Co-funded by the European Union

Network T. Shooting

linkedin.com/in/nazrul13

Troubleshooting Methodology

- Cisco has developed a troubleshooting model to effectively address network issues that will arise and equip you in handling such problems. An important part of troubleshooting is to know how to divide the tasks needed to resolve the issue in a systematic process of elimination. Cisco has broken down the process into eight methodical steps:
  - Define the problem.
  - Gather detailed information.
  - Consider probable cause for the failure.
  - Devise a plan to solve the problem.
  - Implement the plan.
  - Observe the results of the implementation.
  - Repeat the process if the plan does not resolve the problem.
  - Document the changes made to solve the problem.

OSI Model Approach?

- There are several ways to address a network issue using the OSI model, depending on the situation.
  - Top-Down
  - Bottom-Up

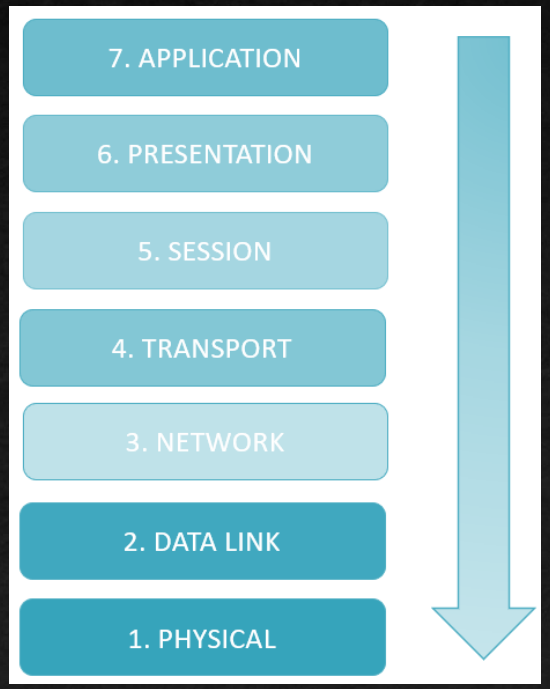

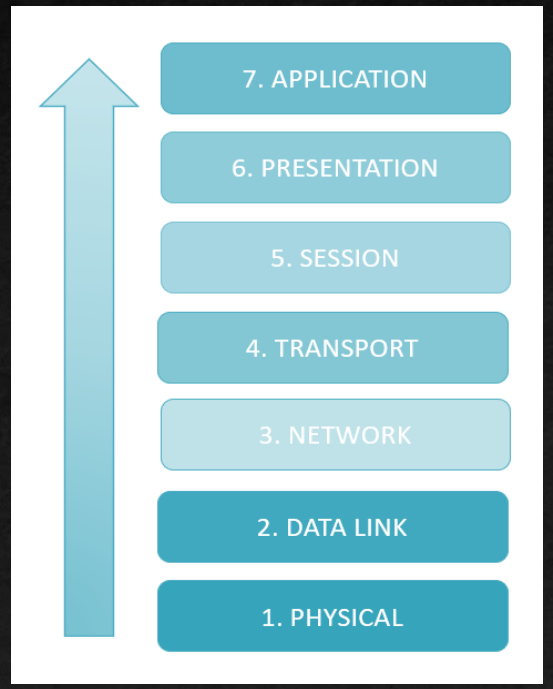

OSI Model Approach?(Cont.)

- Divide and Conquer:

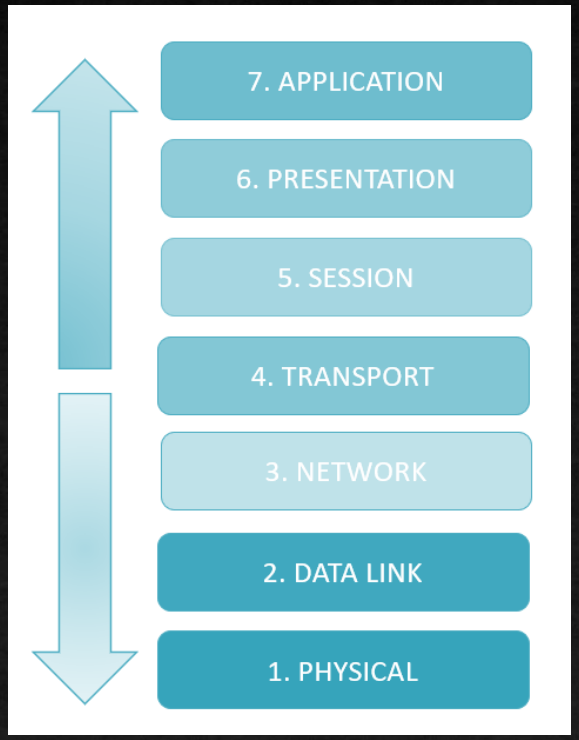

Other troubleshooting methods

- Compare Configurations

A lot of network performance issues are usually caused by human errors, and the initial way to troubleshoot problems is to check if there are configuration changes that have been made in the network. One way of knowing these changes is by implementing the AAA mechanism because such changes are being logged by an AAA server, or you can locally access the logs within the device.

- Trace the Path

One of the most used troubleshooting tools is sending a ping to your destination device. There is another ICMP-based tool that shows you where the ICMP packet stopped in the network, and that is the traceroute. Having to know where your ping stops gives you an advantage in knowing where the issue is happening so you can easily isolate the problem and further analyze the best approach to rectifying the issue.

Other troubleshooting methods

- Swap-out Components

Usually, network outages are caused by hardware failures ranging from a simple ethernet cable wear and tear to full-on equipment failure. When this happens, we have no choice but to replace the defective hardware with a new one to keep the network up and running. This approach is also used to check if there is a specific device that causes the issue in the network and monitors what happens once the swap has been made

Network Troubleshoot

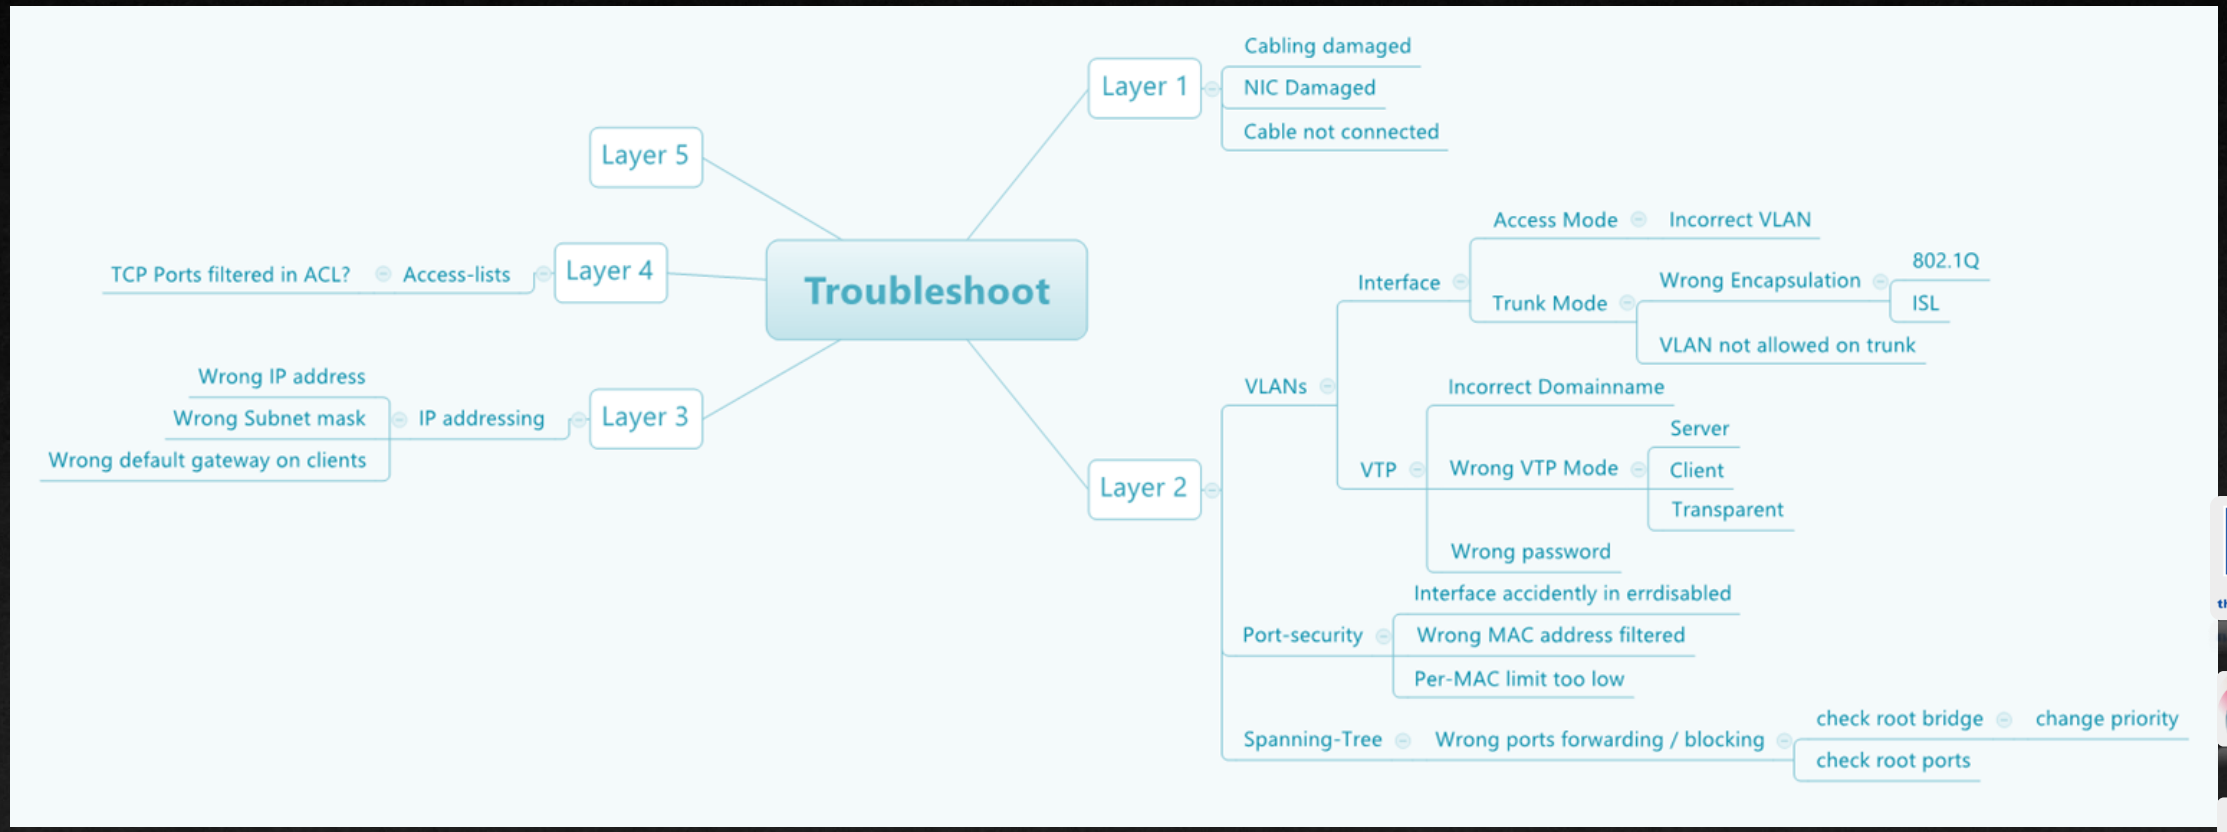

Connectivity Troubleshooting Tools

- There are various troubleshooting tools that are being used to analyze network connection outages or performance issues. Below are some of the most effective tools that we utilize in troubleshooting and can be helpful if we understand how they function.
  - arp
  - ping
  - Traceroute
  - Route
  - Telnet

ICMP (Internet Control Message Protocol)

- ICMP (Internet Control Message Protocol) is a network layer protocol that reports errors and provides information related to IP packet processing. ICMP is used by network devices to send error messages indicating, for example, that a requested service is not available or that a host isn't reachable.
  - ICMP is commonly used by network tools such as ping or traceroute.

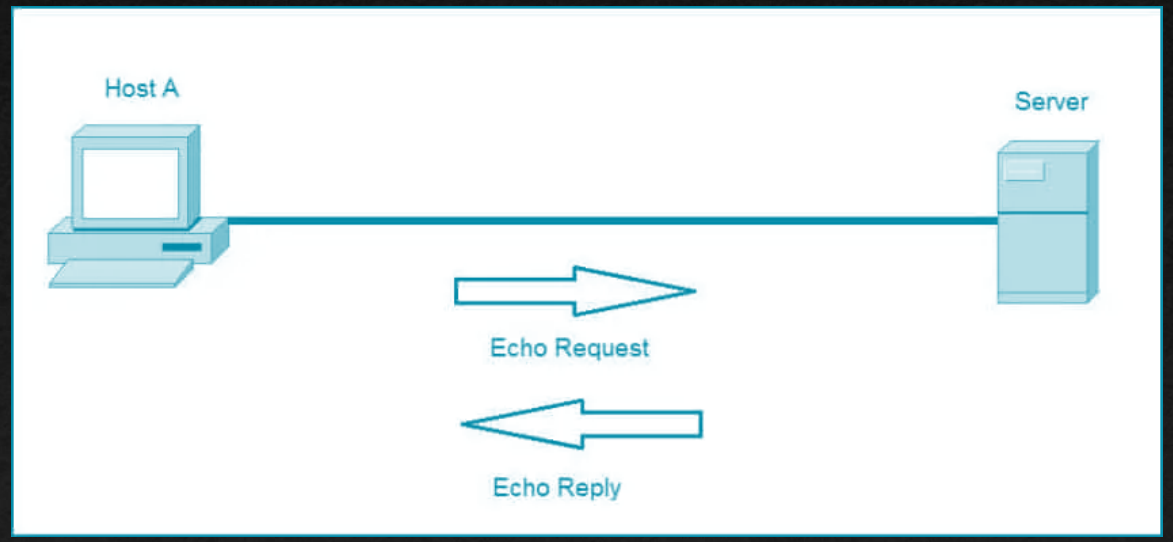

Cisco Ping Return Codes

Code	Meaning	Possible Cause
!	Each exclamation point indicates receipt of an ICMP echo reply.	The ping completed successfully.
.	Each period indicates that the network server timed out while waiting for a reply	This message can indicate many problems: • ping was blocked by an access list or firewall. • A router along the path did not have a route to the destination and did not send an ICMP destination unreachable message. • A physical connectivity problem occurred somewhere along the path.
U	An ICMP unreachable message was received.	A router along the path did not have a route to the destination address.
C	An ICMP source quench message was received.	A device along the path—possibly the destination—may be receiving to much traffic; check input queues.
&	An ICMP time exceeded message was received	A routing loop may have occurred

Ping Code Example (Destination unreachable)

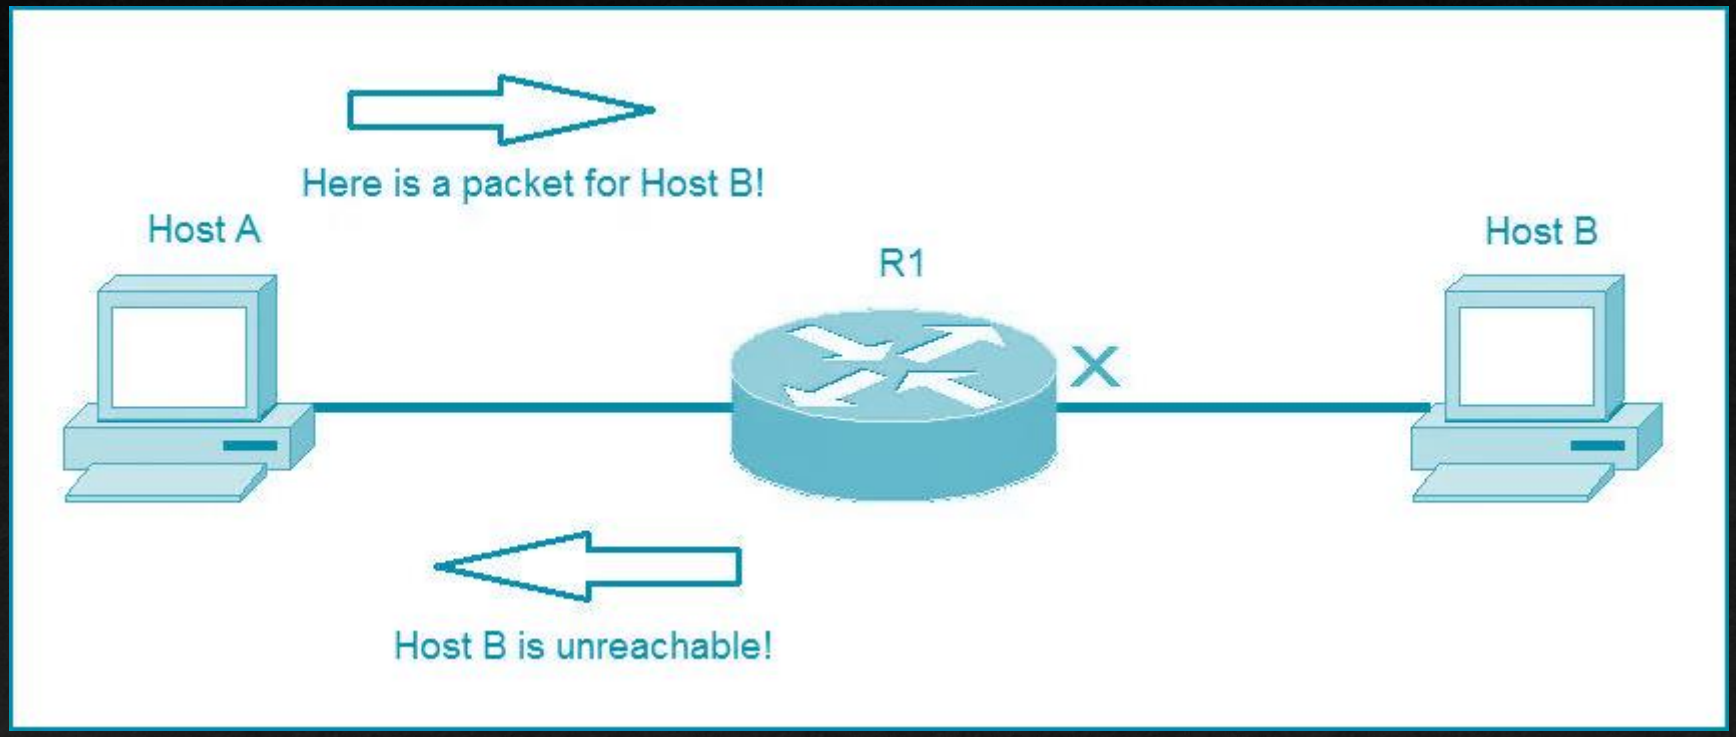

Traceroute

- Traceroute is a command-line interface based tool used to identify the path used by a packet to reach its target. This tool also uses ICMP messages, but unlike ping, it identifies every router in a path taken by the packets. Traceroute is useful when troubleshooting network problems because it can help identify where exactly the problem is. You can figure out which router in the path to an unreachable target should be examined more closely as the probable cause of the network's failure.

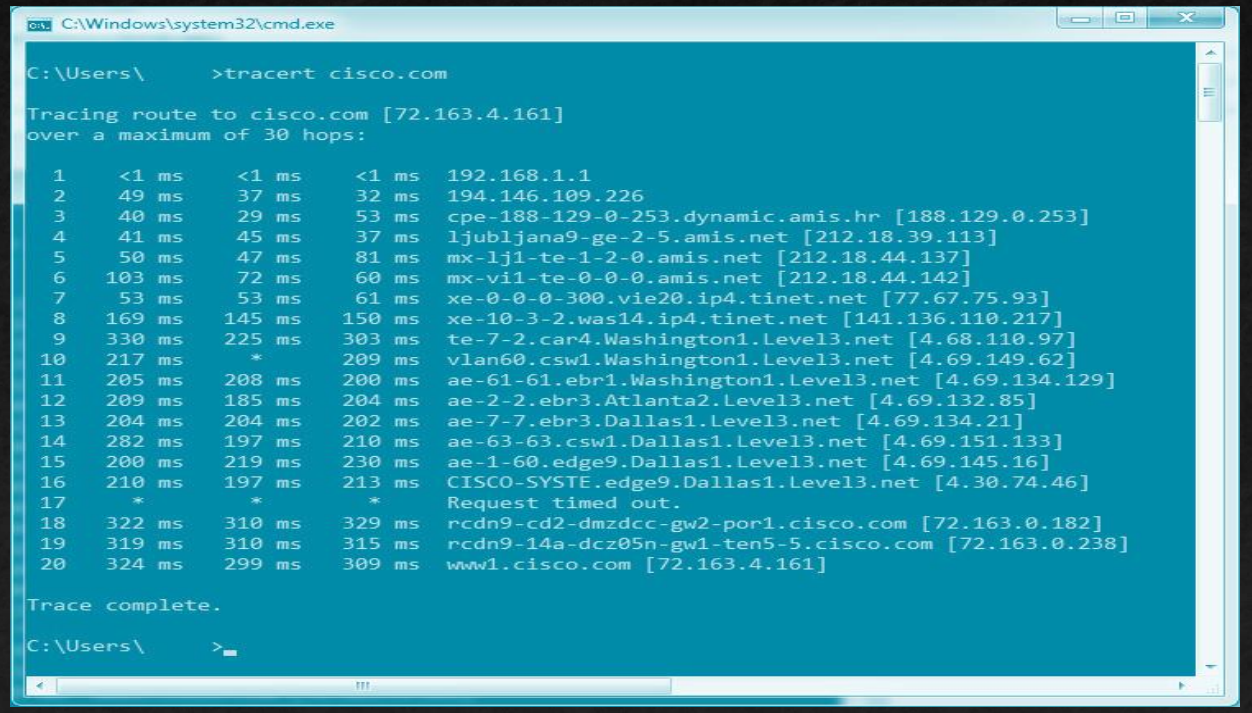

Debugging

- Debugging can be a very powerful part of troubleshooting complex issues in a network.
- One of the most common use cases for debugging is when there is a need to see things at a deeper level (such as when routing protocols are having adjacency issues).
- There is a normal flow that is taken from a troubleshooting perspective, depending on the routing protocol.

```
R1# debug ip ospf adj
OSPF adjacency debugging is on
R1#
19:20:42.559: OSPF-1 ADJ   Et0/1: Rcv DBD from 4.4.4.4 seq 0x247A opt 0x52 flag 0x7
len 32  mtu 1400 state EXCHANGE
19:20:42.559: OSPF-1 ADJ   Et0/1: Nbr 4.4.4.4 has smaller interface MTU
19:20:42.559: OSPF-1 ADJ   Et0/1: Send DBD to 4.4.4.4 seq 0x247A opt 0x52 flag 0x2
len 152
R1#un all
All possible debugging has been turned off
```

Debug IP OSPF ADJ Command

- The debug ip ospf adj command is used to reveal messages that are exchanged during the OSPF adjacency process.
- The output of the debug ip ospf adj command in snapshot clearly states that it received a Database Descriptor packet from the neighbor 4.4.4.4, and that the neighbor 4.4.4.4 has a smaller interface MTU of 1400.

```
R1# debug ip ospf adj
OSPF adjacency debugging is on
R1#
19:20:42.559: OSPF-1 ADJ   Et0/1: Rcv DBD from 4.4.4.4 seq 0x247A opt 0x52 flag 0x7
len 32  mtu 1400 state EXCHANGE
19:20:42.559: OSPF-1 ADJ   Et0/1: Nbr 4.4.4.4 has smaller interface MTU
19:20:42.559: OSPF-1 ADJ   Et0/1: Send DBD to 4.4.4.4 seq 0x247A opt 0x52 flag 0x2
len 152
R1#un all
All possible debugging has been turned off
```

```
R4# debug ip ospf adj
OSPF adjacency debugging is on
R4#
19:28:18.102: OSPF-1 ADJ   Et0/1: Send DBD to 1.1.1.1 seq 0x235C opt 0x52 flag 0x7
len 32
19:28:18.102: OSPF-1 ADJ   Et0/1: Retransmitting DBD to 1.1.1.1 [23]
19:28:18.102: OSPF-1 ADJ   Et0/1: Rcv DBD from 1.1.1.1 seq 0x235C opt 0x52 flag 0x2
len 152  mtu 1500 state EXSTART
19:28:18.102: OSPF-1 ADJ   Et0/1: Nbr 1.1.1.1 has larger interface MTU
R4#un all
All possible debugging has been turned off
```

Debug IP OSPF Hello Command

- The issue with OSPF network type mismatch, which is a very common reason for neighbor adjacency issues. Often this is simply a misconfiguration issue when setting up the network.

```
R1# debug ip ospf hello
OSPF hello debugging is on
R1#
19:47:46.976: OSPF-1 HELLO Et0/0: Send hello to 224.0.0.5 area 0 from 192.168.12.1
19:47:47.431: OSPF-1 HELLO Et0/1: Send hello to 224.0.0.5 area 0 from 192.168.14.1
19:47:48.363: OSPF-1 HELLO Et0/2: Send hello to 224.0.0.5 area 0 from 192.168.17.1
R1#
19:47:50.582: OSPF-1 HELLO Et0/0: Rcv hello from 2.2.2.2 area 0 192.168.12.2
19:47:51.759: OSPF-1 HELLO Et0/2: Rcv hello from 7.7.7.7 area 0 192.168.17.7
R1#
19:47:56.923: OSPF-1 HELLO Et0/0: Send hello to 224.0.0.5 area 0 from 192.168.12.1
19:47:57.235: OSPF-1 HELLO Et0/1: Send hello to 224.0.0.5 area 0 from 192.168.14.1
19:47:58.159: OSPF-1 HELLO Et0/2: Send hello to 224.0.0.5 area 0 from 192.168.17.1
R1#
19:47:59.776: OSPF-1 HELLO Et0/0: Rcv hello from 2.2.2.2 area 0 192.168.12.2
19:48:01.622: OSPF-1 HELLO Et0/2: Rcv hello from 7.7.7.7 area 0 192.168.17.7
R1#un all
All possible debugging has been turned off
```

```
R4# debug ip ospf hello
OSPF hello debugging is on
R4#
19:45:45.127: OSPF-1 HELLO Et0/1: Rcv hello from 1.1.1.1 area 0 192.168.14.1
19:45:45.127: OSPF-1 HELLO Et0/1: Mismatched hello parameters from 192.168.14.1
19:45:45.127: OSPF-1 HELLO Et0/1: Dead R 40 C 120, Hello R 10 C 30
19:45:45.259: OSPF-1 HELLO Et0/3: Rcv hello from 7.7.7.7 area 0 192.168.47.7
R4#
19:45:48.298: OSPF-1 HELLO Et0/0: Send hello to 224.0.0.5 area 0 from 192.168.34.4
19:45:48.602: OSPF-1 HELLO Et0/0: Rcv hello from 3.3.3.3 area 0 192.168.34.3
R4#un all
All possible debugging has been turned off
```

Debugging
Hello and Dead Intervals

- Different network types have different hello intervals and dead intervals.

Network Type	Hello Interval (in seconds)	Dead Interval (in seconds)
Broadcast	10	40
Non-broadcast	30	120
Point-to-point	10	40
Point-to-Multipoint	30	120

Debugging
Output of the Show IP OSPF Interface

- The issue could be simply mismatched network types or mismatched hello or dead intervals.
- The show ip ospf interface command shows what the configured network types and hello and dead intervals are.

```
R4# show ip ospf interface ethernet0/1
Ethernet0/1 is up, line protocol is up
  Internet Address 192.168.14.4/24, Area 0, Attached via Network Statement
  Process ID 1, Router ID 4.4.4.4, Network Type POINT_TO_MULTIPOINT, Cost: 10
  Topology-MTID    Cost    Disabled    Shutdown      Topology Name
        0           10        no          no            Base
  Transmit Delay is 1 sec, State POINT_TO_MULTIPOINT
  Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5
    oob-resync timeout 120
    Hello due in 00:00:05
  Supports Link-local Signaling (LLS)
  Cisco NSF helper support enabled
  IETF NSF helper support enabled
  Index 2/2, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 2
  Last flood scan time is 0 msec, maximum is 1 msec
  Neighbor Count is 0, Adjacent neighbor count is 0
  Suppress hello for 0 neighbor(s)
```

Debugging
Conditional Debugging

- Conditional debugging can be used to limit the scope of the messages that are being returned to the console or syslog server. A great example of this is the debug ip packet command. Issuing this command on a router that is in production could send back a tremendous number of messages.
- One way to alleviate this issue is to attach an access list to the debug command to limit the scope of messages to the source or destination specified within the access list. This can be done using standard or extended access lists.
- The options for the debug ip packet command are as follows:
  - <1-199>: Standard access list
  - <1300-2699>: Access list with expanded range
  - detail: More debugging detail

Troubleshooting OSPF Neighbor Relationships

- The following are some of the reasons an OSP neighbor relationship might not form:
  - Interface is down - The interface must be up/up.
  - Interface not running the OSPF process - If the interface is not enabled for OSPF, it does not send hello packets or form adjacencies.
  - Mismatched timers - Hello and dead timers must match between neighbors.
  - Mismatched area numbers - The two ends of a link must be in the same OSPF area.
  - Mismatched area type - In addition to a normal OSPF area type, an area type could be a stub area or a not-so-stubby area (NSSA). The routers must agree on the type of area they are in.
  - Different subnets - Neighbors must be in the same subnet.
  - Passive interface - The passive interface feature suppresses the sending and receiving of hello packets while still allowing the interface's network to be advertised.
  - Mismatched authentication information - Both OSPF interfaces must be configured for matching authentication
  - ACLs - An ACL may be denying packets to the OSPF multicast address 224.0.0.5.
  - MTU mismatch - The maximum transmission unit of neighboring interfaces must match.
  - Duplicate router IDs - Router IDs must be unique.
  - Mismatched network types - neighbors configured with a different OSPF network type might not form an adjacency.

Troubleshooting EIGRP Neighbor Relationships

- The following are some of the reasons an EIGRP neighbor relationship might not form:
  - Interface is down - The interface must be up/up.
  - Mismatched autonomous system numbers - Both routers need to be using the same autonomous system number.
  - Incorrect network statement - The network statement must identify the IP address of the interface you want to include in the EIGRP process.
  - Mismatched K values - Both routers must be using exactly the same K values.
  - Passive interface - The passive interface feature suppresses the sending and receiving of hello packets while still allowing the interface's network to be advertised.
  - Different subnets - The exchange of hello packets must be done on the same subnet; if it isn't, the hello packets are ignored.
  - Authentication - If authentication is being used, the key ID and key string must match, and the key must be valid (if valid times have been configured).
  - ACLs - An access control list (ACL) may be denying packets to the EIGRP multicast address 224.0.0.10.
  - Timers - Timers do not have to match; however, if they are not configured correctly, neighbor adjacencies could flap.

Troubleshooting BGP Neighbor Relationships

- The following are some of the reasons an BGP neighbor relationship might not form:
  - Interface is down - The interface must be up/up.
  - Layer 3 connectivity is broken - You need to be able to reach the IP address you are trying to form the adjacency with.
  - Path to the neighbor is through the default route - You must be able to reach the neighbor using a route other than the default route.
  - Neighbor does not have a route to the local router - The two routers forming a BGP peering must have routes to each other.
  - Incorrect neighbor statement - The IP address and ASN in the neighbor ip_address remote-as as_number statement must be accurate.
  - ACLs - An access control list (ACL) or a firewall may be blocking TCP (Transmission Control Protocol) port 179.
  - BGP packets sourced from the wrong IP address - The source IP (Internet Protocol) address of an inbound BGP packet must match the local neighbor statement.
  - The TTL (time-to-live) of the BGP packet expires - The peer may be further away than is permitted.
  - Mismatched authentication - The two routers must agree on the authentication parameters.
  - Misconfigured peer group - Peer groups simplify repetitive BGP configurations; however, if not carefully implemented, they can prevent neighbor relationships from forming or routes from being learned.
  - Timers - Timers do not have to match; however, if the minimum holddown from neighbor option is set, it could prevent a neighbor adjacency.

Debugging Simple Network Management Protocol (SNMP)

- The typical tool for reactive alerting from network devices is Simple Network Management Protocol (SNMP).
- SNMP sends unsolicited traps to an SNMP collector or network management system (NMS). These traps are in response to something that happened in the network.
- For example, traps may be generated for link status events, improper user authentication, and power supply failures. These events are defined in the SNMP Management Information Base (MIB).
- The MIB can be thought of as a repository of device parameters that can be used to trigger alerts.

Debugging Syslog

- Devices can generate useful information to the console, to the logging buffer, and to off-box syslog collectors. In fact, all three can be sent the same or different message types.
- It is critical to note that prior to configuring any device to send log information, the date and time of the clock must be properly configured for accurate time.

Level Keyword	Level	Description	Syslog Definition
emergencies	0	System unstable	LOG_EMERG
alerts	1	Immediate action needed	LOG_ALERT
critical	2	Critical conditions	LOG_CRIT
errors	3	Error conditions	LOG_ERR
warnings	4	Warning conditions	LOG_WARNING
notifications	5	Normal but significant conditions	LOG_NOTICE
informational	6	Informational messages only	LOG_INFO
debugging	7	Debugging message	LOG_DEBUG

QnA