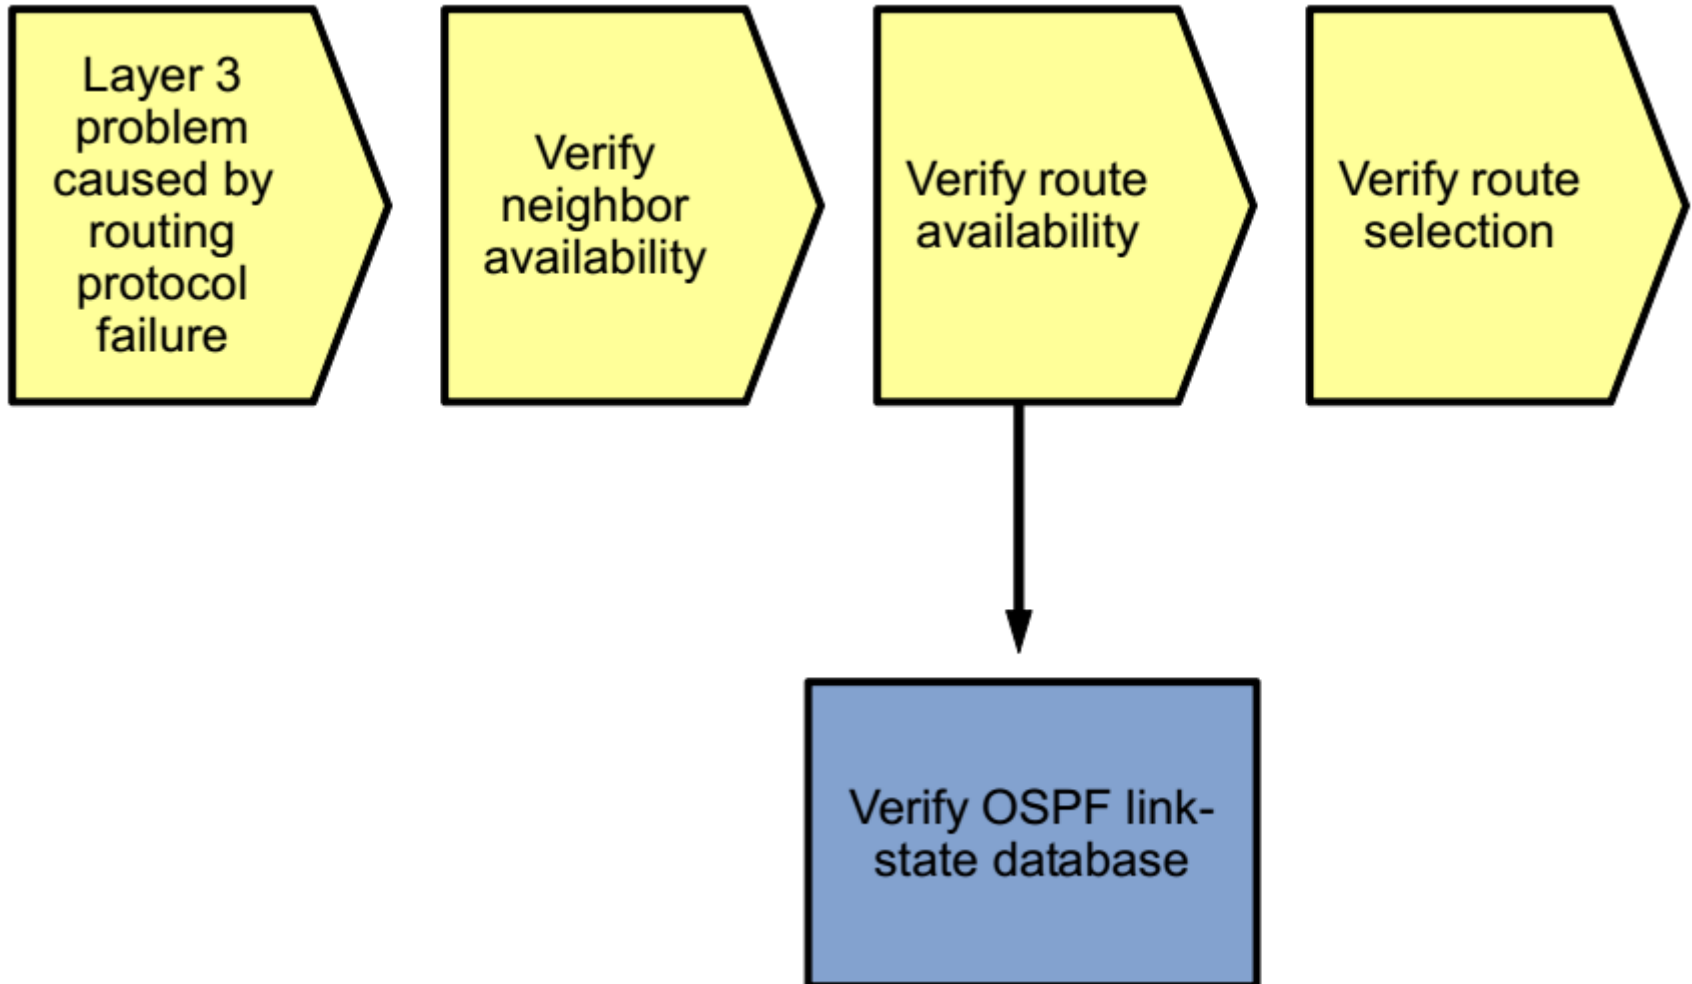
OSPF Troubleshooting Flow ①



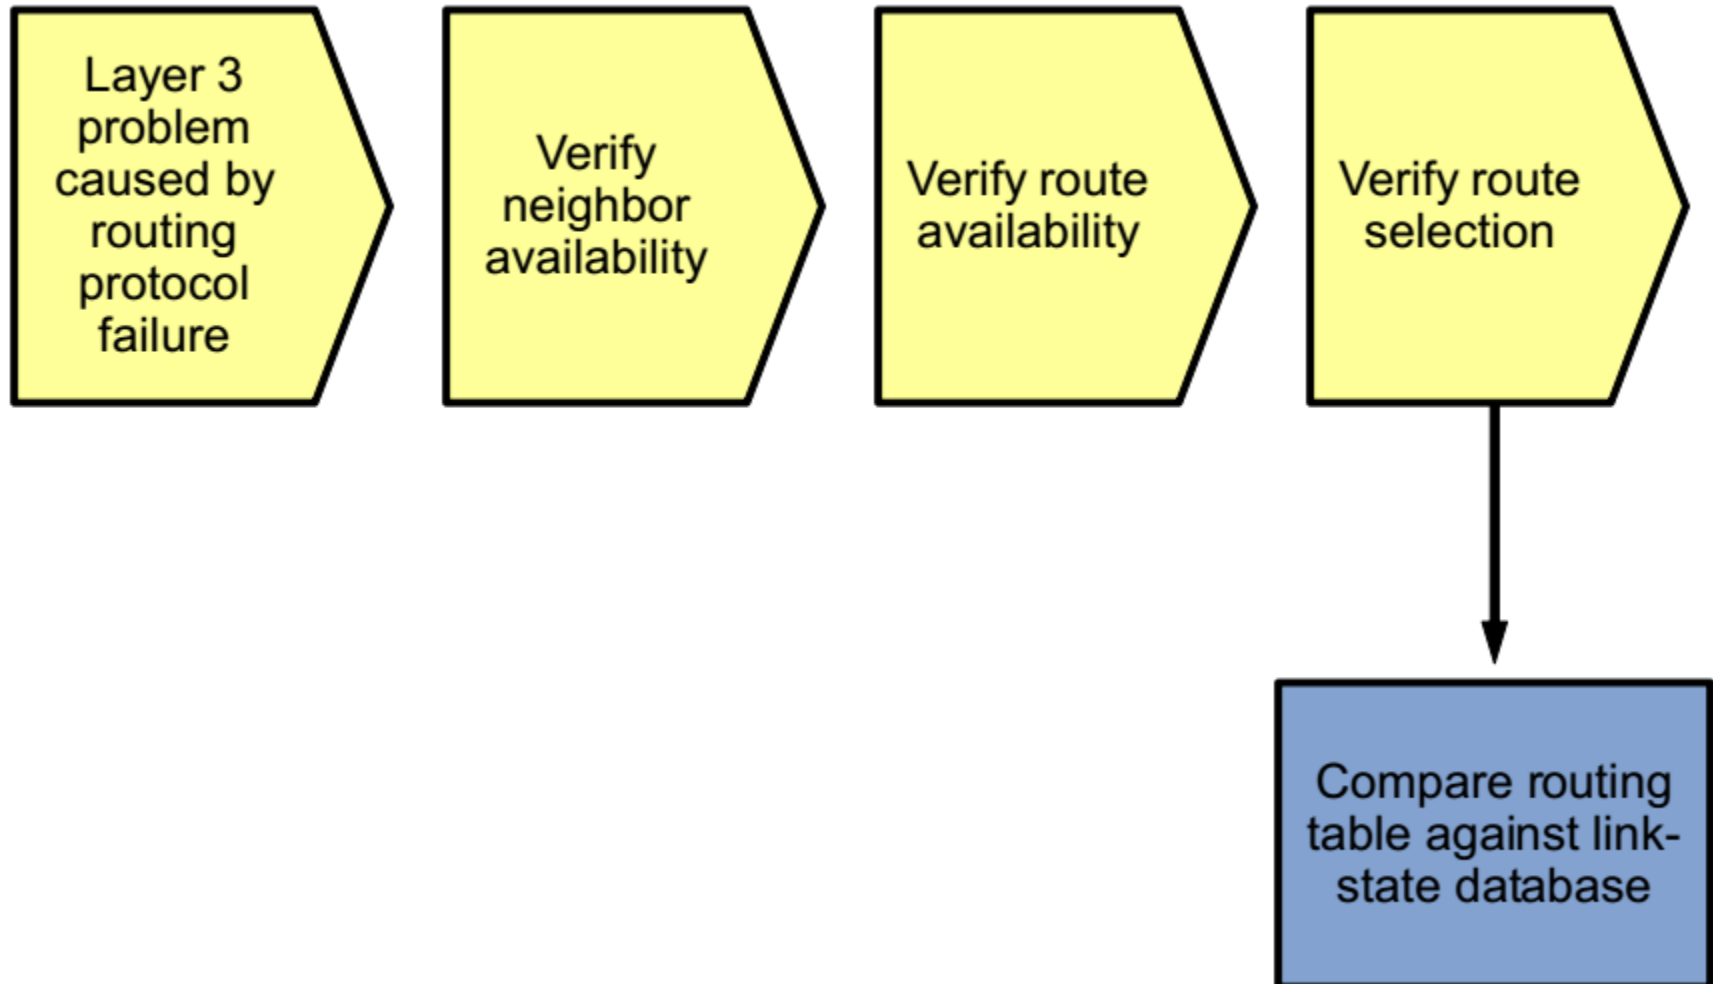
OSPF Troubleshooting Flow ②



OSPF Troubleshooting Flow ③



OSPF Troubleshooting Flow ④

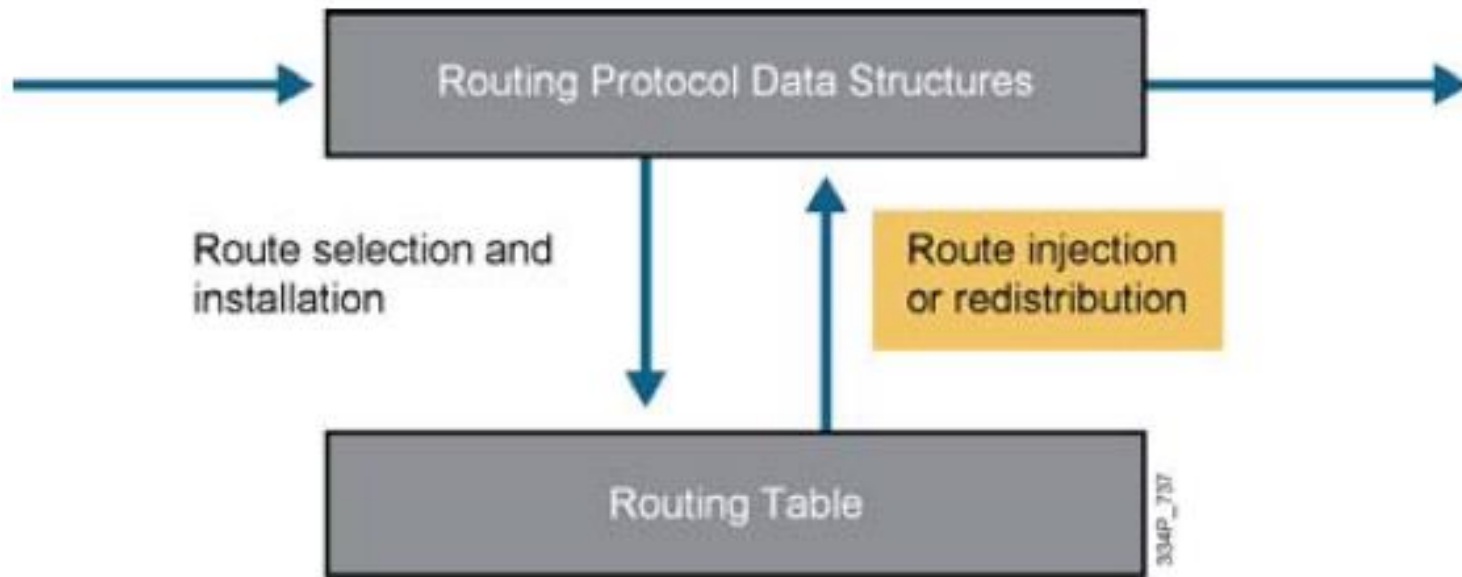


Troubleshooting Redistribution



Route Injection and Redistribution

- Routing protocols can carry two types of routes
 - **Internal**: injected locally based on directly connected subnets
 - **External**: redistributed from other routing source

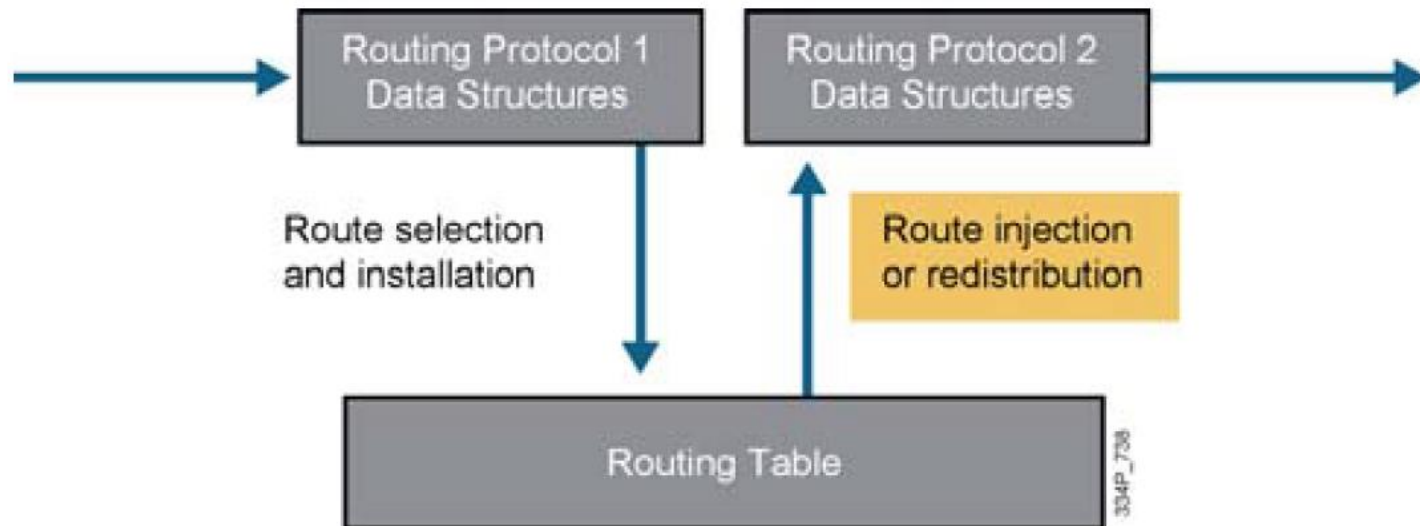


Route Redistribution Review

- Ideally, no more than one interior (intra-AS) routing protocol is used within an organization
- Organizational requirements (mergers, migrations) might dictate the use of multiple routing protocols
- Route redistribution between the different routing protocols may be necessary for IP connectivity between the different parts of the network
- Route redistribution adds an extra layer of complexity to a routed network
- It is important to understand the interactions between multiple routing protocols
- A network support engineer must be able to diagnose and resolve problems such as suboptimal routing and routing feedback that can occur when route redistribution is implemented

Redistribution and Routing Table

- Two important conditions for a prefix learned from one protocol (using redistribution) to be successfully advertised through another protocol:
 - The route needs to be installed in the routing table
 - A proper seed metric is assigned to the redistributed route
- Redistributed routes are taken from routing table
- Routes are not transferred directly between the routing protocol data structures of different protocols
- Route redistribution is **configured under the destination protocol**



Seed Metrics

- All routes exchanged via a routing protocol need to have an associated metric for that protocol
- There are no formulas to transform the metric for one protocol to a metric for a different protocol
- The starting metric or **seed metric** for protocol can be configured
- IF no seed metric is configure THEN a default is used
 - RIP and EIGRP default seed metrics are “unreachable”
 - Redistribution into these protocols will fail without explicit configuration of a seed metric. Knowledge of this is important when troubleshooting redistribution issues.

Common Redistribution Troubleshooting Tasks

- Troubleshooting the source routing protocol
 - Routes can only be redistributed if they are present in the routing table of the redistributing router
 - Confirm that the expected routes are learned on the redistributing router via the source protocol
- Troubleshooting route selection and installation
 - With bidirectional redistribution between routing protocols routing loops can be created
 - Suboptimal routing can occur causing routing instability requiring diagnosis
 - Changing the administrative distance or filtering routes to influence the route selection and installation process can often solve the problem
- Troubleshooting the redistribution process
 - If routes are in the routing table of the redistributing router, but not advertised by the redistributing protocol, verify the configuration of the redistribution process
 - Bad seed metrics, route filtering, or misconfigured routing protocol process or autonomous system numbers are common causes for the redistribution process to fail
- Troubleshooting the destination routing protocol
 - If the routing information is propagated using a protocol's routing update mechanisms, but not properly distributed to all routers in the destination routing domain, troubleshoot the routing exchange mechanisms for the destination protocol
 - Each routing protocol has its own methods of exchanging routing information, including external routing information
 - Determine if external routes are handled differently than internal routes. For example, OSPF external routes do not propagate into stub areas

Route Profiling

- **Route profiling** is a feature that can be helpful in diagnosing suspected route instability
- Entering the **ip route profile** command in global configuration mode enables this feature
- After enabling this feature, the router tracks the number of routing table changes that occurred over 5 second sampling intervals
- The **show ip route profile** command displays the results gathered by this feature
- The output can provide an indication of the overall stability of the routing table, without the need to enable a debug command

The show ip route profile Command

```
R1# show ip route profile
```

```
IP routing table change statistics:
```

```
Frequency of changes in a 5 second sampling interval
```

Change/ interval	Fwd-path change	Prefix add	Nexthop change	Pathcount change	Prefix refresh
0	96	96	96	96	96
1	0	0	0	0	0
2	0	0	0	0	0
3	0	0	0	0	0
4	0	0	0	0	0
5	0	0	0	0	0
10	0	0	0	0	0
15	0	0	0	0	0

```
output truncated
```

Verifying and Troubleshooting Route Propagation

show ip ospf database

- Displays the content of OSPF link-state database.

show ip eigrp topology

- Displays the content of the EIGRP topology table.

show ip route *network mask*

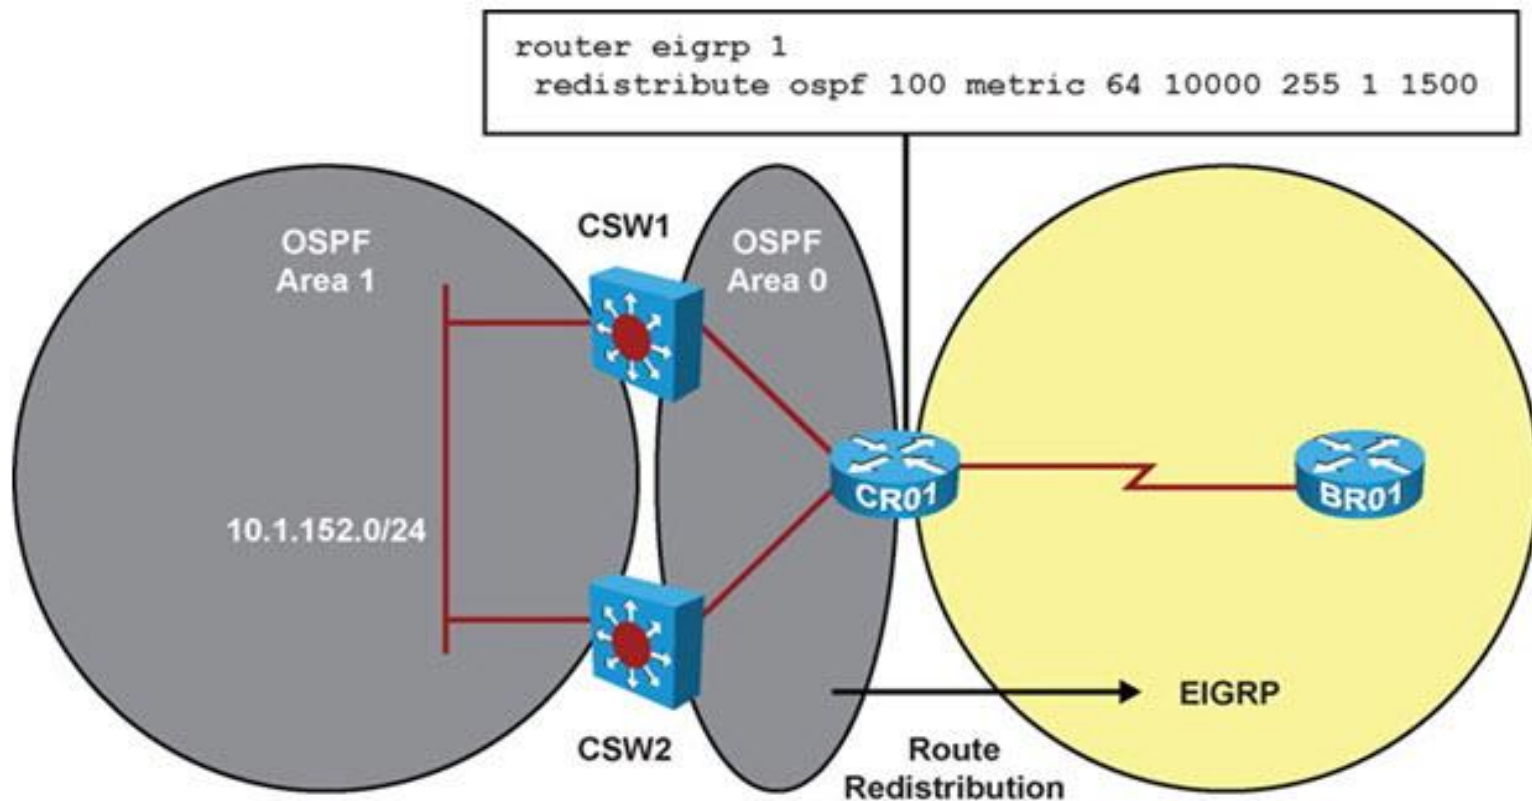
- Displays detailed information about specific routes installed in the routing table.

debug ip routing

- Displays routes being installed or removed from the routing table in real time.
- Can be very powerful when you are troubleshooting routing loops or flapping routes caused by route redistribution.

Redistribution Case-Study ①

- This example illustrates the redistribution process and the commands that can be used to verify it. The case does not revolve around a problem



Redistribution Case-Study ②

Router CRO1's OSPF database is displayed looking for LSA Type-3

```
CRO1# show ip ospf database | begin Summary
      Summary Net Link States (Area 0)
```

Link ID	ADV Router	Age	Seq#	Checksum
10.1.152.0	10.1.220.252	472	0x8000003B	0x00A7D1
10.1.152.0	10.1.220.253	558	0x8000003B	0x00A1D6

```
<output omitted>
```


Redistribution Case-Study ③

- The IP routing table for CR01 includes two OSPF paths to 10.1.152.0/24
- Both paths through switch CSW1 and switch CSW2 have been installed in the routing table because their costs are identical.
- The routing table also shows that this route has been marked for redistribution by EIGRP and the configured EIGRP seed metric is also listed.

```
CR01# show ip route 10.1.152.0 255.255.255.0
Routing entry for 10.1.152.0/24
  Known via "ospf 100", distance 110, metric 2, type inter area
  Redistributing via eigrp 1
  Advertised by eigrp 1 metric 64 10000 255 1 1500
  Last update from 10.1.192.9 on FastEthernet0/1, 00:28:24 ago
  Routing Descriptor Blocks:
    10.1.192.9, from 10.1.220.253, 00:28:24 ago, via FastEthernet0/1
      Route metric is 2, traffic share count is 1
    * 10.1.192.1, from 10.1.220.252, 00:28:24 ago, via FastEthernet0/0
      Route metric is 2, traffic share count is 1
```

Redistribution Case-Study ④

- The EIGRP topology table on router CR01 verifies that the route is being redistributed.
- The route was taken from the routing table and inserted into the topology table as an external route.
- The five components of the configured seed metric are listed.
- The route was originated by the OSPF protocol with process number 100 and was injected into EIGRP by the router with EIGRP router ID 10.1.220.1 (which is the local router, CRO1).

```
CR01# show ip eigrp topology 10.1.152.0 255.255.255.0
IP-EIGRP (AS 1): Topology entry for 10.1.152.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 42560000
  Routing Descriptor Blocks:
    10.1.192.9, from Redistributed, Send flag is 0x0
      Composite metric is (42560000/0), Route is External
      Vector metric:
        Minimum bandwidth is 64 Kbit
        Total delay is 100000 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 0
      External data:
        Originating router is 10.1.220.1 (this system)
        AS number of route is 100
        External protocol is OSPF, external metric is 2
        Administrator tag is 0 (0x00000000)
```

Redistribution Case-Study ⑤

- The external information that router CR01 added to the EIGRP topology table during redistribution, is passed along to router BR01 within the EIGRP routing updates.
- In the output of the topology table on router BR01, the originating router and routing protocol are still visible.

```
BR01# show ip eigrp topology 10.1.152.0 255.255.255.0
IP-EIGRP (AS 1): Topology entry for 10.1.152.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 43072000
  Routing Descriptor Blocks:
    10.1.193.1 (Serial0/0/1), from 10.1.193.1, Send flag is 0x0
      Composite metric is (43072000/42560000), Route is External
      Vector metric:
        Minimum bandwidth is 64 Kbit
        Total delay is 120000 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 1
      External data:
        Originating router is 10.1.220.1
        AS number of route is 100
        External protocol is OSPF, external metric is 2
        Administrator tag is 0 (0x00000000)
```

Redistribution Case-Study ⑥

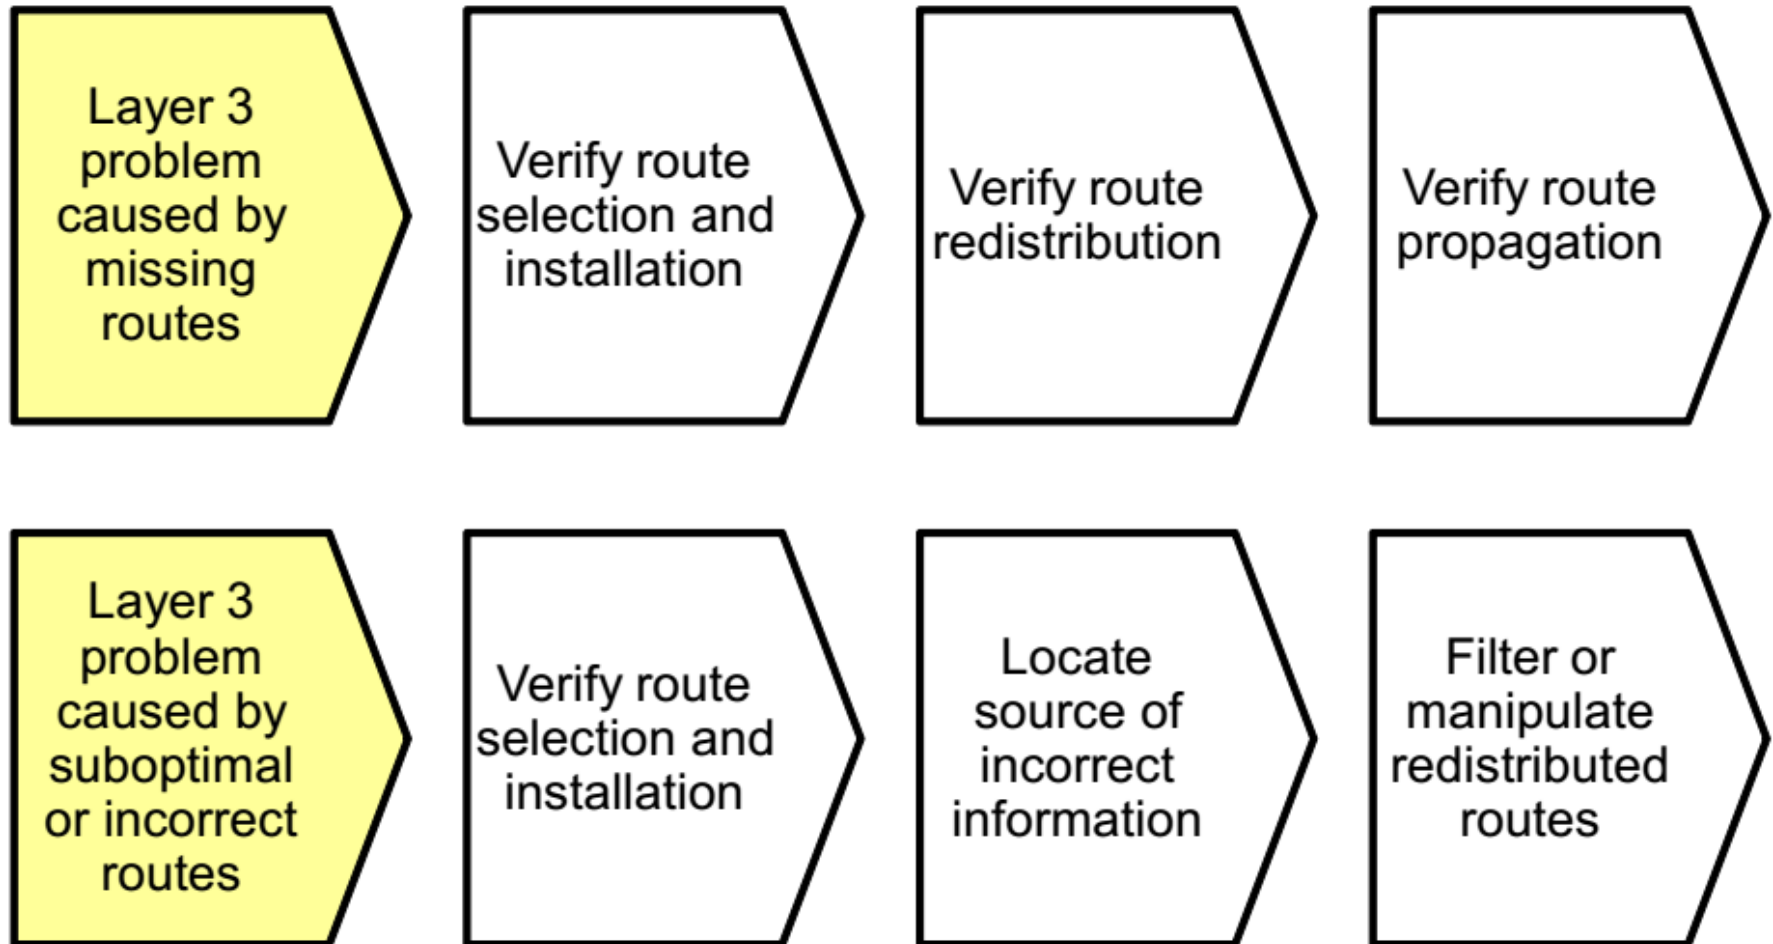
- On router BRO1, EIGRP selects the 10.1.152.0/24 route learned from CR01 and installs it in the IP routing table.
- The route is marked as an EIGRP external route and has a corresponding administrative distance of 170.
- The external information present in the EIGRP topology table, such as the originating router and protocol, is not carried into the routing table.

```
BRO1# show ip route 10.1.152.0 255.255.255.0
Routing entry for 10.1.152.0/24
  Known via "eigrp 1", distance 170, metric 43072000, type external
  Redistributing via eigrp 1
  Last update from 10.1.193.1 on Serial0/0/1, 00:00:35 ago
  Routing Descriptor Blocks:
    * 10.1.193.1, from 10.1.193.1, 00:00:35 ago, via Serial0/0/1
      Route metric is 43072000, traffic share count is 1
      Total delay is 120000 microseconds, minimum bandwidth is 64 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 3/255, Hops 1
```

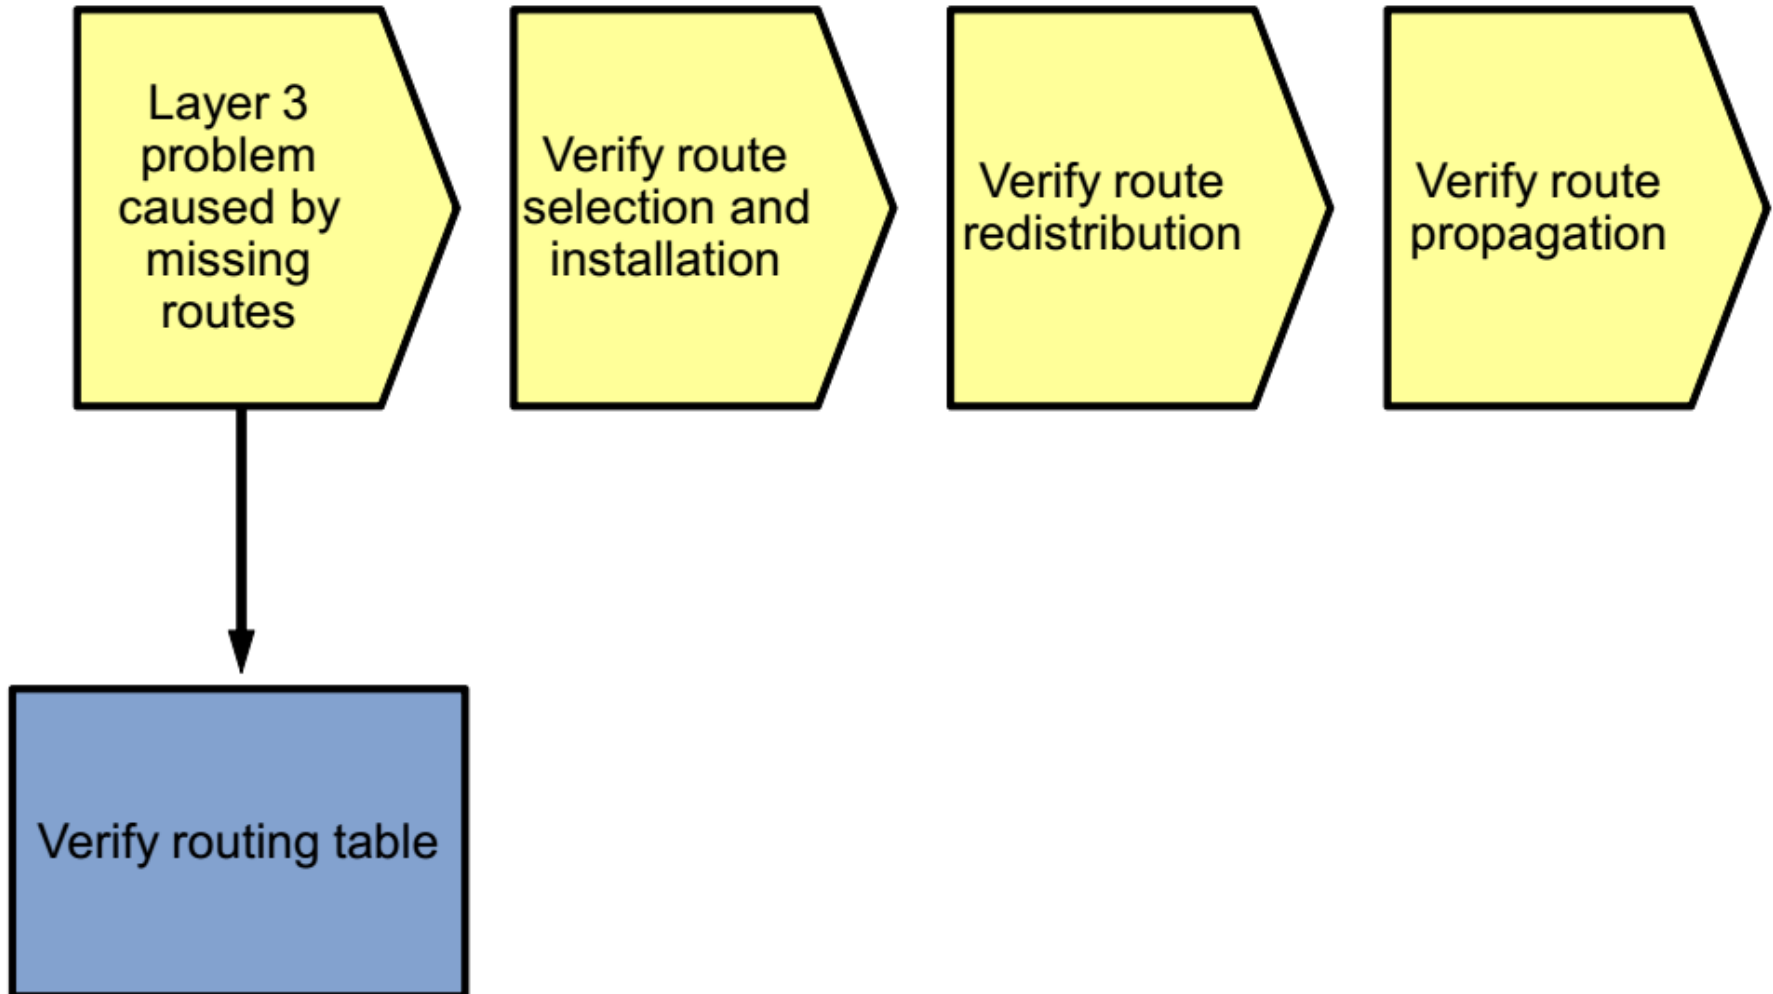
Redistribution Troubleshooting Flow



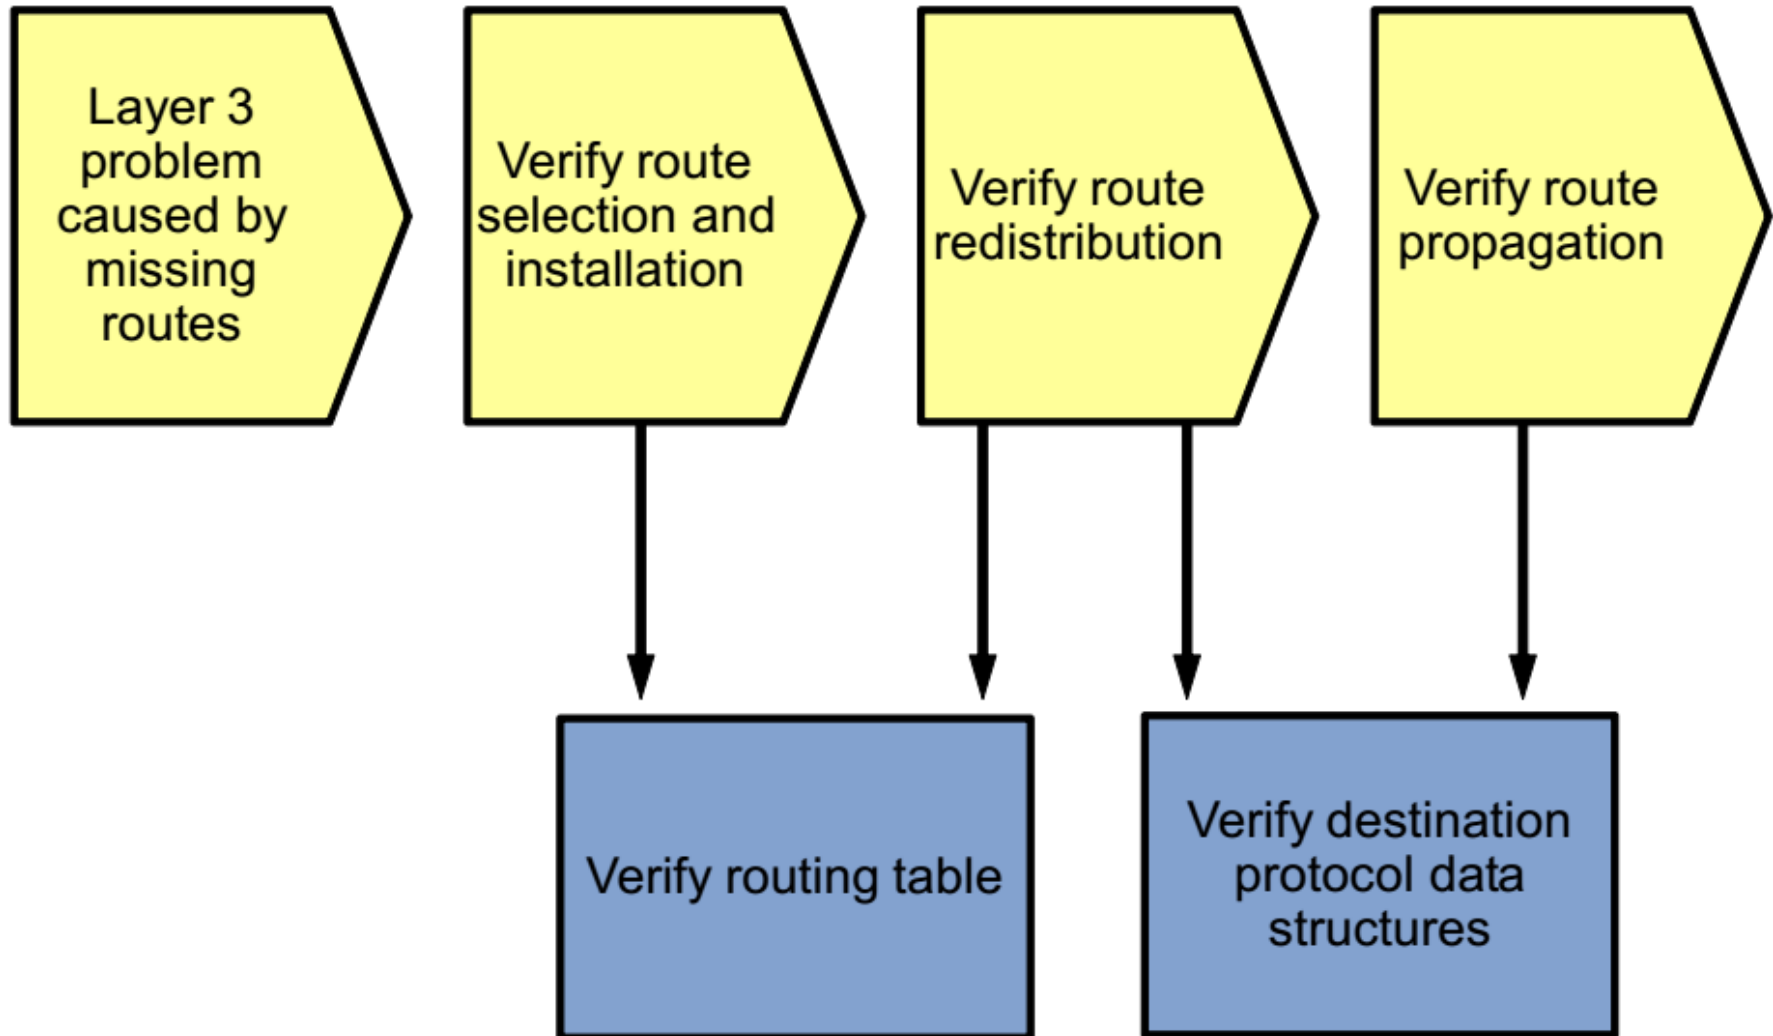
Redistribution Troubleshooting Flow ①



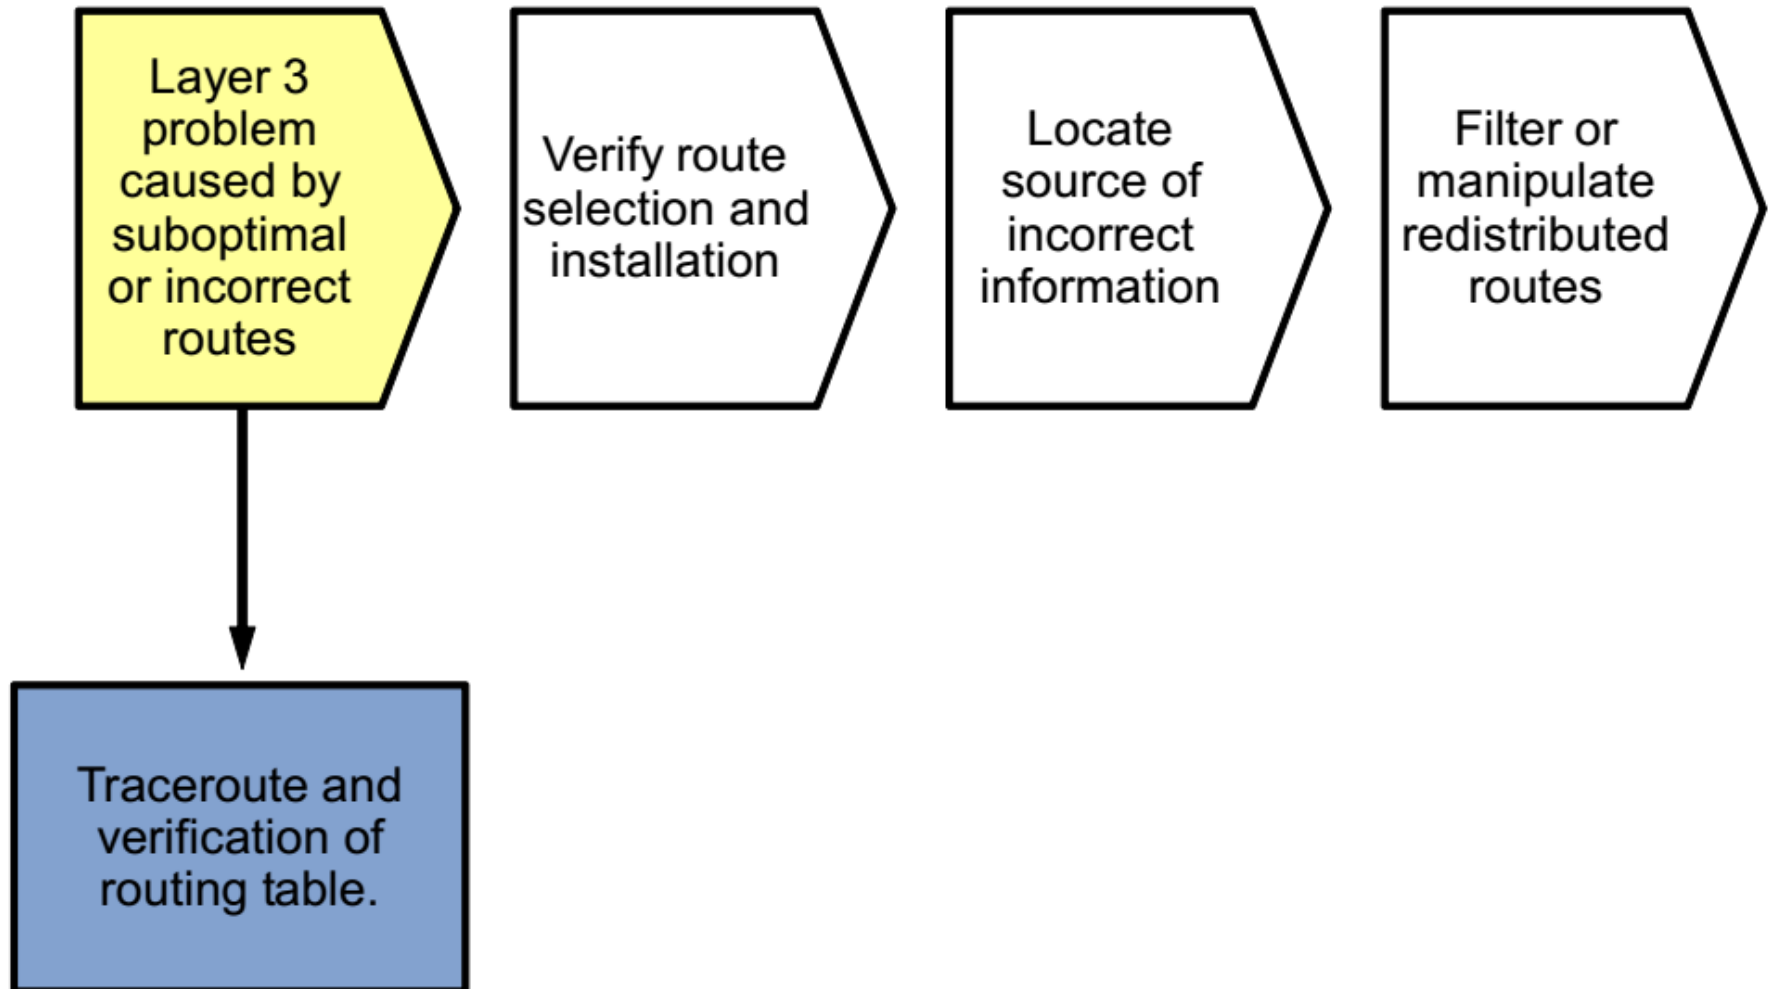
Redistribution Troubleshooting Flow ②



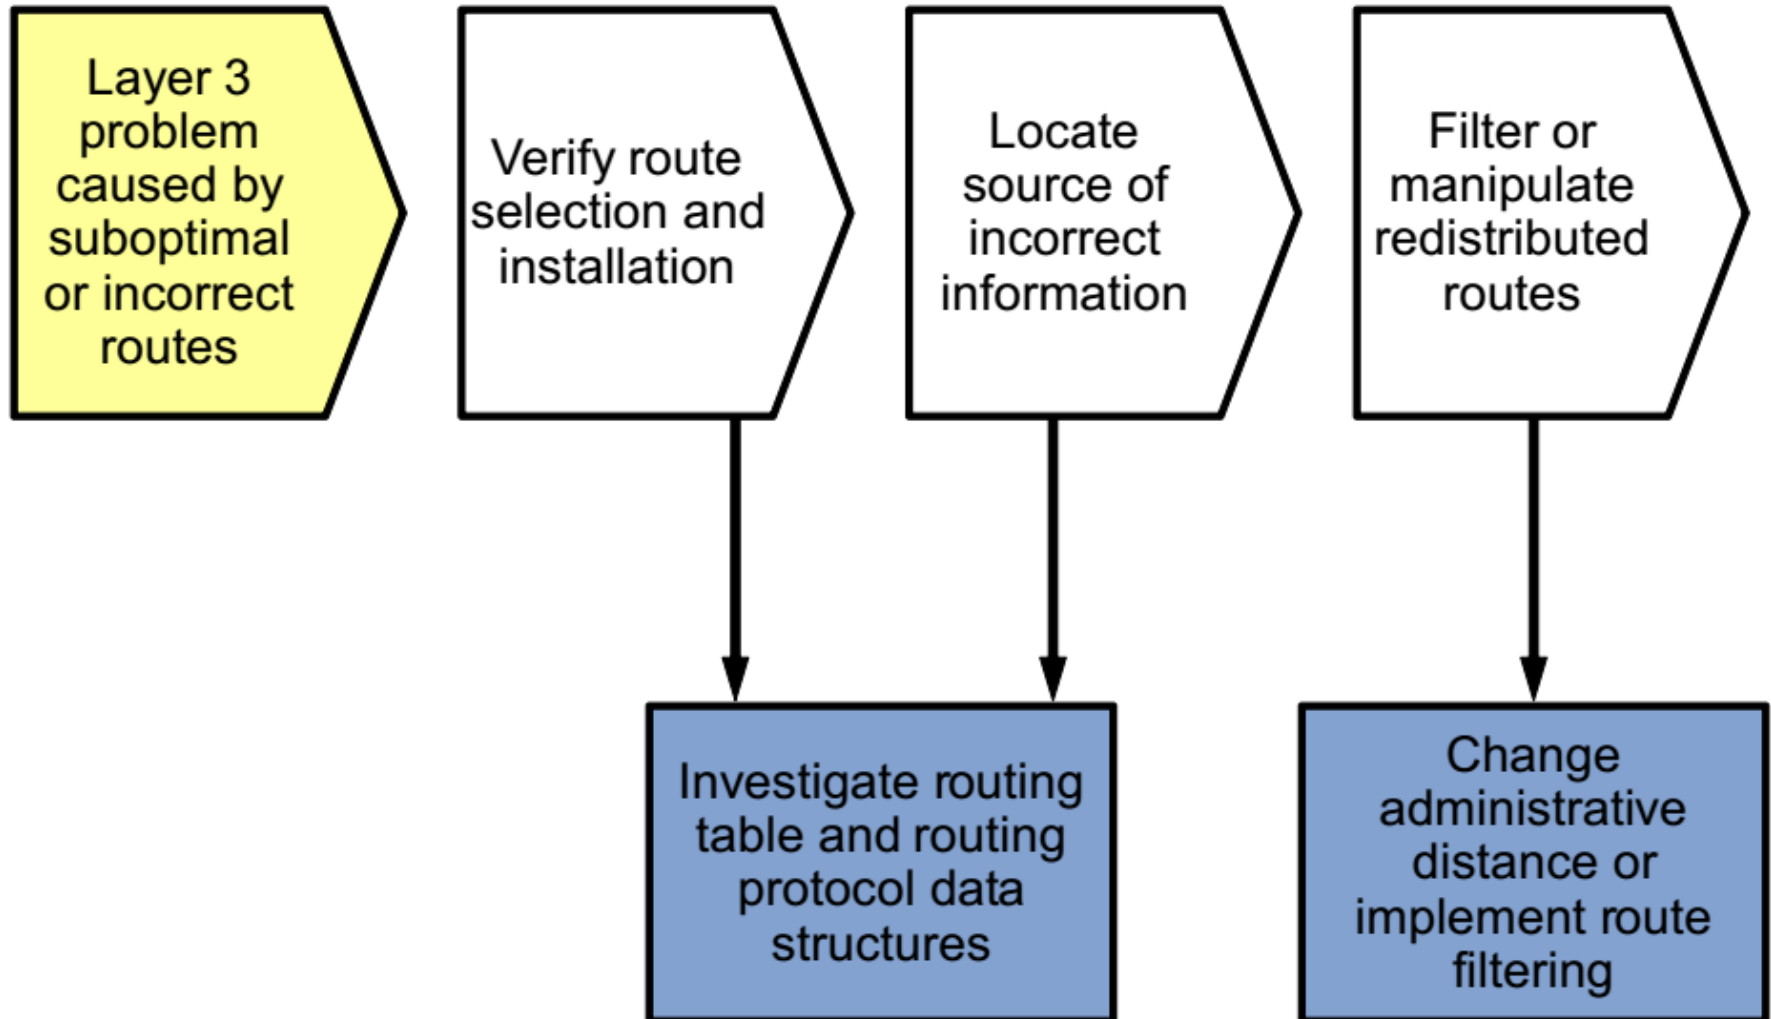
Redistribution Troubleshooting Flow ③



Redistribution Troubleshooting Flow ④



Redistribution Troubleshooting Flow



Troubleshooting EIGRP



IGP Routing Protocol Review

Routing protocols consist of the following elements and processes:

- Reception of routing information from neighbors
- Routing protocol data structures
- Route injection or redistribution
- Route selection and installation
- Transmission of routing information to neighbors

EIGRP Review

EIGRP stores its operational data, configured parameters, and statistics in three main data structures:

- **Interface table:**

- Lists all interfaces enabled for the processing of EIGRP packets.
- Passive interfaces are not listed in this table.

- **Neighbor table:**

- Keeps track of active EIGRP neighbors.
- Based on reception of hello packets.

- **Topology table:**

- Holds routes received from neighboring routers, locally injected, or redistributed into EIGRP.
- For each prefix, EIGRP selects the best path (successor route) for the IP routing table.
- EIGRP's best path selection is based on the Diffusing Update Algorithm (DUAL).
- Multiple paths with the same metric that satisfy the Feasibility Condition can be selected for installation in the routing table.
- Routes with a higher metric are not selected for installation in the routing table, unless unequal cost load balancing has been enabled.

EIGRP Review – Cont.

EIGRP uses an incremental update process.

- When the adjacency is first established, each router sends a full update to its neighbor router.
- All prefixes for which there a successor in the topology table are sent.
- After initial exchange, routing updates will only be sent due to changes on the networks.
- Changes can be caused by:
 - Changes in connectivity (such as loss or discovery of a link or neighbor , modification of an interface metric, changes to the split horizon or next-hop modification on an interface, stub routing activation/deactivation)
 - Addition of new interfaces to EIGRP
 - Implementation of route summarization
 - Implementation of route filtering
 - Implementation of route redistribution

Monitoring EIGRP with show commands

- To gather information from the EIGRP data structures use the following **show** commands:
- **show ip eigrp interfaces:**
 - Displays the list of interfaces that have been activated for EIGRP processing.
- **show ip eigrp neighbors:**
 - Lists all neighbors that have been discovered by this router on its active EIGRP interfaces.
- **show ip eigrp topology:**
 - Displays the content of the EIGRP topology table. To select a specific entry from the table, the network and mask can be provided as an option to the command.


Monitoring EIGRP with debug Commands

To observe the real-time EIGRP information exchange use the following debug commands:

- **debug ip routing:**
 - Not specific to EIGRP.
 - Displays changes made to the routing table, such as installation or removal of routes.
 - Can be useful in diagnosing routing protocol instabilities.
- **debug eigrp packets:**
 - Displays the transmission and reception of EIGRP packets.
 - All packets can be displayed, or packets of a particular type, such as hellos, updates, queries, and replies can be selected.
- **debug ip eigrp:**
 - Displays EIGRP routing events, such as updates, queries, and replies sent to or received from neighbors.
- **debug ip eigrp neighbor *as-number ip-address*:**
 - Limits output to information that is associated with the specified neighbor.
- **debug ip eigrp *as-number network mask*:**
 - Limits output to information that is associated with the network specified by the network and mask options.

EIGRP Troubleshooting Example:

Packets from BR01 to CR01 Lo0 take wrong path.



```
BR01# traceroute 10.1.220.1
```

```
Type escape sequence to abort.
```

```
Tracing the route to crol.mgmt.tshoot.local (10.1.220.1)
```

```
 1 10.1.163.130 0 msec 0 msec 0 msec
 2 10.1.194.5 12 msec 12 msec *
```

```
BR01# ping 10.1.194.1
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 10.1.194.1, timeout is 2 seconds:
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/29/32 ms
```

EIGRP Troubleshooting Example – Cont.

EIGRP **show** commands indicate that there is only one BR01 topology entry for CR01 Lo0 and that BR01 and CR01 are not EIGRP neighbors.

```
BR01# show ip eigrp topology 10.1.220.1 255.255.255.255
```

```
IP-EIGRP (AS 1): Topology entry for 10.1.220.1/32
```

```
State is Passive, Query origin flag is 1, 1 Successor(s), FD is 40642560
```

```
Routing Descriptor Blocks:
```

```
10.1.163.130 (FastEthernet0/1.30), from 10.1.163.130, Send flag is 0x0
```

```
Composite metric is (40642560/40640000), Route is Internal
```

```
Vector metric:
```

```
Minimum bandwidth is 64 Kbit
```

```
Total delay is 25100 microseconds
```

```
Reliability is 255/255
```

```
Load is 1/255
```

```
Minimum MTU is 1500
```

```
Hop count is 2
```

```
BR01# show ip eigrp neighbors
```

```
IP-EIGRP neighbors for process 1
```

H	Address	Interface	Hold (sec)	Uptime	SRTT (ms)	RTO	Q Cnt	Seq Num
0	10.1.163.130	Fa0/1.30	12	00:09:56	4	200	0	585

EIGRP Troubleshooting Example – Cont.

Only the BR01 Fa0/1.30 interface is participating in EIGRP. The **show run** command reveals that network statement for 10.1.194.1 is the problem.

```
BR01# show ip eigrp interfaces
```

```
IP-EIGRP interfaces for process 1
```

Interface	Peers	Xmit Queue Un/Reliable	Mean SRTT	Pacing Time Un/Reliable	Multicast Flow Timer	Pending Routes
Fa0/1.30	1	0/0	4	0/1	50	0

```
BR01# show running-config | section router eigrp
```

```
router eigrp 1
```

```
network 10.1.163.129 0.0.0.0
```

```
network 10.1.194.1 0.0.0.0
```

```
no auto-summary
```

EIGRP Troubleshooting Example – Cont.

After correcting the EIGRP **network** statement, both BR01 interfaces are participating in EIGRP and BR02 and CR01 are BR01 neighbors.

```
BR01# show ip eigrp interfaces
```

```
IP-EIGRP interfaces for process 1
```

Interface	Peers	Xmit Queue Un/Reliable	Mean SRTT	Pacing Time Un/Reliable	Multicast Flow Timer	Pending Routes
Fa0/1.30	1	0/0	1	0/1	50	0
Se0/0/0.111	1	0/0	707	10/380	4592	0

```
BR01# show ip eigrp neighbors
```

```
IP-EIGRP neighbors for process 1
```

H	Address	Interface	Hold (sec)	Uptime	SRTT (ms)	RTO	Q Cnt	Seq Num
1	10.1.194.1	Se0/0/0.111	14	00:10:10	707	4242	0	783
0	10.1.163.130	Fa0/1.30	12	01:34:49	1	200	0	587

EIGRP Troubleshooting Example – Cont.

The new EIGRP Topology table after corrections were made.

```
BRO1# show ip eigrp topology 10.1.220.1 255.255.255.255
IP-EIGRP (AS 1): Topology entry for 10.1.220.1/32
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 40640000
  Routing Descriptor Blocks:
    10.1.194.1 (Serial0/0/0.111), from 10.1.194.1, Send flag is 0x0
      Composite metric is (40640000/128256), Route is Internal
      Vector metric:
        Minimum bandwidth is 64 Kbit
        Total delay is 25000 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 1
    10.1.163.130 (FastEthernet0/1.30), from 10.1.163.130, Send flag is 0x0
      Composite metric is (40642560/40640000), Route is Internal
      Vector metric:
        Minimum bandwidth is 64 Kbit
        Total delay is 25100 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
```


EIGRP Troubleshooting Example – Cont.

The IP routing table after corrections were made.

```
BRO1# show ip route 10.1.220.1 255.255.255.255
Routing entry for 10.1.220.1/32
  Known via "eigrp 1", distance 90, metric 40640000, type internal
  Redistributing via eigrp 1
  Last update from 10.1.194.1 on Serial0/0/0.111, 00:20:55 ago
  Routing Descriptor Blocks:
  * 10.1.194.1, from 10.1.194.1, 00:20:55 ago, via Serial0/0/0.111
    Route metric is 40640000, traffic share count is 1
    Total delay is 25000 microseconds, minimum bandwidth is 64 Kbit
    Reliability 255/255, minimum MTU 1500 bytes
    Loading 1/255, Hops 1
```

EIGRP Troubleshooting Example – Cont.

Traceroute to CR01 Lo0 now shows correct path.



```
BR01# traceroute 10.1.220.1
```

```
Type escape sequence to abort.
```

```
Tracing the route to cro1.mgmt.tshoot.local (10.1.220.1)
```

```
 1 10.1.194.1 16 msec 12 msec *
```

Troubleshooting BGP



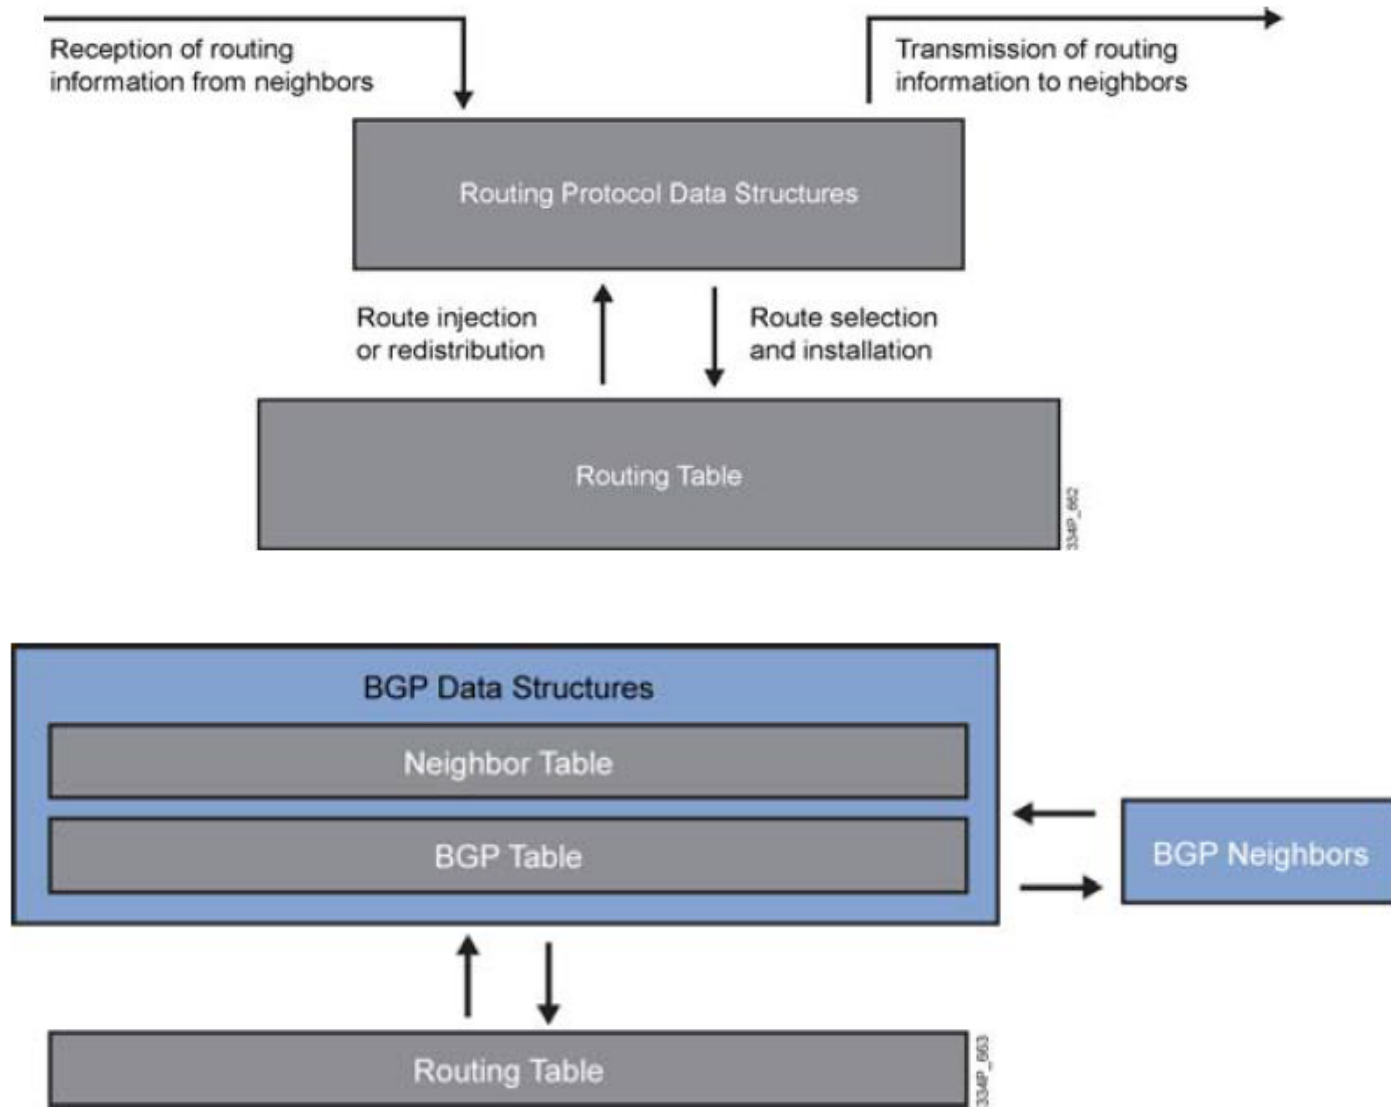
Overview

- BGP is classified as an **exterior gateway protocol** or an inter-autonomous-system (inter-AS) routing protocol
- Exchanges routing information with external networks (other autonomous systems), such as ISPs
- Plays a different role in enterprise networks as compared to IGPs, such as EIGRP or OSPF
 - Not used to find the best paths within the enterprise network
 - Used to **implement routing policies to control the flow of traffic** to and from external networks
- BGP is a routing protocol similar to IGPs in that it:
 - 1) exchanges information about reachability of prefixes with other BGP routers;
 - 2) selects the best path for each of the prefixes that it has learned about;
 - 3) offers the best paths to the routing table.

Route Processing ①

- Differentiate inter-AS neighborship (**internal BGP = iBGP**) and peering with router in different AS (**external BGP = eBGP**)
- Reception of routing information from neighbors
 - Neighbors need not be directly connected
 - They are manually configured, not discovered through a hello protocol
 - Two BGP routers that exchange information are commonly referred to as **peers** or **neighbors**
- A **TCP session on port 179** is established between neighbors to exchange routing information
- Routing protocol data structures
 - **Neighbor table**: Keeps track of the state of configured neighbors
 - **BGP table**: Stores all the prefixes, including those received from the neighbors (subtables like Adj-RIB-In/Out, Loc-RIB)

Route Processing ②



Route Processing ③

- BGP does not automatically inject any routes into the BGP table
 - Import of local networks is done via **network** command or redistribution
 - In both cases, a prefix needs to be present in the IP routing table before it can be advertised to BGP neighbors.
- Routes learned from neighbors are placed in the BGP table and can be advertised out to other BGP neighbors
- Routes learned from iBGP neighbors are subject to the synchronization rule, unless synchronization is off
 - Synchronization means that networks for iBGP need to be present in routing table and learned first via IGP protocol
 - Nowadays is synchronization disabled by default

Route Processing ④

- Route selection and installation
 - BGP compares paths received from different neighbors and **selects the best one for each prefix**
 - Similar to other routing protocols, BGP offers the paths that it selected as best, to the IP routing table
 - On Cisco routers, BGP **routes learned via eBGP have an administrative distance of 20 while the iBGP-learned routes have an administrative distance of 200**
- Transmission of routing information to neighbors
 - Several rules, such as the one commonly referred to as **iBGP split-horizon rule**, govern the advertisement of BGP routes to neighbors
 - Access lists, prefix lists, and route maps may be applied to filter and manipulate the prefixes and their attributes before exchanging them with a neighbor
 - This type of filtering and manipulation can be performed either before transmitting or when receiving information from a neighbor

Neighbor Table

- Lists all neighbors that have been configured on a router
- Stores information such as:
 - Configured autonomous system (AS) number of the neighbor
 - Whether the neighbor is an internal or an external peer
 - The state of the session
 - Capabilities of the peer
 - How long the neighbor has been up/down for (uptime)
 - How many prefixes were exchanged with the neighbor.

BGP Table

- a.k.a. **BGP Routing Information Base (RIB)** stores all the locally injected routes, plus all routes that were received from all the router's peers, each one with appropriate attributes
 - Next hop
 - AS path
 - Local preference
 - Origin
 - Multi-exit discriminator (MED) or metric
 - Origin code
 - Community attributes
- For each prefix, the BGP best path selection algorithm will assess the usability of the available paths (reachability of next-hop, synchronization)
 - The best path is installed in the routing table and advertised to other BGP peers
 - The **BGP Multipath** feature allows additional paths to be installed in the IP routing table

Peering

- BGP **neighbor is explicitly configured** manually by specifying its IP address and AS number within BGP router configuration specifying IP address and ASN
 - **Configuration is dual** for peering routers
 - A BGP router attempts to establish a TCP session to the neighbor IP address on TCP port 179 using closest interface
 - Peer accepts incoming TCP sessions to port 179, as long as the source IP address matches one of its configured neighbor IP addresses
- After a TCP session has been successfully established with the neighbor, the two routers send BGP “OPEN” messages to exchange basic parameters and capabilities, such as:
 - Autonomous system number
 - Router ID
 - Hold time
 - Supported address families
 - Another capabilities (e.g. Route Refresh, Outbound Route Filtering)
- During this phase, each router will compare the neighbor's claimed AS number to the AS number its administrator has entered for the neighbor
 - **IF** these numbers do not match **THEN** the session is reset, and the relation is not established

Common Peering Problems

- No IP connectivity between the local BGP router and the configured peer's IP addresses
 - Because BGP peers are not necessarily directly connected, both routers need to have an IP path to the configured neighbor IP address in their routing table **other than a default route**
- The source IP address used by the router that initiates the session does not match the configured neighbor IP address on the receiving router
- The AS number of a BGP router (specified in its OPEN message) does not match the AS number its neighbor has configured for it (and expects from it)
- In case of eBGP peering, neighbor's IP addresses is not directly connected to any adjacent network

Monitoring BGP with show Commands ①

show ip bgp summary

- Displays essential BGP parameters, such as:
 - Router ID
 - AS number of the router
 - Statistics on the memory usage of the BGP process
 - Brief overview of the configured neighbors and the state of the relationship with each
- Used to quickly check the status of the relationship with one or more neighbors, or how long the relation has been down/up

show ip bgp neighbors

- Lists all configured neighbors and their current operational state, configured parameters, and extensive information about each neighbor.
- Output can be limited to a specific neighbor by using the **show ip bgp neighbors *ip-address*** command

Monitoring BGP with show Commands ②

```
Router# show ip bgp summary
```

```
BGP router identifier 172.16.1.1, local AS number 100
BGP table version is 199, main routing table version 199
37 network entries using 2850 bytes of memory
59 path entries using 5713 bytes of memory
18 BGP path attribute entries using 936 bytes of memory
2 multipath network entries and 4 multipath paths
10 BGP AS-PATH entries using 240 bytes of memory
7 BGP community entries using 168 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
90 BGP advertise-bit cache entries using 1784 bytes of memory
36 received paths for inbound soft reconfiguration
BGP using 34249 total bytes of memory
Dampening enabled. 4 history paths, 0 dampened paths
BGP activity 37/2849 prefixes, 60/1 paths, scan interval 15 secs
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.100.1.1	4	200	26	22	199	0	0	00:14:23	23
10.200.1.1	4	300	21	51	199	0	0	00:13:40	0

```
Device# show ip bgp neighbors 10.10.10.1
```

```
BGP neighbor is 10.10.10.1, remote AS 10, internal link
Fall over configured for session
BFD is configured. BFD peer is Up. Using BFD to detect fast fallover (single-hop) with c-bit check-control-plane-failure.
Inherits from template cbit-tps for session parameters
BGP version 4, remote router ID 10.7.7.7
BGP state = Established, up for 00:03:55
Last read 00:00:02, last write 00:00:21, hold time is 180, keepalive interval is 60 seconds
Neighbor sessions:
  1 active, is not multisession capable (disabled)
Neighbor capabilities:
  Route refresh: advertised and received(new)
  Four-octets ASN Capability: advertised and received
  Address family IPv4 Unicast: advertised and received
  Enhanced Refresh Capability: advertised and received
  Multisession Capability:
  Stateful switchover support enabled: NO for session 1
```

Monitoring BGP with show Commands ③

show ip bgp

- Displays the content of the BGP table
- To select a specific entry from the table, provide the network and mask of the selected prefix as an option
- Useful during troubleshooting to verify:
 - What paths are present
 - What their attributes are
 - Why certain paths are selected as best
- Does not reveal all of the attributes of the BGP paths

show ip bgp *prefix netmask*

- Display all of the attributes for a specific BGP prefix

Monitoring BGP with show Commands ④

```
Router# show ip bgp
```

```
BGP table version is 6, local router ID is 10.0.96.2
```

```
Status codes:  s suppressed, d damped, h history, * valid, > best,  
                i - internal,  
                r RIB-failure, S Stale, m multipath, b backup-path,  
                x best-external, f RT-Filter, a additional-path
```

```
Origin codes:  i - IGP, e - EGP, ? - incomplete
```

```
RPKI validation codes: V valid, I invalid, N Not found
```

	Network	Next Hop	Metric	LocPrf	Weight	Path
*	10.0.0.1	10.0.0.3	0		0	3 ?
*>		10.0.3.5	0		0	4 ?
Nr	10.0.0.0/8	10.0.0.3	0		0	3 ?
Nr>		10.0.3.5	0		0	4 ?
Nr>	10.0.0.0/24	10.0.0.3	0		0	3 ?
V*>	10.0.2.0/24	0.0.0.0	0		32768	i
Vr>	10.0.3.0/24	10.0.3.5	0		0	4 ?

Monitoring BGP with debug Commands

`debug ip bgp`


- Displays significant BGP related events, most notably the phases of establishing a BGP peering relationship
- Does not display the content of the BGP updates

`debug ip bgp updates`


- Displays the transmission and reception of BGP updates
- Output can be limited to a specific neighbor and specific prefixes by use of extra options
- Issuing the command `debug ip bgp ip-address updates access-list` limits the output of the command to only updates received from or sent to the neighbor specified by the `ip-address` option and only for those networks that match the access list specified by the `access-list` option
- Useful during troubleshooting because you can see which router is or is not sending what updates

BGP Troubleshooting Example ①

- Packets from IR01 incorrectly take route through IR02 toward ISP1



BGP Troubleshooting Example ②



```
IRO1# trace 192.168.224.1
```

```
Type escape sequence to abort.
```

```
Tracing the route to 192.168.224.1
```

```
 1 10.1.192.20 4 msec 0 msec 0 msec
 2 172.24.244.86 [AS 64566] 4 msec 0 msec 4 msec
 3 192.168.100.1 [AS 65486] 0 msec 4 msec 0 msec
 4 192.168.224.1 [AS 65525] 0 msec * 0 msec
```


BGP Troubleshooting Example ③

- The results from the **ping** command show that the link from IR01 to ISP1 access router is up

```
IR01# ping 192.168.224.254
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 192.168.224.254, timeout is 2 seconds:
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

BGP Troubleshooting Example ④

- In the IR01 routing table, the route to IP address 192.168.224.1 is via network 192.168.224.0/19
- This route is an internal BGP path. Its source is the router with IP address 10.1.220.4, which is IRO2.

```
IR01# show ip route 192.168.224.1
Routing entry for 192.168.224.0/19, supernet
  Known via "bgp 64568", distance 200, metric 0
  Tag 64566, type internal
  Redistributing via eigrp 1
  Last update from 172.24.244.86 00:24:22 ago
  Routing Descriptor Blocks:
  * 172.24.244.86, from 10.1.220.4, 00:24:22 ago
    Route metric is 0, traffic share count is 1
    AS Hops 2
    Route tag 64566
```

BGP Troubleshooting Example ⑤

- IR01 BGP table shows only one path to 192.168.224.0/19

```
IR01# show ip bgp 192.168.224.1
BGP routing table entry for 192.168.224.0/19, version 12
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Not advertised to any peer
  64566 65525
    172.24.244.86 (metric 30720) from 10.1.220.4 (10.1.220.4)
      Origin IGP, metric 0, localpref 100, valid, internal, best
```

BGP Troubleshooting Example ⑥

- IR01 BGP neighbor 192.168.224.244 (ISP1 access router) is in active state.

```
IR01# show ip bgp summary
```

```
BGP router identifier 10.1.220.3, local AS number 64568
```

```
BGP table version is 14, main routing table version 14
```

```
6 network entries using 702 bytes of memory
```

```
7 path entries using 364 bytes of memory
```

```
6/4 BGP path/bestpath attribute entries using 744 bytes of memory
```

```
3 BGP AS-PATH entries using 72 bytes of memory
```

```
0 BGP route-map cache entries using 0 bytes of memory
```

```
0 BGP filter-list cache entries using 0 bytes of memory
```

```
BGP using 1882 total bytes of memory
```

```
BGP activity 6/0 prefixes, 13/6 paths, scan interval 60 secs
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	
State/PfxRcd									
10.1.220.4	4	64568	82	80	14	0	0	01:12:02	6
192.168.224.244	4	65525	0	0	0	0	0	never	Active


BGP Troubleshooting Example ⑦

- The **ping** from IR01 to the configured BGP neighbor's address fails

```
IR01# ping 192.168.224.244
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.224.244, timeout is 2
seconds:
.....
Success rate is 0 percent (0/5)
```

BGP Troubleshooting Example ⑧

- Correcting BGP neighbor ISP1 address on IR01



```
IR01(config)# router bgp 64568
IR01(config-router)# no neighbor 192.168.224.244
IR01(config-router)# neighbor 192.168.224.254 remote-as 65525
```

BGP Troubleshooting Example ⑨

- After the correction was made on IR01, the neighbor state for 192.168.224.254 is established
- A path to 192.168.224.0/19 is received from this neighbor

```
IR01# show ip bgp summary | begin Neighbor
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.1.220.4	4	64568	146	146	19	0	0	02:15:17	5
192.168.224.254	4	65525	14	12	19	0	0	00:03:23	5

```
IR01# show ip bgp 192.168.224.0
```

```
BGP routing table entry for 192.168.224.0/19, version 17
```

```
Paths: (1 available, best #1, table Default-IP-Routing-Table)
```

```
  Advertised to update-groups:
```

```
    2
```

```
  65525
```

```
    192.168.224.254 from 192.168.224.254 (192.168.100.1)
```

```
      Origin IGP, metric 0, localpref 100, valid, external, best
```

BGP Troubleshooting Example ⑩

- IR02's BGP table shows two paths to 192.168.224.0/19
- Only the best path via IR01 will be used

```
IR02# show ip bgp 192.168.224.0
BGP routing table entry for 192.168.224.0/19, version 24
Paths: (2 available, best #1, table Default-IP-Routing-Table)
  Advertised to update-groups:
    1
  65525
    192.168.224.254 (metric 30720) from 10.1.220.3 (10.1.220.3)
      Origin IGP, metric 0, localpref 100, valid, internal, best
  64566 65525
    172.24.244.86 from 172.24.244.86 (172.24.240.1)
      Origin IGP, localpref 100, valid, external
```


BGP Troubleshooting Example ⑪

- IR01's IP routing table shows one path to 192.168.224.0/19
- The **traceroute** command shows that the path is through ISP1

```
IR01# show ip route 192.168.224.1
Routing entry for 192.168.224.0/19, supernet
  Known via "bgp 64568", distance 20, metric 0
  Tag 65525, type external
  Redistributing via eigrp 1
  Last update from 192.168.224.254 00:49:55 ago
  Routing Descriptor Blocks:
  * 192.168.224.254, from 192.168.224.254, 00:49:55 ago
    Route metric is 0, traffic share count is 1
    AS Hops 1
    Route tag 65525
```

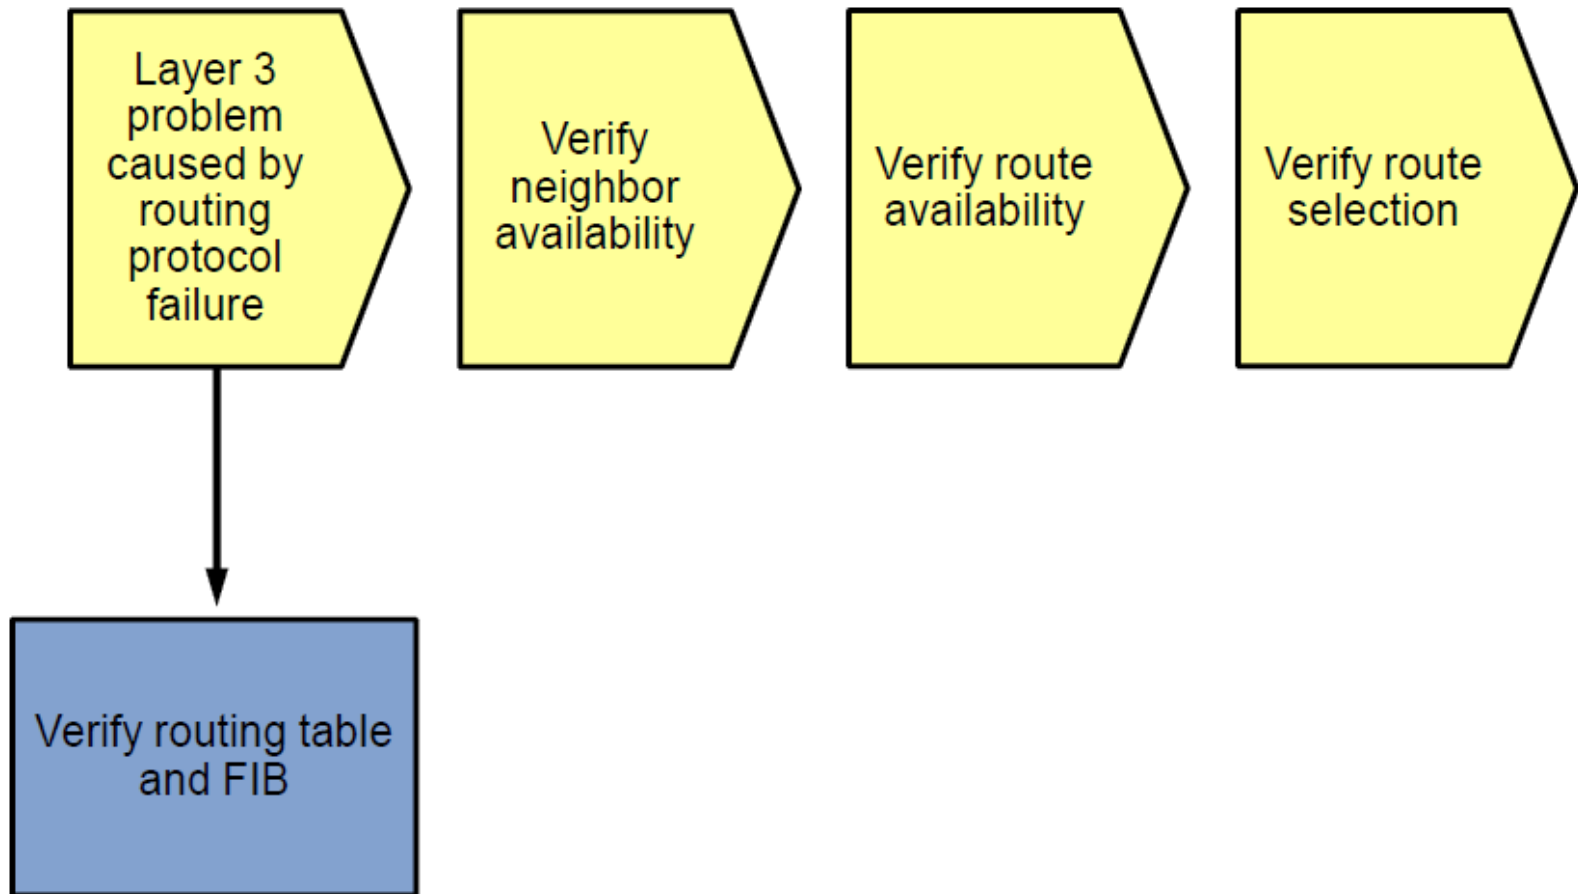
```
IR01# traceroute 192.168.224.1
Type escape sequence to abort.
Tracing the route to 192.168.224.1

 1 192.168.224.254 [AS 65525] 0 msec 0 msec 4 msec
 2 192.168.224.1 [AS 65525] 0 msec 0 msec *
```

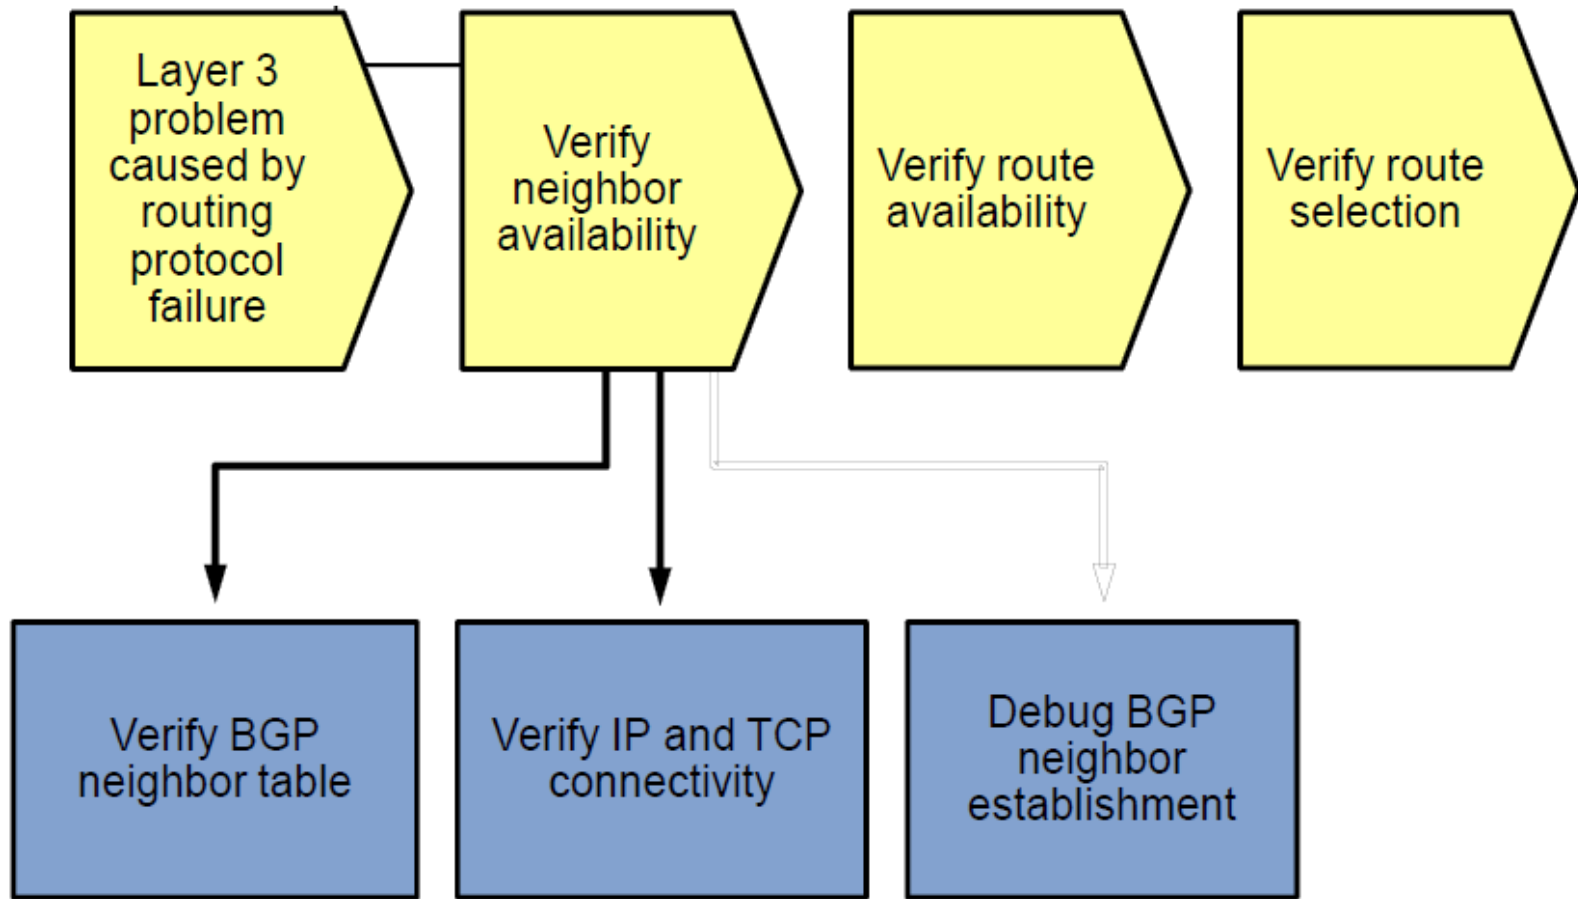
BGP Troubleshooting Flow



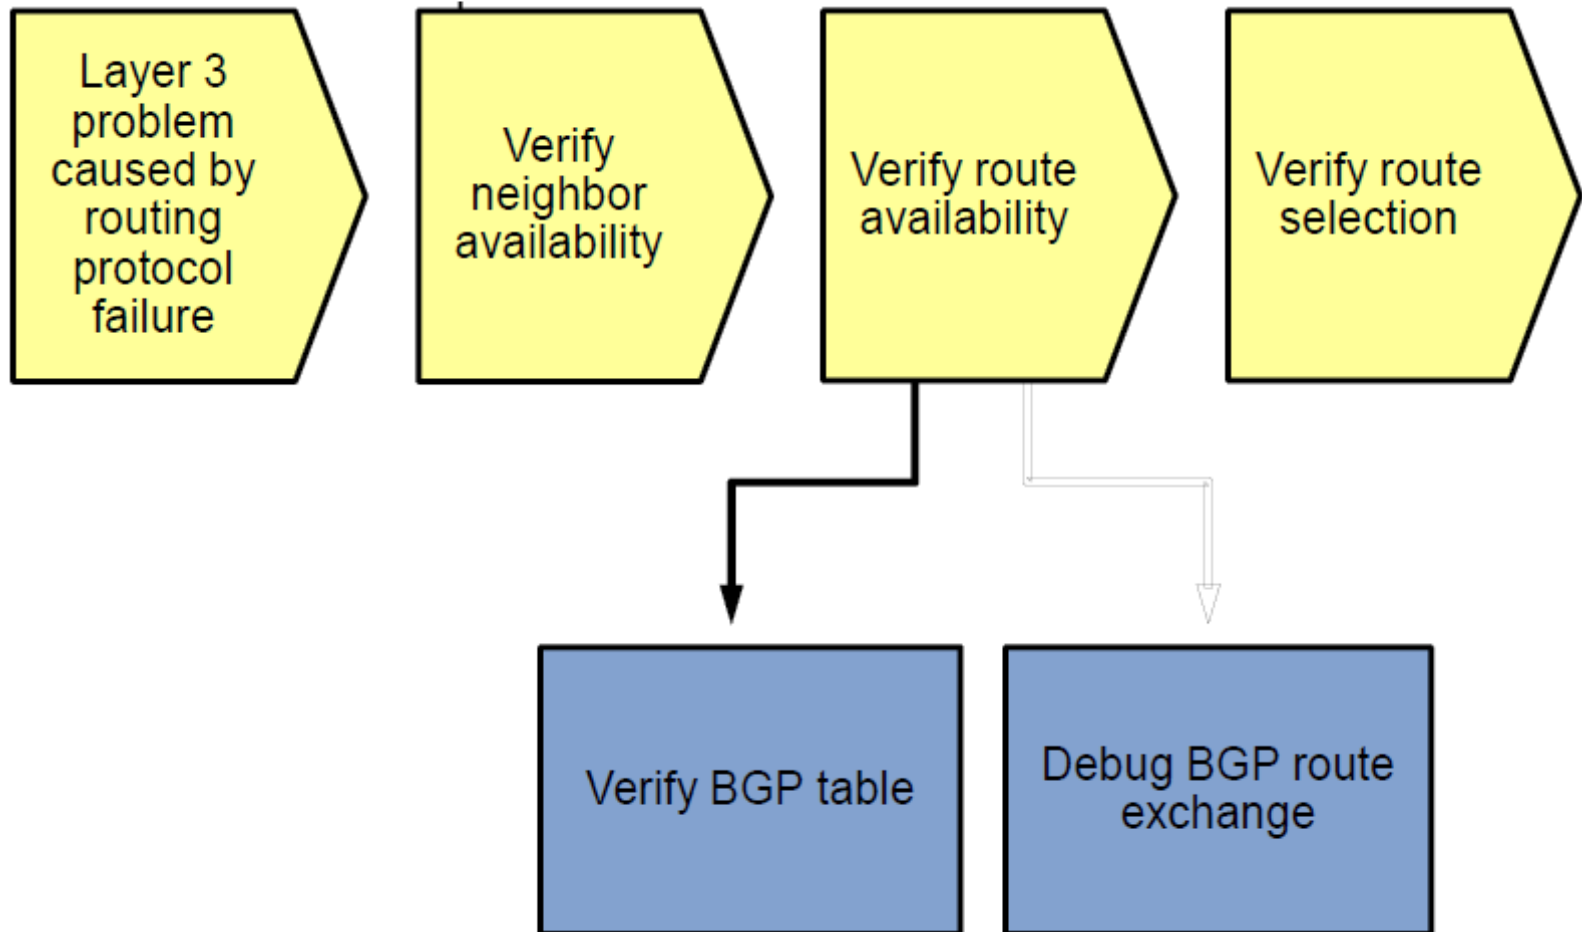
BGP Troubleshooting Flow ①



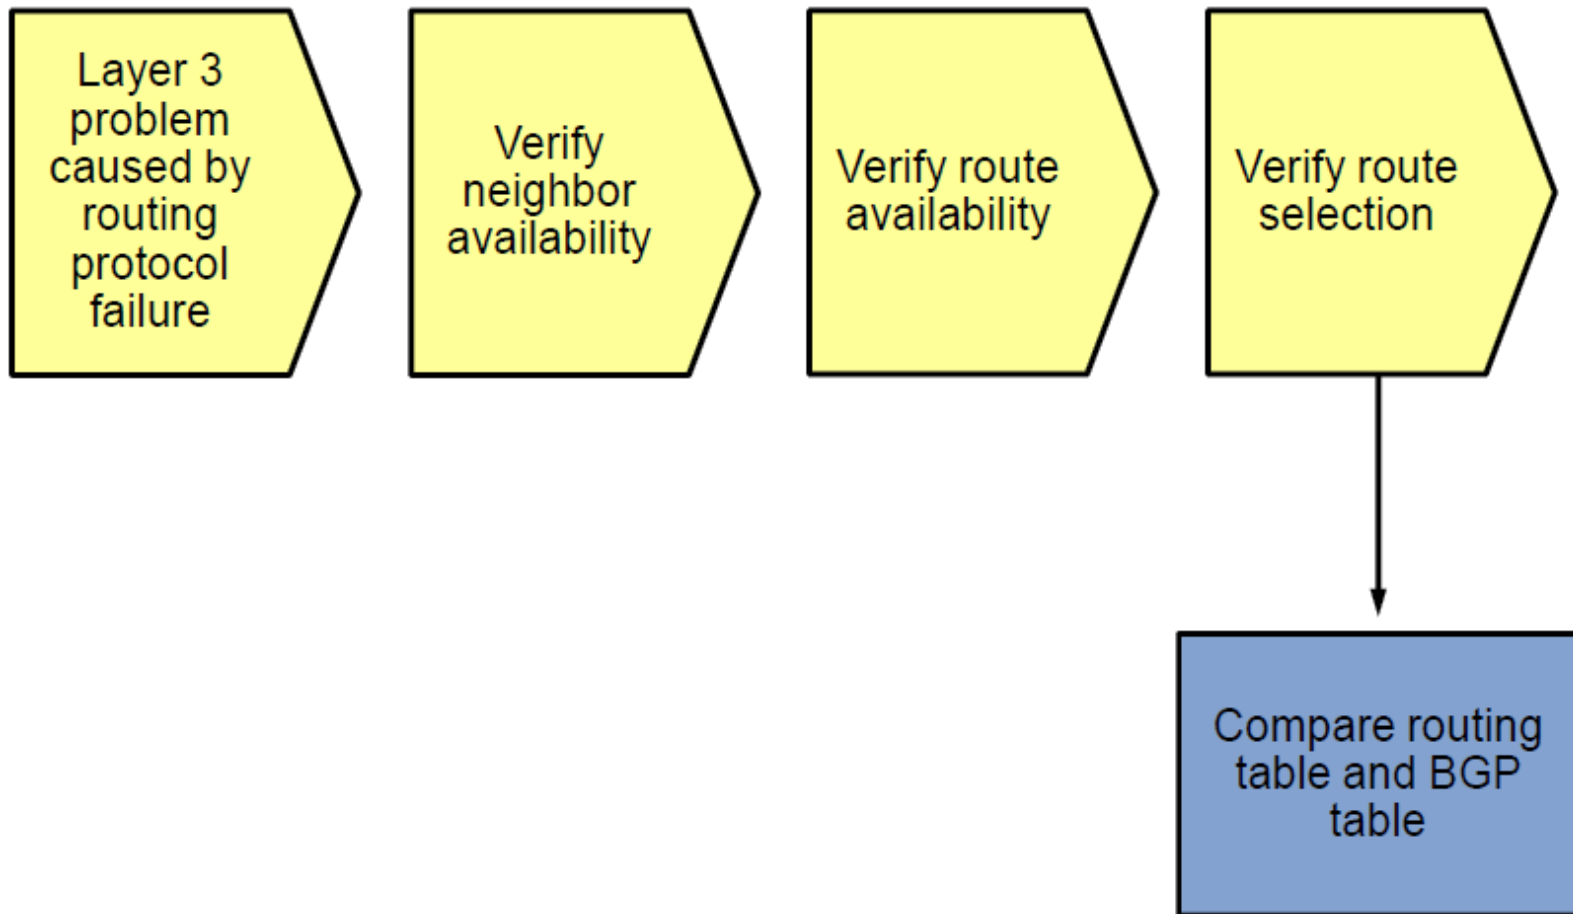
BGP Troubleshooting Flow ②



BGP Troubleshooting Flow ③



BGP Troubleshooting Flow ④



Summary



Routing Protocol Similarities

- Virtually all routing protocols include the following elements and processes:
 - Reception of routing information from neighbors
 - Routing protocol data structures
 - Route injection or redistribution
 - Route selection and installation
 - Transmission of routing information to neighbors

Other Useful Show Commands

show ip route *ip-address*: Displays best route that matches the IP address.

show ip route *network mask*: Displays routing table entry that is an exact match for the network and mask.

show ip route *network mask longer-prefixes*: Displays routing table prefixes within the network and mask parameters.

show ip cef *ip-address*: Displays information necessary to forward packets in the FIB for the prefix (not routing protocol related information).

show ip cef *network mask*: Displays the FIB table entry that matches the network and mask.

show ip cef exact-route *source destination*: Displays adjacency used to forward a packet with specified source and destination IP addresses.

show ip arp: Verifies IP address to Ethernet MAC address mappings resolved by ARP.

show frame-relay map: Lists mappings of next-hop IP addresses on multipoint (sub-) interfaces to the DLCI of the corresponding PVC. It also lists DLCIs manually associated to specific point-to-point sub-interfaces.

EIGRP Troubleshooting Commands

- **show ip eigrp interfaces**: Displays the list of interfaces that have been activated for EIGRP
- **show ip eigrp neighbors**: Lists all neighbors that have been discovered by the local router
- **show ip eigrp topology**: Displays the content of the EIGRP topology table
- **debug ip routing**:
 - Generic IP debugging command
 - Displays any changes that are made to the routing table, such as installation or removal of routes
 - Can be useful to diagnose routing protocol instabilities
- **debug eigrp packets**: Displays the transmission and reception of all EIGRP packets
- **debug ip eigrp**: Displays EIGRP routing events, such as updates, queries, and replies sent to or received from neighbors
- **debug ip eigrp neighbor as-number ip-address**: Limits output to information that is associated with the specified neighbor
- **debug ip eigrp as-number network mask**: Limits output to information that is associated with the network specified by the network and mask options.

OSPF Troubleshooting Commands

- **show ip ospf interface**: Displays the interfaces that have been activated for OSPF.
- **show ip ospf neighbor**: Lists neighbors discovered by this router and shows their current state.
- **show ip ospf database**: Displays the content of the OSPF link state database.
- **show ip ospf statistics**: Shows how often and when the SPF algorithm was last executed. This command can be helpful when diagnosing routing instability.
- **debug ip routing**: Displays any changes that are made to the routing table.
- **debug ip ospf packet**: Displays the transmission and reception of OSPF packets.
- **debug ip ospf events**: Displays OSPF events.
- **debug ip ospf adj**: Displays events that are related to the adjacency building process.
- **debug ip ospf monitor**: Monitors when the SPF algorithm is scheduled to run and displays the triggering LSA and a summary of the results after the SPF algorithm has completed.

BGP Troubleshooting Commands

- **show ip bgp summary**: Displays BGP parameters, such as the router ID and AS number of the router, statistics on BGP process memory usage, and an overview of the configured neighbors and the state of the relationship with each one.
- **show ip bgp neighbors**: Lists configured neighbors and their current operational state, configured parameters, and statistics for each neighbor. The output can be limited to a specific neighbor by using the **show ip bgp neighbors ip-address** command.
- **show ip bgp**: Displays the contents of the BGP table. To select a specific entry from the table, specify the network and mask of a prefix as an option to the command.
- **debug ip bgp**: Displays BGP related events, most notably the subsequent phases of establishing a BGP peering relationship. This command does not display the content of the BGP updates and is a relatively safe to use.
- **debug ip bgp updates**: Displays the transmission and reception of BGP updates. The output of this debug can be limited to a specific neighbor and specific prefixes by use of extra options.



Slides adapted by [Vladimír Veselý](#) and [Matěj Grégr](#)
partially from official course materials
but the most of the credit goes to CCIE#23527 Ing. Peter Palúch, Ph.D.

The last update: 2017-03-16