

Högskolan i Halmstad
Sektionen för Informationsvetenskap, Data- Och Elektroteknik (IDÉ)
Olga Torstensson

**Written Exam in
Advanced Applied Routing
March 19, 2009**

Allowed aid in addition to the attached formulae:
Writing material.

Welcome to the exam!

READ THIS FIRST:

Motivate all answers. Insufficient motivation can give reduced points even if the answer is correct. If required, you are allowed to make own (reasonable) assumptions. You are allowed to answer in either ENGLISH or SWEDISH but do not mix languages in the same answer.

GOOD LUCK!

Number of exercises: 10

Maximal number of points: 60

The grade limits 30p to pass the Exam (Grade 3), 42p for Grade 4 and 54p for Grade 5.

Assignment 1: Select one of two (12 p)

Choose **one** of the following assignments. Appropriate length of an answer/description is 1-2 pages including figures. Write clear and concise. It's more important that what you write is coherent, logical and correct than everything in the subject being included. In other words, it's more important to show that you have an overall understanding than to just mention a lot of less important details. Please use examples when appropriate.

A. BGP Database and BGP Message Type

A router running BGP keeps several tables to store BGP information that it receives from and sends to other routers. These tables include a neighbor table, a BGP table (also called a forwarding database or topology database), and an IP routing table.

- *neighbor table (list of BGP neighbors)*
- *BGP table (list of networks learned from neighbors, can be multiple path to destination path, contains BGP attributes for each path)*
- *IP routing table (list of best paths to destination networks)*

For BGP to establish an adjacency, you must configure it explicitly for each neighbor. BGP uses TCP as its transport protocol (port 179). It forms a TCP connection with each of the configured neighbors and keeps track of the state of these relationships by periodically sending a BGP TCP keepalive message.

Routers that run a BGP routing process are often referred to as BGP speakers. Two BGP speakers that form a TCP connection between one another for the purpose of exchanging routing information are referred to as neighbors or peers. After establishing an adjacency, the neighbors exchange the BGP routes that are in their IP routing table. Each router collects these routes from each neighbor that successfully establishes an adjacency and then places them in its BGP forwarding database. All routes that have been learned from each neighbor are placed into the BGP forwarding database. The best routes for each network are selected from the BGP forwarding database using the BGP route selection process and then offered to the IP routing table. Each router compares the offered BGP routes to any other possible paths to those networks, and the best route, based on administrative distance, is installed in the IP routing table.

The four BGP message types are open, keepalive, update, and notification.

After a TCP connection is established, the first message sent by each side is an open message. If the open message is acceptable, the side that receives the message sends a keepalive message confirming the open message. After the receiving side confirms the open message and establishes the BGP connection, the BGP peers can exchange any update, keepalive, and notification messages. BGP peers initially exchange their full BGP routing tables. Incremental updates are sent only after topology changes in the network. BGP peers send keepalive messages to ensure that the connection between the BGP peers still exists. They send notification packets in response to errors or special conditions.

Here are more details about the different types of BGP messages:

- **Open message:** *An open message includes the following information:*
 - **Version number:** *The highest common version that both routers support is used. All BGP implementations today use BGP4.*
 - **AS number:** *The AS number of the local router. The peer router verifies this information. If it is not the AS number that is expected, the BGP session is torn down.*

- **Hold time:** Maximum number of seconds that can elapse between the successive keepalive and update messages from the sender. On receipt of an open message, the router calculates the value of the hold timer by using whichever is smaller: its configured hold time or the hold time that was received in the open message.
- **BGP router ID:** 32-bit field indicating the BGP ID of the sender. The BGP ID is an IP address that is assigned to that router, and it is determined at startup. The BGP router ID is chosen in the same way that the OSPF router ID is chosen: it is the highest active IP address on the router, unless a loopback interface with an IP address exists. In this case, the router ID is the highest loopback IP address. The router ID can also be statically configured.
- **Optional parameters:** These parameters are Type, Length, and Value (TLV)-encoded. An example of an optional parameter is session authentication.
- **Keepalive message:** BGP keepalive messages are exchanged between BGP peers often enough to keep the hold timer from expiring. If the negotiated hold-time interval is 0, periodic keepalive messages are not sent. A keepalive message consists of only a message header.
- **Update message:** A BGP update message has information on one path only; multiple paths require multiple update messages. All the attributes in the update message refer to that path, and the networks are those that can be reached through it. An update message can include the following fields:
 - **Withdrawn routes:** This list displays IP address prefixes for routes that are withdrawn from service, if any.
 - **Path attributes:** These attributes include the AS path, origin, local preference, and so on (as described later in this module). Each path attribute includes the attribute TLV. The attribute type consists of the attribute flags, followed by the attribute type code.
 - **Network-layer reachability information:** This field contains a list of IP address prefixes that are reachable by this path.
- **Notification message:** A BGP notification message is sent when an error condition is detected. The BGP connection is closed immediately after this is sent. Notification messages include an error code, an error subcode, and data related to the error. Figure displays the field for error codes that can be used to troubleshoot BGP connections.

B. BGP Attributes. Selecting a BGP Path.

BGP routers send BGP update messages about destination networks to other BGP routers. The update messages contain one or more routes and a set of BGP metrics, which are called path attributes, attached to the routes. An attribute is either well-known or optional, mandatory or discretionary, and transitive or nontransitive. An attribute may also be partial.

Not all combinations of these characteristics are valid. Path attributes fall into the following four categories:

- Well-known mandatory
- Well-known discretionary
- Optional transitive
- Optional nontransitive

Only optional transitive attributes can be marked as partial. All BGP routers must recognize a well-known attribute and propagate it to the other BGP neighbors.

Well-known attributes are either mandatory or discretionary. A well-known mandatory attribute must be present in all BGP updates. A well-known discretionary attribute does not have to be present in all BGP updates.

Attributes that are not well-known are called optional. BGP routers do not have to support an optional attribute. Optional attributes are either transitive or nontransitive.

The following statements apply to optional attributes:

- BGP routers that implement the optional attribute may propagate it to the other BGP neighbors, based on its meaning.
- BGP routers that do not implement an optional transitive attribute should pass it to other BGP routers untouched and mark the attribute as partial.
- BGP routers that do not implement an optional nontransitive attribute must delete the attribute and must not pass it to other BGP routers.

Well-known mandatory attributes

- Autonomous system path
- Next hop
- Origin

Well-known discretionary attributes

- Local preference
- Atomic aggregate

Optional transitive attribute

- Aggregator

Optional nontransitive attribute

- Multi-exit discriminator (MED)

Multiple paths may exist to reach a given network. As paths for the network are evaluated, those determined not to be the best path are eliminated from the selection criteria but are kept in the BGP forwarding table (which can be displayed using the `show ip bgp` command) in the event that the best path becomes inaccessible. BGP is not designed to perform load balancing. Paths are chosen because of policy, not based on bandwidth. The BGP selection process eliminates any multiple paths until a single best path is left. The best path is submitted to the routing table manager process and is evaluated against any other routing protocols that can also reach that network. The route from the source with the lowest administrative distance is installed in the routing table.

After BGP receives updates about different destinations from different autonomous systems, it chooses the best path to reach a specific destination. The decision process is based on the BGP attributes. BGP considers only synchronized routes with no autonomous system loops and a valid next hop.

The following process summarizes how BGP chooses the best route on a Cisco router:

1. Prefer the route with the highest weight. (The weight attribute is proprietary to Cisco and is local to the router only.)
2. If multiple routes have the same weight, prefer the route with the highest local preference value. (The local preference is used within an autonomous system.)
3. If multiple routes have the same local preference, prefer the route that the local router originated. A locally originated route has a next hop of 0.0.0.0 in the BGP table.

4. *If none of the routes were locally originated, prefer the route with the shortest autonomous system path.*
5. *If the autonomous system path length is the same, prefer the lowest origin code (IGP < EGP < incomplete).*
6. *If all origin codes are the same, prefer the path with the lowest MED. (The MED is exchanged between autonomous systems.) The MED comparison is made only if the neighboring autonomous system is the same for all routes considered, unless the*

bgp always-compare-med command is enabled.

7. *If the routes have the same MED, prefer external paths to internal paths.*
8. *If synchronization is disabled and only internal paths remain, prefer the path through the closest IGP neighbor, which means that the router prefers the shortest internal path within the autonomous system to reach the destination (the shortest path to the BGP next hop).*
9. *For EBGP paths, select the oldest route to minimize the effect of routes going up and down (flapping).*
10. *Prefer the route with the lowest neighbor BGP router ID value.*
11. *If the BGP router IDs are the same, prefer the router with the lowest neighbor IP address.*

Only the best path is entered in the routing table and propagated to the BGP neighbors of the router.

Assignment 2: IPv6 (10p)

Compare IPv6 with IPv4 and describe transition from IPv4 to IPv6. Explain IPv6 Address Types.

IPv6 was developed to overcome the limitations of the current standard, IP version 4 (IPv4). IPv4 allows end systems to communicate and forms the foundation of the Internet as we know it today. However, one of the major shortcomings of IPv4 is its limited amount of address space. The explosion of new IP-enabled devices and the growth of undeveloped regions have fueled the need for more addresses.

IP version 6 (IPv6) combines expanded addressing with a more efficient and feature-rich header to meet the demands for scalable networks in the future. IPv6 satisfies the increasingly complex requirements of hierarchical addressing that IP version 4 (IPv4) does not provide. One key benefit is that IPv6 can recreate end-to-end communications without the need for Network Address Translation (NAT)—a requirement for a new generation of shared-experience and real-time applications.

IPv6's generous 128-bit address space, it can generate a virtually unlimited stock of addresses—enough to allocate to everyone on the planet.

IPv6 is a powerful enhancement to IPv4, and several IPv6 features offer functional improvements:

- **Larger address space:** *Offers improved global reachability and flexibility; the aggregation of prefixes that are announced in routing tables; multihoming to several Internet service providers (ISPs); autoconfiguration that can include link-layer addresses in the address space; plug-and-play options; public-to private readdressing end to end without address translation; and simplified mechanisms for address renumbering and modification.*

- **Simpler header:** Provides better routing efficiency; no broadcasts and thus no potential threat of broadcast storms; no requirement for processing checksums; simpler and more efficient extension header mechanisms; and flow labels for per-flow processing with no need to open the transport inner packet to identify the various traffic flows.
- **Mobility and security:** Ensures compliance with mobile IP and IPsec standards functionality; mobility is built in, so any IPv6 node can use it when necessary; and enables people to move around in networks with mobile network devices—with many having wireless connectivity.

Transition richness: You can incorporate existing IPv4 capabilities in IPv6 in the following ways:

- Configure a dual stack with both IPv4 and IPv6 on the interface of a network device.
- Use the technique IPv6 over IPv4 (also called 6to4 tunneling), which uses an IPv4 tunnel to carry IPv6 traffic. This method (RFC 3056) replaces IPv4-compatible tunneling (RFC 2893). Cisco IOS Software Release 12.3(2)T (and later) also allows protocol translation (NAT-PT) between IPv6 and IPv4. This translation allows direct communication between hosts speaking different protocols.

The IPv6 addressing structure is defined in multiple RFCs, including RFC 3513 and the new RFC 4291 (obsoletes RFC 3513). Each RFC defines three types of IPv6 addresses:

- Unicast address
- Multicast address
- Anycast address

Unicast Address

A unicast address identifies a single device. A packet sent to a unicast address is delivered to the interface identified by that address.

There are two types of unicast addresses:

- **Link-local unicast address:** Scope is configured to single link. The address is unique only on this link, and it is not routable off the link.
- **Global unicast address:** Globally unique, so it can be routed globally with no modification. A global address has an unlimited scope on the worldwide Internet. Packets with global source and destination addresses are routed to their target destination by the routers on the Internet.

All interfaces are required to have at least one link-local unicast address. However, a fundamental feature of IPv6 is that a single interface may also have multiple IPv6 addresses of any type (unicast, anycast, and multicast).

Multicast Address

IPv6 does not have broadcast addresses. Broadcasting in IPv4 results in several problems: It generates a number of interrupts in every computer on the network and, in some cases, triggers malfunctions that can completely halt an entire network. This disastrous network event is called a “broadcast storm.”

Broadcasts are replaced by multicast addresses. Multicast enables efficient network operation by using functionally specific multicast groups to send requests to a limited number of computers on the network. A packet sent to a multicast address is delivered to all interfaces identified by that address. The range of multicast addresses in IPv6 is larger than in IPv4. For the foreseeable future, allocation of multicast groups is not being limited.

Anycast Address

IPv6 also defines a new type of address called anycast. An anycast address identifies a list of devices or nodes; therefore, an anycast address identifies multiple interfaces.

A packet sent to an anycast address is delivered to the closest interface, as defined by the routing protocols in use.

Anycast addresses are syntactically indistinguishable from global unicast addresses, because anycast addresses are allocated from the global unicast address space.

Special addresses

There are a number of addresses with special meaning in IPv6

Assignment 3: OSPF (10p)

Explain hierarchical routing, OSPF Areas and Type of OSPF Routers. What is the difference between E1 and E2 OSPF route?

To reduce the SPF calculations, link-state routing protocols can partition networks into sub-domains called areas. An area is a logical collection of OSPF networks, routers, and links that have the same area identification. A router within an area maintains a topological database for the area to which it belongs. Therefore the number of routers and LSAs that flood the area are also smaller. This keeps the LSDB for an area smaller and, as a result, the SPF calculations are less resource intensive. In addition, the router does not have detailed information about network topology outside of its area, thereby reducing the size of its database. Areas limit the scope of route information distribution and reduce the number of routes to propagate. The LSDBs of routers within the same area must be synchronized and be exactly the same. However, route summarization and filtering is possible between different areas.

Link-state routing protocols use a two-layer area hierarchy:

- **Transit area:** Interconnects OSPF area types within a single domain. Generally, end users are not found within a transit area. OSPF area 0, also known as the backbone area, is a transit area.
- **Regular area:** Connects users and resources. Regular areas are usually set up along functional or geographical groupings. By default, a regular area does not allow traffic from another area to use its links to reach other areas. All traffic from other areas must cross a transit area. A regular area, or nonbackbone area, can have a number of subtypes, including a standard area, stub area, totally stubby area, and not-so-stubby area (NSSA).

The four different types of OSPF routers are:

- **Internal routers:** Routers that have all their interfaces in the same area and have identical LSDBs.
- **Backbone routers:** Routers that sit on the perimeter of the backbone area and have at least one interface connected to area 0. Backbone routers maintain OSPF routing

information using the same procedures and algorithms as internal routers.

- **Area border routers:** Routers that have interfaces attached to multiple areas, maintain separate LSDBs for each area to which they connect, and route traffic destined to or arriving from other areas. Area border routers (ABRs) are exit points for the area, which means that routing information destined for another area can get there only via the ABR of the local area.

ABRs can be configured to summarize the routing information from the LSDBs of their attached areas. ABRs distribute the routing information into the backbone. The backbone routers then forward the information to the other ABRs. In a multiarea network, an area can have one or more ABRs.

- **Autonomous System Boundary Routers:** Routers that have at least one interface attached to an external internetwork (another autonomous system), such as a non-OSPF network. Autonomous system boundary routers (ASBRs) can import non-OSPF network information to the OSPF network and vice versa; this process is called route redistribution.

E1 type 1 External Routing, Networks outside of the autonomous system of the router

E2 type2 External Routing, Advertised by way of external LSAs

Assignment 4: Routing tables (7p) –

A network has three routers. Three routing tables are listed below. Based on the routing tables, draw the network connectivity.

R1# show ip route

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP

D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area

* - candidate default, U - per-user static route, o - ODR

P - periodic downloaded static route

Gateway of last resort is not set

```
D EX 192.168.40.0/24 [165/40537600] via 172.16.12.2, 00:11:55, Serial0/0
  172.16.0.0/24 is subnetted, 6 subnets
D   172.16.23.0 [95/41024000] via 172.16.12.2, 00:11:55, Serial0/0
C   172.16.12.0 is directly connected, Serial0/0
C   172.16.1.0 is directly connected, Loopback0
D   172.16.2.0 [95/40640000] via 172.16.12.2, 00:11:55, Serial0/0
D EX 172.16.3.0 [165/40537600] via 172.16.12.2, 00:11:55, Serial0/0
D   172.16.100.0 [95/40640000] via 172.16.12.2, 00:11:55, Serial0/0
D EX 192.168.20.0/24 [165/40537600] via 172.16.12.2, 00:11:55, Serial0/0
C   192.168.51.0/24 is directly connected, Loopback51
C   192.168.50.0/24 is directly connected, Loopback50
D EX 192.168.35.0/24 [165/40537600] via 172.16.12.2, 00:11:55, Serial0/0
C   192.168.49.0/24 is directly connected, Loopback49
C   192.168.70.0/24 is directly connected, Loopback70
C   192.168.48.0/24 is directly connected, Loopback48
D EX 192.168.8.0/22 [165/40537600] via 172.16.12.2, 00:11:57, Serial0/0
D   192.168.48.0/23 is a summary, 02:48:33, Null0
D EX 192.168.48.0/22 [165/40537600] via 172.16.12.2, 00:11:47, Serial0/0
```

R2# show ip route

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP

D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area

* - candidate default, U - per-user static route, o - ODR

P - periodic downloaded static route

Gateway of last resort is not set

```
O 192.168.30.0/24 [110/782] via 172.16.23.3, 00:54:45, Serial0/1
O 192.168.25.0/24 [110/782] via 172.16.23.3, 00:54:45, Serial0/1
O 192.168.40.0/24 [110/782] via 172.16.23.3, 00:54:45, Serial0/1
 172.16.0.0/24 is subnetted, 6 subnets
C   172.16.23.0 is directly connected, Serial0/1
C   172.16.12.0 is directly connected, Serial0/0
D   172.16.1.0 [90/40640000] via 172.16.12.1, 00:09:02, Serial0/0
C   172.16.2.0 is directly connected, Loopback0
O   172.16.3.0 [110/782] via 172.16.23.3, 00:54:46, Serial0/1
C   172.16.100.0 is directly connected, Loopback100
O 192.168.20.0/24 [110/782] via 172.16.23.3, 00:54:46, Serial0/1
D 192.168.51.0/24 [90/40640000] via 172.16.12.1, 00:09:03, Serial0/0
D 192.168.50.0/24 [90/40640000] via 172.16.12.1, 00:09:03, Serial0/0
O 192.168.35.0/24 [110/782] via 172.16.23.3, 00:54:48, Serial0/1
D 192.168.70.0/24 [90/40640000] via 172.16.12.1, 00:09:05, Serial0/0
O IA 192.168.8.0/22 [110/782] via 172.16.23.3, 00:54:48, Serial0/1
D 192.168.48.0/23 [90/40640000] via 172.16.12.1, 00:09:05, Serial0/0
O 192.168.48.0/22 is a summary, 00:08:56, Null0
```

R3#show ip route

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP

Gateway of last resort is not set

```
C 192.168.30.0/24 is directly connected, Loopback30
C 192.168.8.0/24 is directly connected, Loopback8
C 192.168.25.0/24 is directly connected, Loopback25
C 192.168.9.0/24 is directly connected, Loopback9
C 192.168.10.0/24 is directly connected, Loopback10
C 192.168.40.0/24 is directly connected, Loopback40
 172.16.0.0/24 is subnetted, 6 subnets
C   172.16.23.0 is directly connected, Serial0/1
O E2 172.16.12.0 [175/20] via 172.16.23.2, 00:06:54, Serial0/1
O E2 172.16.1.0 [175/20] via 172.16.23.2, 00:06:54, Serial0/1
O E2 172.16.2.0 [175/20] via 172.16.23.2, 00:06:54, Serial0/1
C   172.16.3.0 is directly connected, Loopback0
O IA 172.16.100.0 [115/782] via 172.16.23.2, 00:06:54, Serial0/1
C 192.168.11.0/24 is directly connected, Loopback11
C 192.168.20.0/24 is directly connected, Loopback20
C 192.168.35.0/24 is directly connected, Loopback35
O E2 192.168.70.0/24 [175/20] via 172.16.23.2, 00:06:55, Serial0/1
O 192.168.8.0/22 is a summary, 00:06:55, Null0
O E2 192.168.48.0/22 [175/20] via 172.16.23.2, 00:06:55, Serial0/1
```

Solution: Network connectivity from LAB 5.2

Assignment 5: BGP (6p)

Describe the difference between the operation of IBGP and EBGP areas (make picture).
Explain synchronization rule.

*Recall that when BGP is running between routers in different autonomous systems, it is called EBGP. Generally, routers running EBGP are directly connected to each other. For two routers to exchange BGP routing updates, the TCP-reliable transport layer on each side must successfully pass the TCP three-way handshake before the BGP session can be established. Therefore, the IP address used in the BGP **neighbor** command must be reachable without using an IGP, which can be accomplished by pointing at an address that is reachable through a directly connected network or by using static routes to that IP address.*

*When BGP runs between routers within the same autonomous system, it is called IBGP. IBGP exchanges BGP information so that all BGP speakers have the same BGP routing information about outside autonomous systems. Routers running IBGP do not have to be directly connected to each other as long as they can reach each other so that TCP handshaking can be performed to set up the BGP neighbor relationships. The IBGP neighbor can be reached by a directly connected network, static routes, or by the internal routing protocol. Because multiple paths generally exist within an autonomous system to reach the other IBGP routers, a loopback address is usually used in the BGP **neighbor** command to establish the IBGP sessions.*

The BGP synchronization rule states that a BGP router should not use, or advertise to an external neighbor, a route that is learned from IBGP unless that route is local or the router learns it from the IGP. In other words, BGP and the IGP must be synchronized before BGP can use networks that are learned from an IBGP neighbor.

Assignment 6: IS-IS (3p)

Describe features of IS-IS.

The features of IS-IS include the following:

- *Hierarchical routing*
- *Classless behavior*
- *Rapid flooding of new information*
- *Fast convergence*
- *Very scalable*
- *Flexible timer tuning*
- *Protocol independent*

Assignment 7: IS-IS (3p)

In IS-IS, what does Level 2 routing mean?

Level 2: *Learn about paths between areas (interarea).*

Assignment 8: Multicasting (3p)

How to determine the Multicast Ethernet Address?

Multicast MAC addresses always begin with the low-order bit (0x01) in the first octet. Specifically, the 0x01005e prefix (plus the next lower bit, which is zero) has been reserved for mapping Layer 3 IP multicast addresses into Layer 2 MAC addresses. The complete multicast MAC address range is from 0100.5e00.0000 through 0100.5e7f.ffff.

This makes the first 25 bits of the MAC address fixed (24 bits plus the zero bit) and allows for the last 23 bits of the MAC address to correspond to the last 23 bits in the IP multicast group address. The translation between IP multicast and MAC address is achieved by the mapping of the low-order 23 bits of the IP (Layer 3) multicast address into the low-order 23 bits of the IEEE (Layer 2) MAC address.

Assignment 9: Route optimization (3p)

Which factors have most impact on route redistribution?

Factors that have the most impact on redistribution include:

- *Metrics*
- *Administrative distance*
- *Classful/classless capabilities of the protocols*

Assignment 10: Multicasting (3p)

Explain the difference between the forwarding of a unicast IP packet and the forwarding of a multicast IP packet.

Multicast may be used to send the same data packets to multiple receivers. By sending the data packets to multiple receivers, the packets are not duplicated for every receiver but are sent in a single stream.

To send data to multiple destinations using unicast, the sender has to send the same data flow to each receiver separately. The sender has to make copies of the same packet and send them once for each receiver.