minutes. RIP routers go through a period of a hold-down and garbage collection and slowly time-out information that has not been received recently. This is inappropriate in large environments and could cause routing inconsistencies.

RIP has no concept of network delays and link costs. Routing decisions are based on hop counts. The path with the lowest hop count to the destination is always preferred even if the longer path has a better aggregate link bandwidth and less delays.

RIP networks are flat networks. There is no concept of areas or boundaries. With the introduction of classless routing and the intelligent use of aggregation and summarization, RIP networks seem to have fallen behind.

Some enhancements were introduced in g -1trodrs gn of clP neclen dlP n2RIP r2 addrs reshe

reached according to the cost calculated in the tree.

Areas and Border Routers

As previously mentioned, OSPF uses flooding to exchange link-state updates between routers. Any change in routing information is flooded to all routers in the network. Areas are introduced to put a boundary on the explosion of link-state updates. Flooding and calculation of the Dijkstra algorithm on a router is limited to changes within an area. All routers within an area have the exact link-state database. Routers that belong to multiple areas, and connect these areas to the

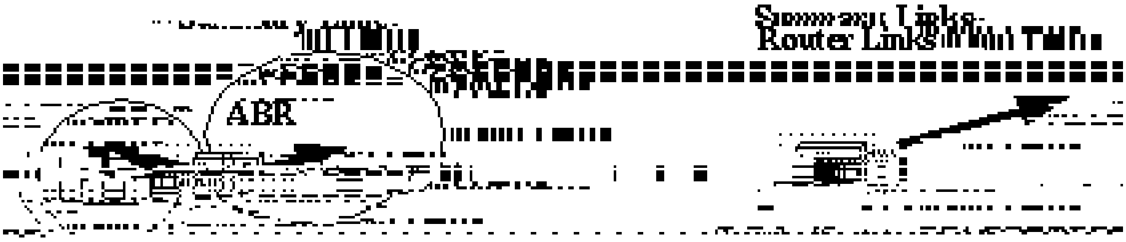

As indicated above, the router links are an indication of the state of the interfaces on a router belonging to a certain area. Each router will generate a router link for all of its interfaces. Summary links are generated by ABRs; this is how network reachability information s; reaissemined by woreenj 0 0 [illegible] Area 0.0.0.0

area 0 authentication message-digest

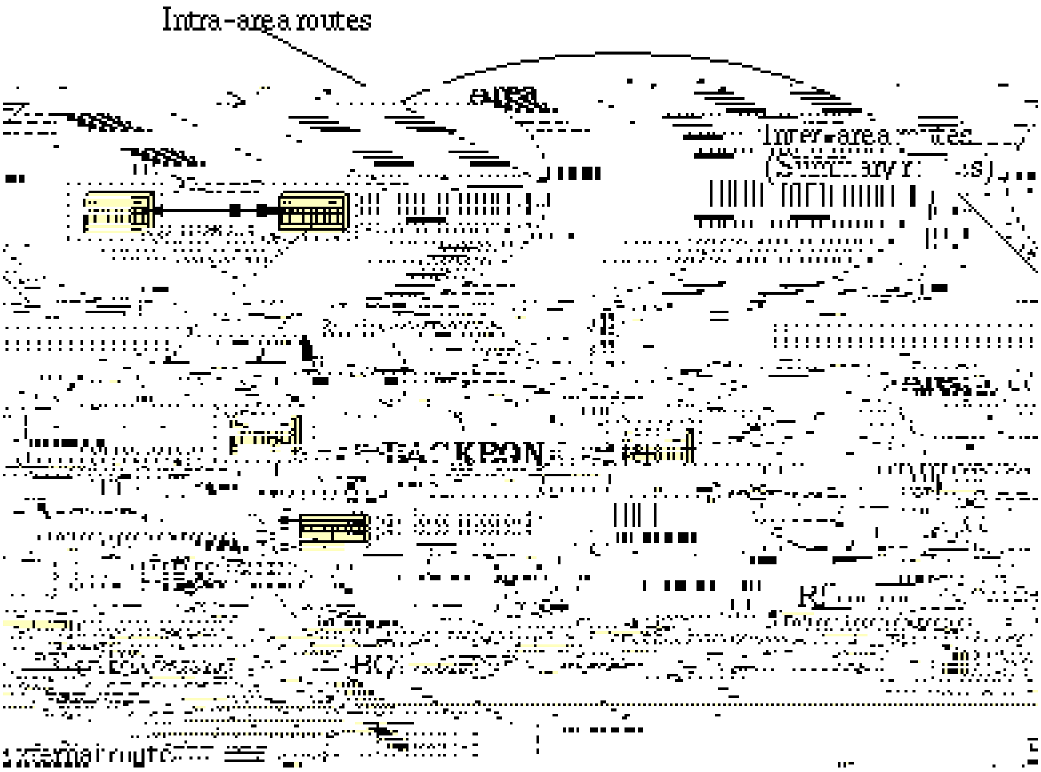

Linking an area that does not have a physical connection to the backbone.

Patching the backbone in case discontinuity of area 0 occurs.

Areas Not Physically Connected to Area 0

As mentioned earlier, area 0 has to be at the center of all other areas. In some rare case where it is impossible to have an area physically connected to the backbone, a virtual link is used. The virtual link will provide the disconnected area a logical path to the backbone. The virtual link has to be established between two ABRs that have a common area, with one ABR connected to the backbone. This is illustrated in the following example:

In this example, area 1 does not have a direct physical connection into area 0. A virtual link has to be configured between RTA and RTB. Area 2 is to be used as a transit area and RTB is the entry point into area 0. This way RTA and area 1 will have a logical connection to the backbone. In order to configure a virtual link, use the **area** *<area-id>* **virtual-link** *<RID>* router OSPF sub-command on both RTA and RTB, where area-id is the transit area. In the above diagram, this is area 2. The RID is the router-id. The OSPF router-id is usually the highest IP address on the box, or the highest loopback address if one exists. The router-id is only calculated at boot time or anytime the OSPF process is restarted. To find the router-id, use the **show ip ospf interface** command. Assuming that 1.1.1.1 and 2.2.2.2 are the respective RIDs of RTA and RTB, the OSPF configuration for both routers would be:

```
RTA#
router ospf 10
area 2 virtual-link 2.2.2.2


RTB#
router ospf 10
area 2 virtual-link 1.1.1.1
```

Partitioning the Backbone

OSPF allows for linking discontinuous parts of the backbone using a virtual link. In some cases, different area 0s need to be linked together. This can occur if, for example, a company is trying to merge two separate OSPF networks into one network with a common area 0. In other instances, virtual-links are add the box, or theOSPF allows for linking discon1p,0Nwmpossi-link failule toints example:

In the above diagram, RTA and RTB have the same interface priority but RTB has a higher RID.

```
 ip address 203.250.14.2 255.255.255.0


router ospf 10
 network 203.250.0.0 0.0.255.255 area 0.0.0.0
```

The above is a simple example that demonstrates a couple of commands that are very useful in debuggin

debuggin

The abov1 is a s8.75 6rg (show50.1 0 -1in0 Tface4 Td (g14.4 Td/F3 is a si<in0 Tface>4 Td-225 66rou because of the higher RID. In the same way, RTA was elected as BDR. RTD and RTB are neither a DR or BDR and their state is DROTHER.

Also note the neighbor count and the adjacent count. RTD has three neighbors and is adjacent to two of them, the DR and the BDR. RTF has three neighbors and is adjacent to all of them becauseTd 0 -1 diagram where DR selection is very important:

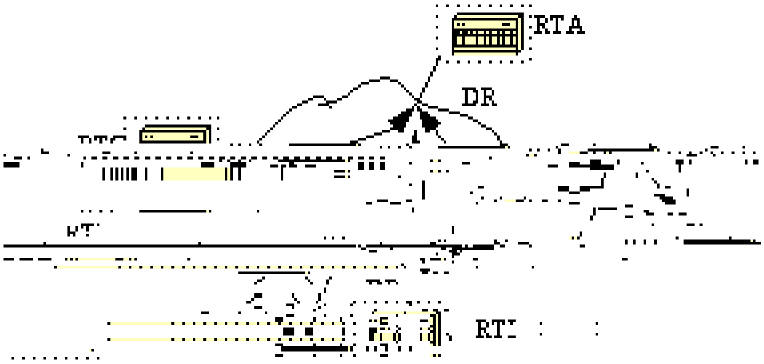

In the above diagram, it is essential for RTA's interface to the cloud to be elected DR. This is because RTA is the only router that has full connectivity to other routers. The election of the DR could be influenced by setting the ospf priority on the interfaces. Routers that do not nee Q qdrG-/ers.s gr

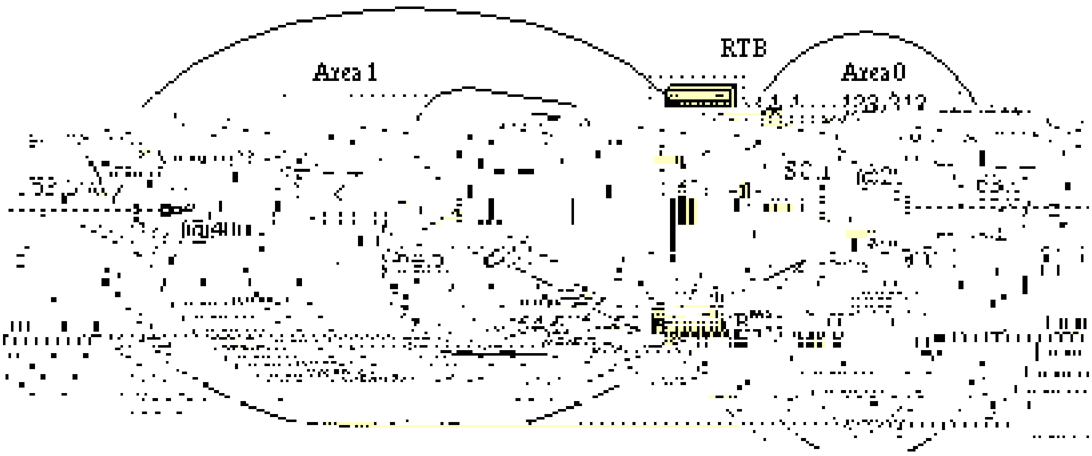

RTB
Area 1
Area 0

```
RTB#

interface Serial0
 ip address 128.213.10.2 255.255.255.0
 encapsulation frame-relay
 ip ospf network point-to-multipoint

interface Serial1
 ip address 123.212.1.1 255.255.255.0

router ospf 10
network 128.213.0.0 0.0.255.255 area 1
network 123.212.0.0 0.0.255.255 area 0
```

Note that no static frame relay map statements were configured; this is because Inverse ARP takes care of the DLCI to IP address mapping. Let us look at some of **show ip ospf interface** and **show ip ospf route** outputs:

```
RTA#show ip ospf interface s0 Serial0 is up, line protocol is up Internet Address 128.213.10.1
255.255.255.0, Area 0 Process ID 10, Router ID 200.200.10.1, Network Type POINT_TO_MULTIPOINT,
Cost: 64 Transmit Delay is 1 sec, State POINT_TO_MULTIPOINT, Timer intervals configured, Hello
```

static mapping on RTC to be able to reach next hop 128.213.10.2.

Broadcast Interfaces

This approach is a workaround for using the "neighbor" command which statically lists all existing neighbors. The interface will be logically set to broadcast and will behave as if the router were connected to a LAN. DR and BDR electu 12manWegic -1rhbosd lecb(wspproal care shouldc -1takenere)T

Prior to IOS 12.0, Cisco was compliant with the then-current RFC 1583. As of IOS 12.0, Cisco changed the behavior of OSPF to be compliant with the new standard, RFC 2328. This situation created the possibility of sub-optimal routing if all of the ABRs in an area were not upgraded to the new code at the same time. In order to address this potential problem, a command has been added to the OSPF configuration of Cisco IOS that allows you to selectively disable compatibility with RFC 2328. The new configuration command is under

Also, an ASBR cannot be internal to a stub area. These restrictions are made because a stub area is mainly configured not to carry external routes and any of the above situations cause external links to be injected in that area. The backbone, of course, cannot be configured as stub.

All OSPF routers inside a stub area have to be configured as stub routers. This is because whenever an area is configured as stub, all interfaces that belong to that area will start exchanging Hello packets with a flag that indicates that the interface is stub. Actually this is just a bit in the Hello packet (E bit) that gets set to 0. All routers that have a common segment have to agree on

i S oioutes and add kdefaultof c0.0.0.0are mdd koly coutes anjected in o that area

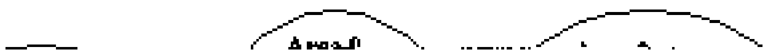

subnets

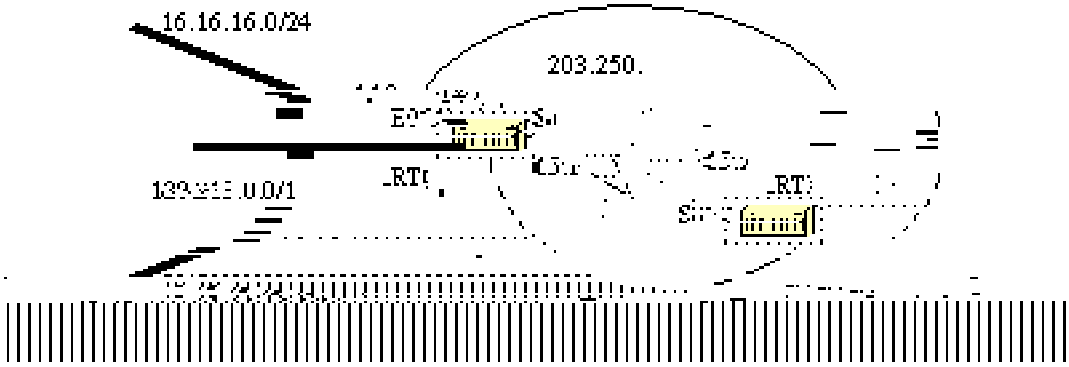
16.16.16.0/24
203.250.

```
connected, Serial0 O IA 203.250.14.0 [110/74] via 203.250.15.1, 00:00:02, Serial0 O E2
128.213.0.0 [110/50] via 203.250.15.1, 00:00:02, Serial0
```

Note that 16.16.16.0 has shown up now and the cost to external routes is 50. Since the external routes are of type 2 (E2), the internal cost has not been added. Suppose now, we change the type to E1:

ro static metric is metric-of ty1 subnets(50)Tj /F4 9 T RTE#10/50

```
C 203.250.15.0 255.255.255.252 is directly connected, Serial0 O 203.250.15.64 255.255.255.192 [110/74] via 203.250.15.1, 00:15:55, Serial0 RTA#show ip route Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default Gateway of last resort is not set 203.250.15.0 255.255.255.192 is subnetted, 1 subnets C 203.250.15.64 is directly connected, Ethernet0
```

Note that RTE has recognized that 203.250.15.0 has two subnets while RTA thinks that it has only one subnet (the one configured on the interface). Information about subnet 203.250.15.0 255.255.255.252 is lost in the RIP domain. In order to reach that subnet, a static route needs to be configured on RTA:

```
RTA#
 interface Ethernet0
  ip address 203.250.15.68 36 502.16 m 180.3 502.16 l S 0 G 1 w qc,
anetwork 203.250.15.0

ip route 203.250.15.0 255.255.255.0 203.250.15.67
```

This way RTA will be able to reach the other subnets.

Mutual Redistribution

Mutual redistribution between protocols should be done very carefully and in a controlled manner. Incorrect configuration could lead to potential looping of routing information. A rule of thumb for

msamecrotocols. Passiv interfacesand iistributie lstr should be dapplid on the iedistributiongTj 0 0 Td 0 -1

```
RTE#
 interface Ethernet0
  ip address 203.250.15.130 255.255.255.192

 interface Serial0
  ip address 203.250.15.2 255.255.255.192

 router rip
  network 203.250.15.0


RTC#
 interface Ethernet0
  ip address 203.250.15.67 255.255.255.192

 interface Serial1
  ip address 203.250.15.1 255.255.255.192

 router ospAj 03pce Serial1
  network 203.250.15.0




  ip address 2050.15.0


RTC#

RTC#
 interface Eternet0
  ip address 23.250.15.67 255.255.2558192

 router ospAj03pce Serial1
  network 203.50.15.0
```

whether the ASBR has a default route. The latter can be set by adding the keyword **always**. You should be careful when using the **always** keyword. If your router advertises a default (0.0.0.0) inside the domain and does not have a default itself or a path to reach the destinations, routing will be broken.

The metric and metric type are the cost and type (E1 or E2) assigned to the default route. The route map specifies the set of conditions that need to be satisfied in order for the default to be generated.

Assume that RTE i/F2 12 Tf alws2n
route map esorition203.250.15.2. 12C0 0 Tea depropag0 0ult to be

people have different approaches to designing OSPF networks. The important thing to remember is that any protocol can fail under pressure. The idea is not to challenge the protocol but rather to work with it in order to get the best behavior. The following are a list of things to consider.

Number of Routers per Area

The maximum number of routers per area depends on several factors, including the following:

- What kind of area do you have?
- What kind of CPU power do you have in that area?
- What kind of media?
- Will you be running OSPF in NBMA mode?

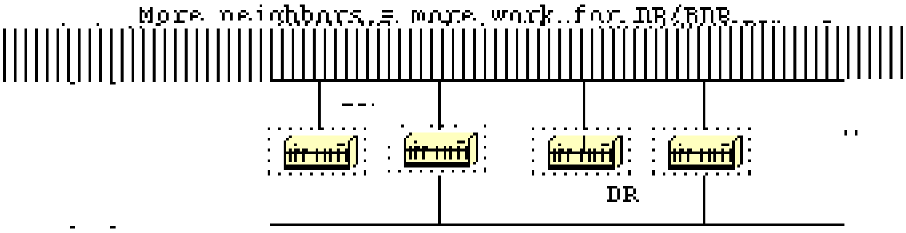

Number of Areas per ABR

ABRs will keep a copy of the database for all areas they service. If a router is connected to five areas for example, it will have to keep a list of five different databases. The number of areas per

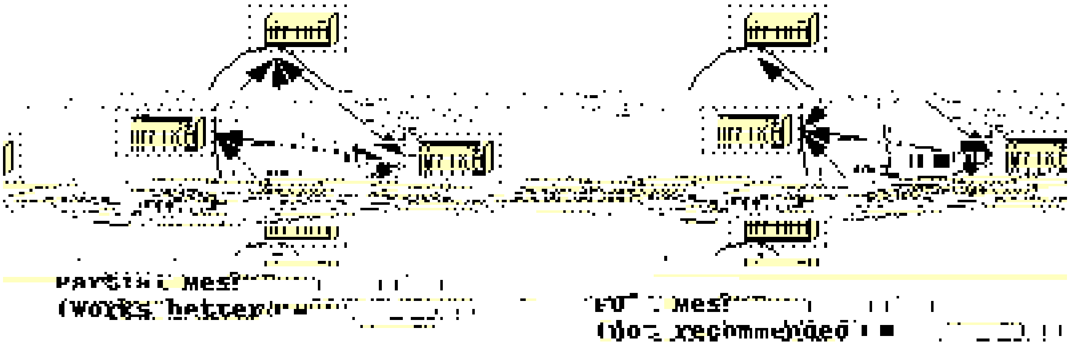

Memory Issues

initial **sequence number**. The sequence number is used to detect old or duplicate Link-State Advertisements (LSA).

In the **Exchange** state, Database Description Packets (DD) will get exchanged. These are abbreviated link-state advertisements in the form of link-state headers. The header supplies

Point-to-point links: These could be physical or logical (subinterfaces) point-to-point serial link connections. These links could be numbered (an IP address is configured on the link) or unnumbered.

Transit links: These are interfaces connected to networks that have more than one router attached, hence the name transit.

Virtual links: These are logical links that connect areas that do not have physical connections to the backbone. Virtual links are treated as numbered point-to-point links.

The **link-ID**

```
RTC# ip subnet-zero interface Ethernet0 ip address 203.250.15.67 255.255.255.192 interface
```

Link State ID: 203.250.15.67 Advertising Router: 203.250.15.67 LS Seq Number: 80000008 Checksum:

```
Summary ASB Link States (Area 0)

Link ID         ADV Router      Age    Seq#       Checksum
```

metrics: 0 TOS 0 Metrics: 10 LS age: 1575 Options: (No TOS-capability) LS Type: Router Links Link State ID: 203.250.15.67 Advertising Router: 203.250.15.67 LS Seq Number: 80000028 Checksum: 0x5666 Length: 36 Area Border Router Number of Links: 1 Link connected to: a Transit Network (Link ID) Designated Router address: 203.250.15.68 (Link Data) Router Interface address: 203.250.15.67 Number of TOS metrics: 0 TOS 0 Metrics: 10 RTC#**show ip ospf database network** OSPF Router with ID (203.250.15.67) (Process ID 10) Net Link States (Area 0) Routing Bit Set on this LSA LS age: 1725 Options: (No TOS-capability) LS Type: Network Links Link State ID: 203.250.15.68 (address of Designated Router) Advertising Router: 203.250.13.41 LS Seq Number: 80000026 Checksum: 0x6CDA Length: 32 Network Mask: 255.255.255.192 Attached Router: 203.250.13.41 Attached Router: 203.250.15.67 RTC#**show ip ospf database summary** OSPF Router with ID (203.250.15.67) (Process ID 10) Summary Net Link States (Area 1) LS age: 8 Options: (No TOS-capability) LS Type: Summary Links(Network) Link State ID: 203.250.13.41 (summary Network Number) Advertising Router: 203.250.15.67 LS Seq Number: 80000029 Checksum: 0x42D1 Length: 28 Network Mask: 255.255.255.255 TOS: 0 Metric: 11 LS age: 26 Options: (No TOS-capability) LS Type: Summary Links(Network) Link State ID: 203.250.15.64 (summary Network Number) Advertising Router: 203.250.15.67 LS Seq Number: 80000030 Checksum: 0xB182 Length: 28 Network Mask: 255.255.255.192 TOS: 0 Metric: 10 LS age: 47 Options: (No TOS-capability) LS Type: Summary Links(Network) Link

Summary Links(Network) Link State ID: 203.250.15.64 (summary Netw) Router) Adv8E0ising Router:
203.250.15.67 LS Seq Number: 80000030 Checksum: 0xB182 Attached Router: (summary Netwo4k Number) Adv3Dtising R
Network) Link Router with (Process ID 30) Net Link Stat50.15.67 LS Seq Number: 80ype:

Some special IP multicast addresses are reserved for OSPF:

224.0.0.5: All OSPF routers should be able to transmit and listen to this address.

224.0.0.6: All DR and BDR routers should be able to transmit and listen to this address.

The mapping between IP multicast addresses and MAC addresses has the following rule:

For multiaccess networks that support multicast, the low order 23 bits of the IP address are used as the low order bits of the MAC multicast address 01-005E-00-00- 00. For example:

224.0.0.5 would be mapped to 01-00-5E-00-00-05

224.0.0.6 would be mapped to 01-00-5E-00-00-06

OSPF uses broadcast on Token Ring networks.

Appendix C: Variable Length Subnet Masks (VLSM)

The following is a binary/decimal conversion chart:

1080 Tc (8)S9 8Tm 0 Tc (m 0 Tc (8)Ti 10 0 R1 64.7 877m46Tm 618.7 8m 249.056618856618202.m029 6 249.0566188568164.7 87.7 618Tm 0029 6 249.0566188568164.7 87.7 618Tm 0029 6 l 175 5.7798 6l 80.9024966.246.189.16 175Td 10 0 R1 58 08Tm 0 Tc -6.670 0 R1 58 08Tm 8 08Tm 0 Tc -6.670 0 6 249.0566188568120 0 R1 79.8m75896m R1 58 08m7589Tm 0 Tc (0)Ti 10 6246.17856246Tm 0 Tc -6.670 RTd 10 06Tm 0 Tc -6.670 RTd9896m9Tm 0 Tc -6.670 RTd E 0 Tc -6.670 RTd E 0 Tc -6.670 RTd E 0 Tc Tc -6.670 RTd E 0 8Tm 0 T488556.1849 79..8202.996618.7 825.51 -6.670 RTd 10 0 R1 0 R1 7Tm 0 Tc RTd E 0 RT 10 0 R1 446. 7RTd E 0 RT 10

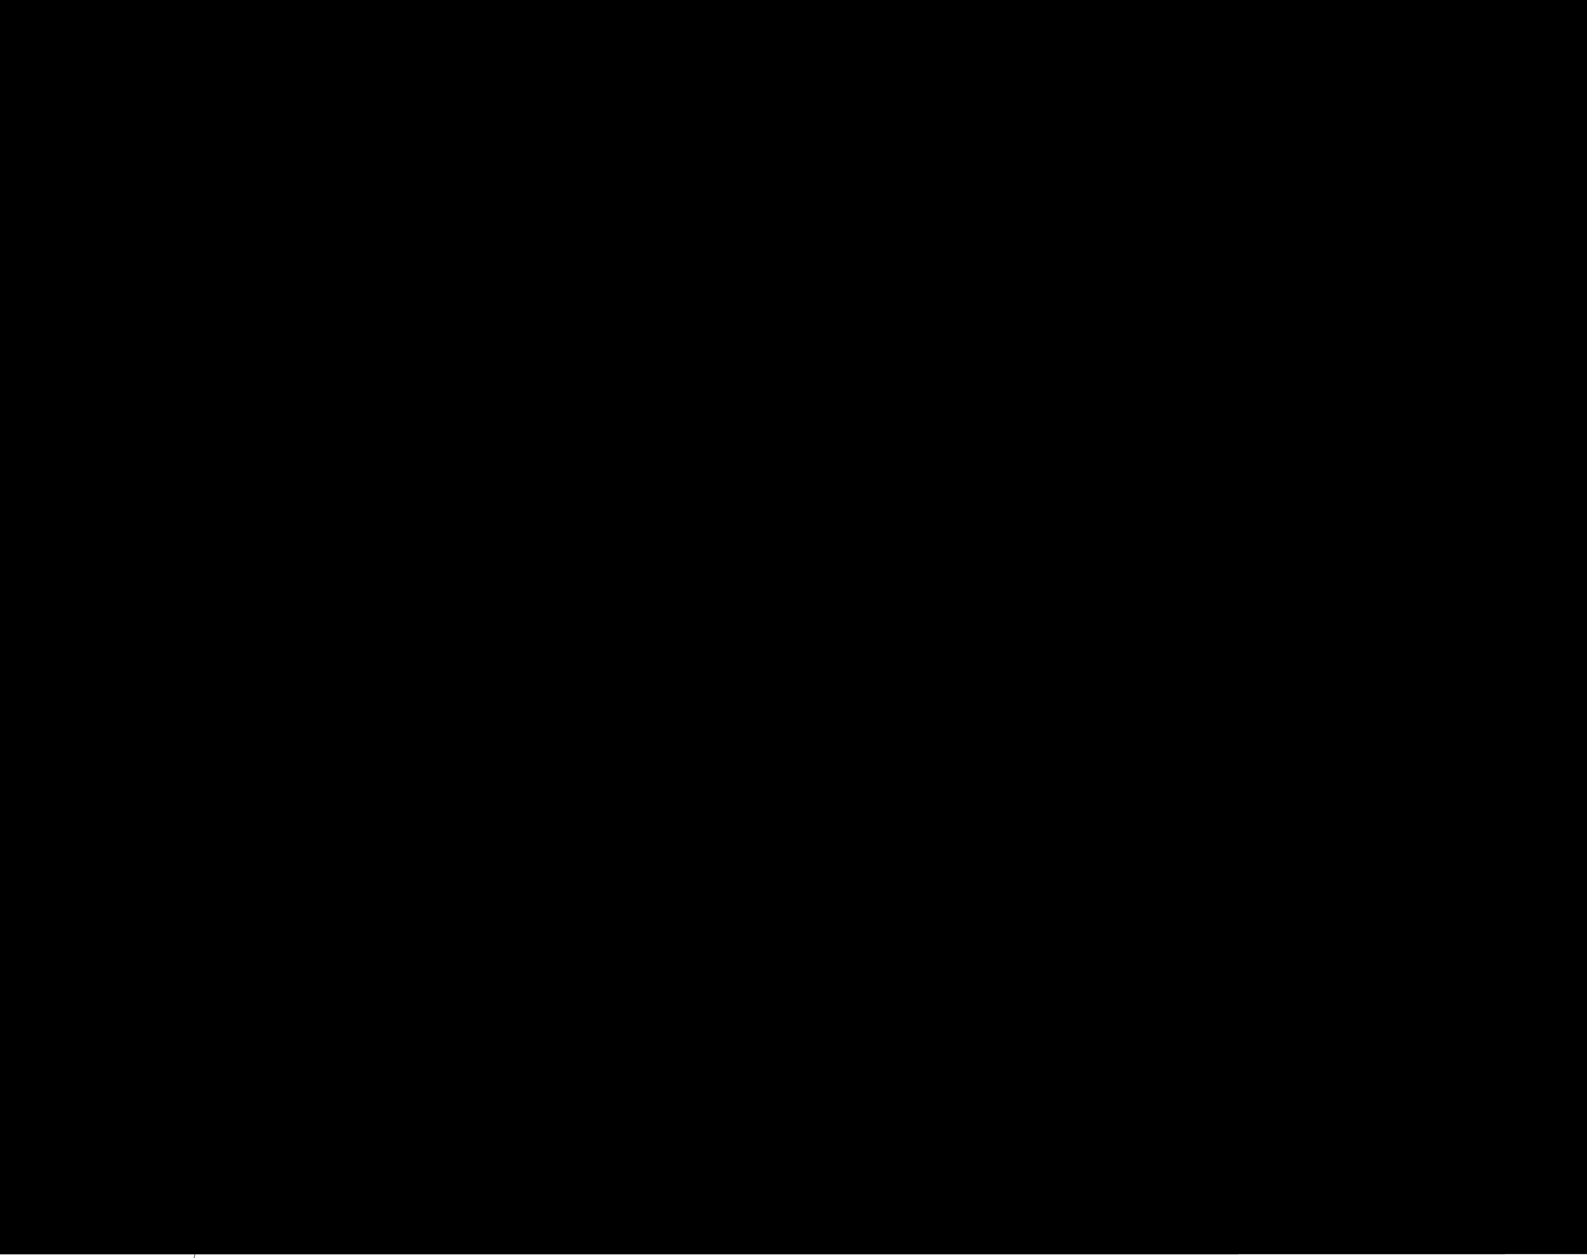

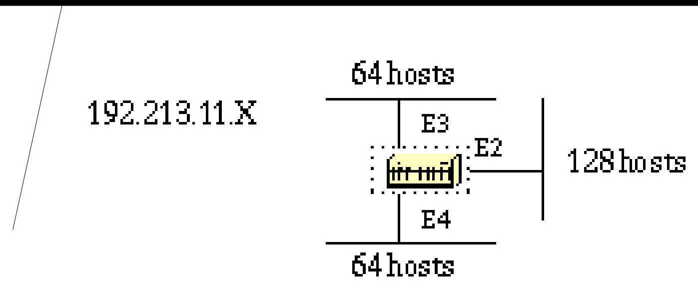

There are a handful of subnet masks that can be used; note that a mask should have a contiguous number of ones that start from the left and the rest of the bits are all 0s.

```
-252 (1111 1100) The address space is divided into 64.
 -248 (1111 1000) The address space is divided into 32.
 -240 (1111 0000) The address space is divided into 16.
 -224 (1110 0000) The address space is divided into 8.
 -192 (1100 0000) The address space is divided into 4.
 -128 (1000 0000) The address space is divided into 2.
```

Without VLSM you have the choice to use mask 255.255.255.128 and divide the addresses into 2 subnets with 128 hosts each or use 255.255.255.192 and divide the space into 4 subnets with 64 hosts each. This does not meet the requirement. If you use multiple masks, you can use mask 128