


NSTC: N42

IP Network Design Considerations: An OSPF Tutorial for z/OS CS

The Architecture

Mike Fox (mjfox@us.ibm.com)
Gwendolyn J. Dente (gdente@us.ibm.com)



Orlando, FL September 23 - September 27, 2002

Wednesday, September 25: 14:00- 15:15

Thursday, September 26: 11:30- 12:45

© IBM Corporation 2002

Abstract

> Prerequisites:

- > Knowledge of Basic Routing Concepts for Static IP Routing
- > Experience with implementing MVS or OS/3900-based TCP/IP


> Abstract: OSPF as part of OMPROUTE is the strategic direction for implementing a dynamic IP routing protocol in CS for OS/390. The link-state algorithms of OSPF bring great advantages to an IP network over those offered by the distance vector algorithms of RIP. As a result, many customers are choosing to move away from static routing and from the dynamic RIP routing of OROUTED to an OSPF implementation with OMPROUTE, available since OS/390 V2R6. To do this can be relatively simple if you understand the basic architecture of OSPF and if you understand how this architecture influences your implementation and coding choices in OS/390 (z/OS).

- > This session is the first of two dealing with OSPF in CS for OS/390 or z/OS. This first session in the two-part series introduces you to concepts like "autonomous systems," different

**Attend N64: High Availability Connections to the Mainframe
(Speaker: Mike Blair, Cisco Systems)**

you design the integration of OS/390 (z/OS) into an OSPF network and showing you how to code for several network design options.

> Acknowledgements: Many thanks to Alan Packett of the IBM Design and Development groups in Raleigh for their suggestions and several visuals. Many thanks to Mike Law of IBM Integrated Test Services for the CISCO examples.



Agenda

1. **Why Dynamic Routing and Why OSPF?**
2. **OSPF Tutorial**
 1. General Terminology
 2. OSPF and Area Types
 3. Route Aggregation
 4. Link State Routing Protocol
 5. OSPF Packet Types
 6. Designated Router
 7. Adjacencies
 8. Flooding
 9. Route Computation
 10. Link State Advertisements (LSAs)
3. **Bibliography**






1. Why Dynamic Routing?




© IBM Corporation 2002

Where does Dynamic Routing Help?

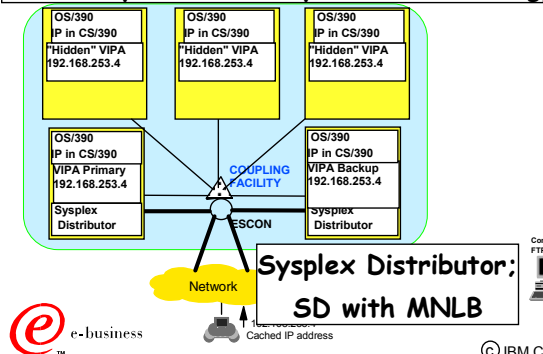
Availability: Dynamic Re-Route




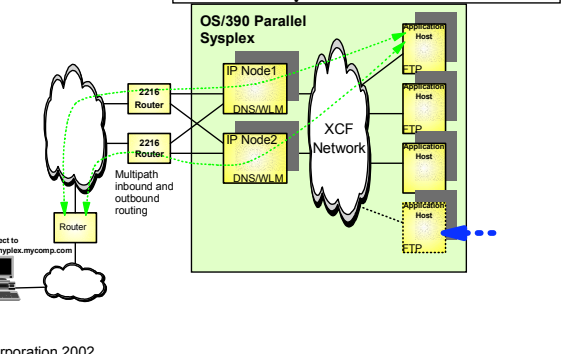
Automated Recovery & Hot Backup: VIPA Dynamics



Scalability & Availability & Load Balancing



Scalability: Seamless Growth




© IBM Corporation 2002

1. IP Non-disruptive Reroute
2. VIPA Dynamics with VIPA Takeover (V2R8) and non-disruptive Giveback in V2R10.
3. Scalability to add or subtract OS/390 images "seamlessly" - Dynamic routing allows route discovery to the new image.
4. "Sysplex Distributor" is the name given to an enhancement in V2R10 which builds on VIPA Takeover and Sysplex technology to perform IND-like function on an OS/390 TCP stack. This function will require additional storage and CPU capacity on the distributing stack, but will give additional benefits over existing IND implementations:
 1. WLM can be consulted real-time on every connection request, rather than polling periodically, so that up-to-date load information is used to select the target node/stack.
 2. Beginning in z/OS V1R2, Sysplex Distributor may be combined with the MultiNode Load Balancing (MNLB) of CISCO by coding a Service Manager in the Sysplex Distributor node.
5. There are distance vector dynamic routing protocols, like RIPV1 and RIPV2.
 1. Every 30 seconds, a router will send its full set of distance vectors to neighboring routers, i.e., a full table worth of routing data. When a router receives a set of distance vectors, it will have to recalculate the shortest path to each destination. A distance of 16 constitutes an invalid/unreachable destination in RIP.
6. There are link-state routing protocols, as with OSPF.
 1. Every router has a full map of the entire network topology, that is, all routers have an identical copy of the network database. Changes to the network topology (due to link-state changes or outages) cause a database update for those changes only (not for an entire table of data) to be propagated throughout the network. When a change is received at a router, the router recomputes its shortest path to all destinations. There is virtually no limit to the network design with OSPF.
 1. The next page's notes show you a table comparing various dynamic routing protocols.

Comparisons of Dynamic Routing Protocols

Characteristics	RIP V1	RIP V2	OSPF	EIGRP
Algorithm	Distance vector	Distance vector	Link state	Distance
Network load (note 1)	May be high	May be high	Low	Low
CPU processing requirement (note 1)	Low	Low	May be high	Low
IP network design restrictions	Many	Some	0 - Few	0 - Few
Convergence time	LT/EQ 180 sec.	LT/EQ 180 sec.	0 - immed.	0 - 5 sec.
Information exchange	Broadcast	Broadcast / Multicast	Multicast	Multicast
Multipath Outbound	No	No	Yes	Yes
Support on OS/390 or z/OS	TCP/IP V2	TCP/IP V3	OS/390 V2R6	N/a



Note 1: Depends on network size and (in)stability.

© IBM Corporation 2002

1. These three are the most widely used interior routing protocols today. Most TCP/IP products support RIP Version 1.
2. There are a few, very special situations, where a distance vector algorithm relies on the ability to 'count to infinity' in order to fully converge after a topology change.
3. Most current RIP implementations have implemented techniques to minimize the probability for those situations, but they cannot be fully eliminated.
4. When 'counting to infinity' occurs, the convergence time may be higher than 180 seconds for RIP.
5. RIP V1 is class-full
6. RIP V2 is class-less.


2. OSPF Tutorial



1. Both RFC 2328 and RFC 1583 define OSPF. However, CS for OS/390 OMPROUTE supports 1583+ and not RFC2328. Another RFC, #1587, describes a special type of OSPF area called the Not So Stubby Area. OMPROUTE does not support RFC 1587 or the draft document that proposes additional capability for NSSAs: draft-ietf-ospf-nssa-update-10.txt.
2. This tutorial focuses on OSPF as described in RFC 1583.
3. Below is a table to compare various routing protocols with each other.

Characteristics	RIP V1	RIP V2	OSPF	EIGRP
Algorithm (Distance Vector or Link State)	Distance	Distance	LinkSt.	Distance
Network load	May be high	May be high	Low	Low
CPU requirement	Low	Low	May be high	Low
IP net.design	Many	Some	0-Few	0-Few
Restrictions				
Convergence time	</= 180 sec.	</= 180 sec.	0-immed.	0-5 sec.
Information Exchange	Broadcast	Broad-/Multicast	Multicast	Multicast
Multipath Outbound	No	No	Yes	Yes
Support on OS/390	TCP/IP V2	TCP/IP V3	OS/390 R6	N/A

General Terminology: AS, BGP, ASBR




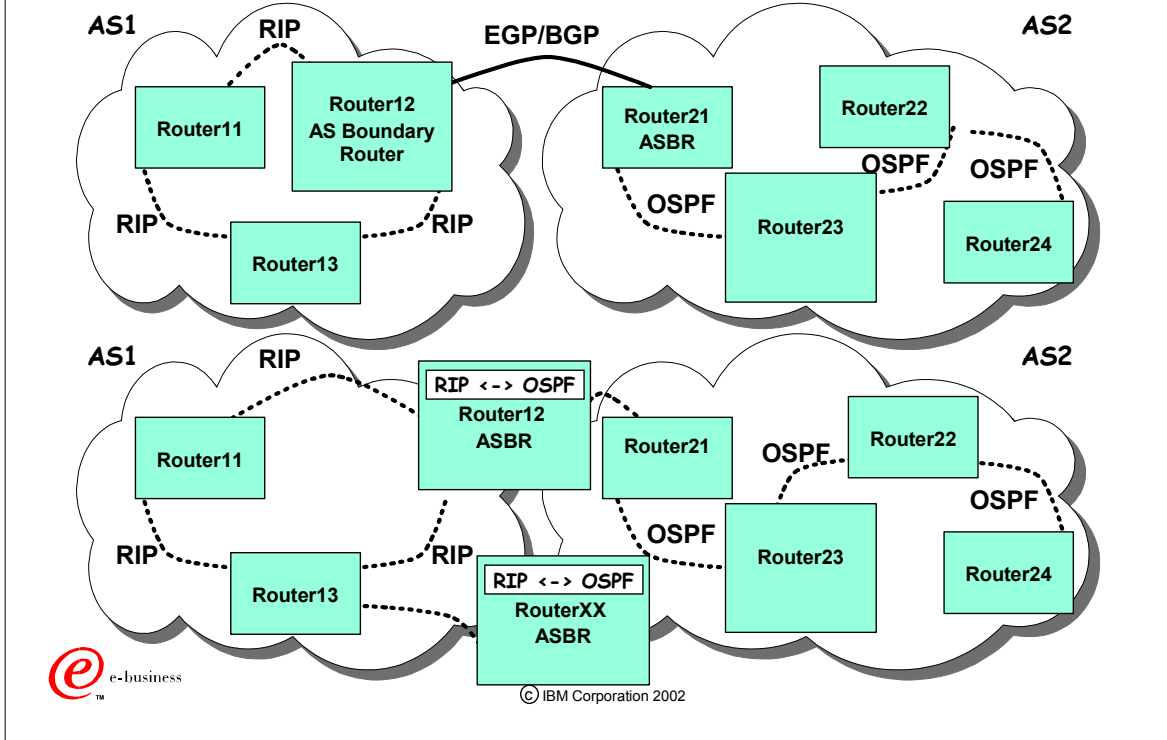
AS = Autonomous System	A collection of routers running a common routing protocol
IGP = Interior Gateway or Routing Protocol	A common routing protocol that is used among routers of an autonomous system
EGP = Exterior Gateway Protocol or BGP = Border Gateway Protocol	Protocol that may be used among multiple Autonomous Systems; alternatively use a router that speaks both IGPs. Not supported in CS for OS/390.
ASBR = AS Boundary/Border Router	A router that speaks multiple IGP routing protocols and can import/export or redistribute routes among differing IGPs, or a router that uses BGP to communicate with another router using BGP. An ASBR may sit in each AS if it is using BGP.



© IBM Corporation 2002

1. We begin here a discussion of the terminology of OSPF (Open Shortest Path First). We start with a description of only a few terms and proceed to more explanations on subsequent pages.
2. Autonomous System (AS)
 1. A group of routers running a common routing protocol. The routers connect multiple networks.
 1. Within an AS, all the routers run the same IGP.
3. ASs are interconnected by means of routers running what were formerly called External Gateway Protocols (EGPs). An enhancement to EGP, called Border Gateway Protocol (BGP), was introduced in the 1990s to provide routing loop avoidance and other improvements. The routers that interconnect Autonomous Systems are called AS Boundary Routers or AS Border Routers. (NOTE: AS Boundary Router is the preferred term so as not to be confused with another router type called an "Area Border Router.")
 1. OMPROUTE in CS for OS/390 does not support EGP or BGP; it uses another technique to interconnect Autonomous Systems. (See next page.)

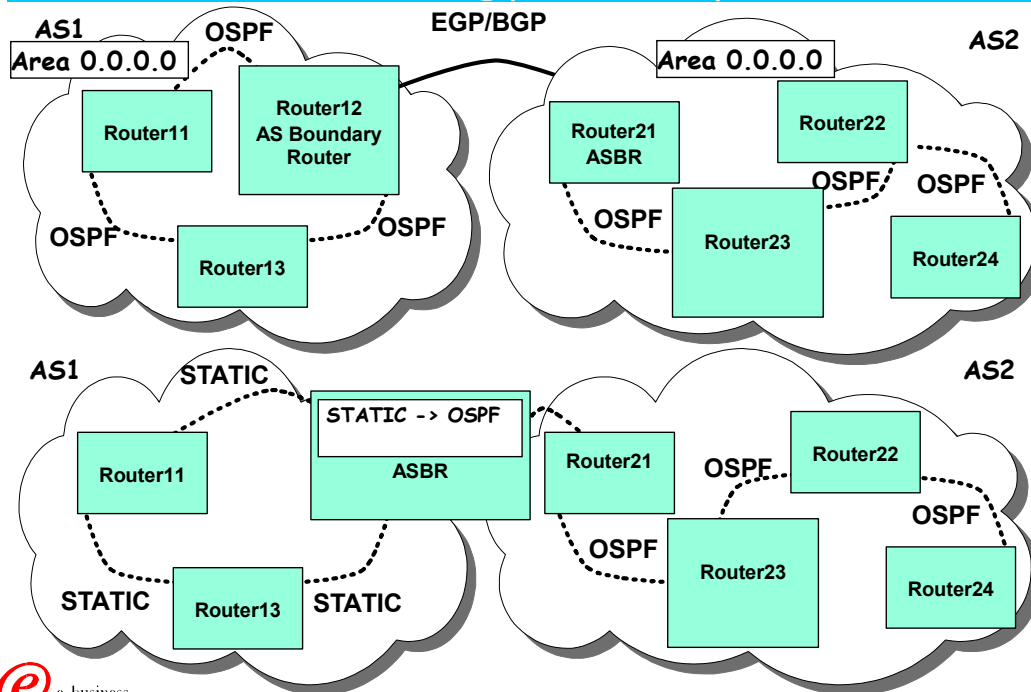
Terminology: Example 1



© IBM Corporation 2002

1. Common Interior Gateway Protocols in use today are:
 1. RIP
 2. OSPF
 3. EIGRP
 4. IGRP
 5. HELLO protocol, a protocol different from the HELLO protocol of OSPF
2. There are various flavors of BGPs in use today. BGP may interconnect two ASs of different Protocols (as you see here with RIP and OSPF) or two ASs of the same Protocol (e.g., OSPF ASx with OSPF ASy). We show you an example of this type of configuration on a subsequent page. OMPROUTE does not support the use of EGPs or BGPs, so they will not be covered in this pitch.
3. More commonly a router that communicates in several routing protocols is used to span the Autonomous Systems. A single AS Boundary Router that "redistributes" or "imports" the routing protocol of one IGP into the routing protocol of another can interconnect two AS's. This is the approach used by OMPROUTE when OMPROUTE represents an ASBR.
 1. The diagram in the bottom half of the visual depicts such a scenario.
 1. In this multiprotocol scenario, RIP routes with their metrics are imported into OSPF; OSPF routes with their Costs are imported into RIP.
4. Since metrics and costs mean different things, you must establish a route precedence methodology to determine which type of route should be preferred in the case of multiple equal-cost routes of different flavors.
5. It goes beyond the scope of this tutorial to examine these issues. For more information, please consult the IP Configuration Guide and the Information APARs on RETAIN regarding "route precedence" and "Type 1" and "Type 2" external routes.
6. If you have redundant ASBRs, as you see in the bottom half of this visual, configure them to include route tags so that the receiving AS does not attempt to re-distribute the same routes back into the AS from which they came. The route tag is set in the LSA Type 5 (External) packet. Route tags are not built by OMPROUTE, but they are passed on with the OSPF packets for the benefit of other ASBRs that may need to interpret them.

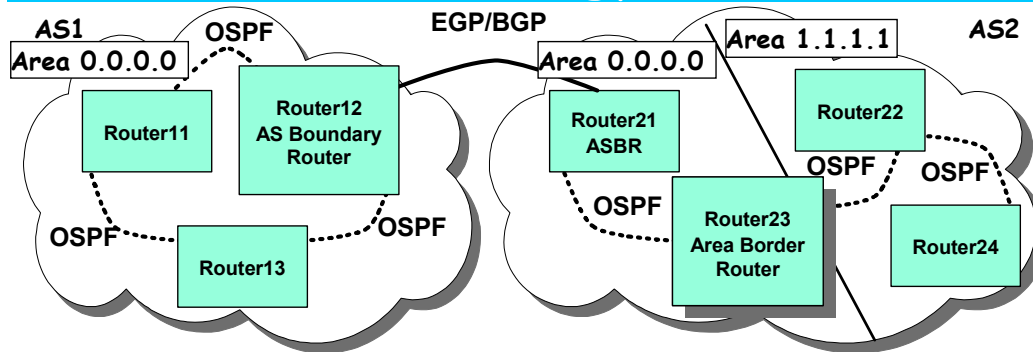
Terminology: Example 2




© IBM Corporation 2002

1. In this example you see two diagrams.
 1. The one in the top half of the visual shows how OSPF can be the IGP in each AS that is interconnected with AS Boundary Routers (ASBRs).
 1. In such a scenario, a network must employ ASBRs in each AS which is communicating with another using a Border Gateway Protocol. The single ASBR model that spans both AS's does not apply to a scenario in which multiple AS's all use OSPF as the IGP.
 2. The second diagram in the bottom half of the visual depicts an Autonomous System (AS) that uses Static Routing definitions to interface with an Autonomous System (AS) with OSPF dynamic routing protocol.
 1. A router that sits between such AS's is as much an ASBR as any other ASBR we have seen in previous diagrams.

General Terminology: Area, ABR

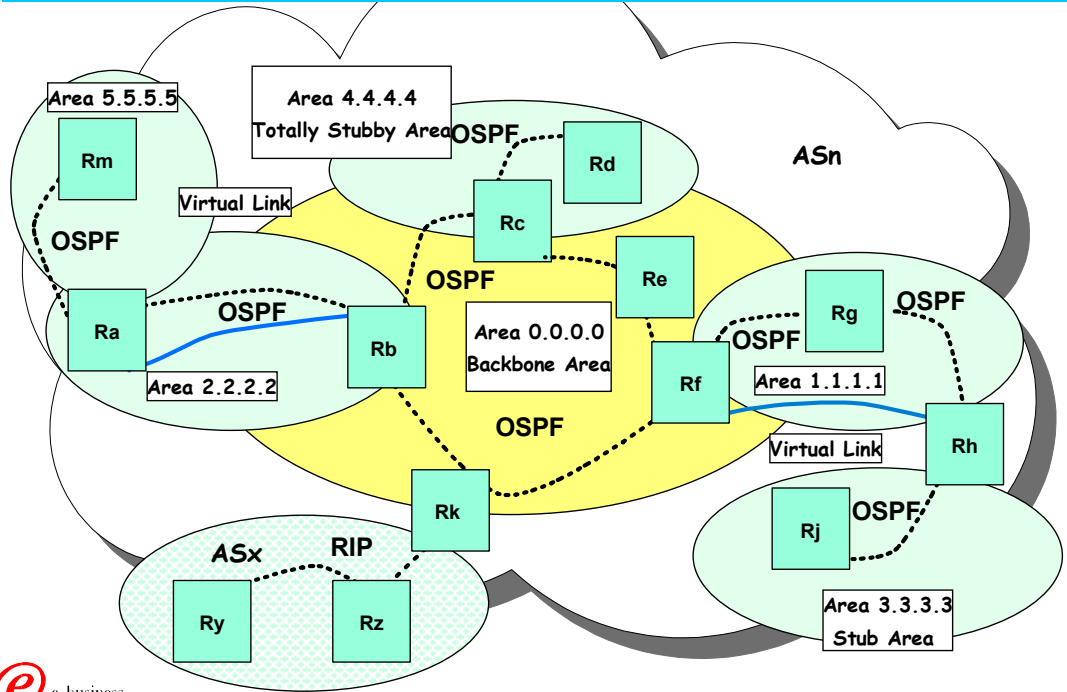


Area	<p>A subdivision of an AS</p> <ul style="list-style-type: none"> • reduces the router overhead required to keep all routers in the AS informed of the network topology; • saves storage, memory, cycles in routers that do not span the areas
ABR = Area Border Router	A router that spans multiple areas of an AS



1. We now continue our discussion of more OSPF terminology.
2. We start with the concept of "OSPF Area."
 1. If an AS is large, it may be subdivided into multiple Areas. An area is identified by a 4-octet number. You see this concept depicted in AS2: there is an Area 0.0.0.0 and another Area 1.1.1.1.
 1. Areas are interconnected by means of Area Border Routers (ABRs) which run Interior Gateway Protocols (IGPs). In our example, Router 23 represents the ABR between Area 0.0.0.0 (also called simply "Area 0"), the backbone area, and Area 1.1.1.1, a non-backbone area.
 2. The topology of an area is hidden from the rest of the Autonomous System, thus reducing routing overhead. Fewer routing updates need to be sent and smaller routing trees need to be computed and maintained, leading to a reduction in storage and CPU consumption.
 3. We discuss the different types of areas on subsequent pages.

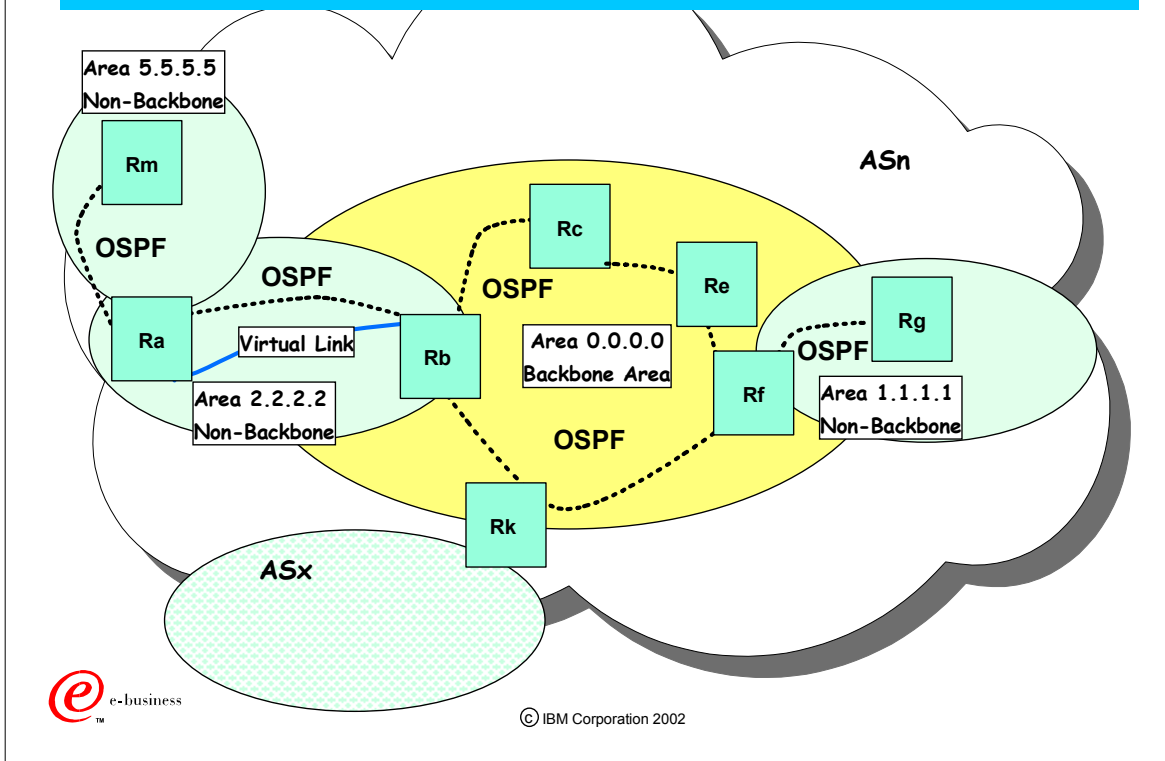
OSPF and Areas: Overview



© IBM Corporation 2002

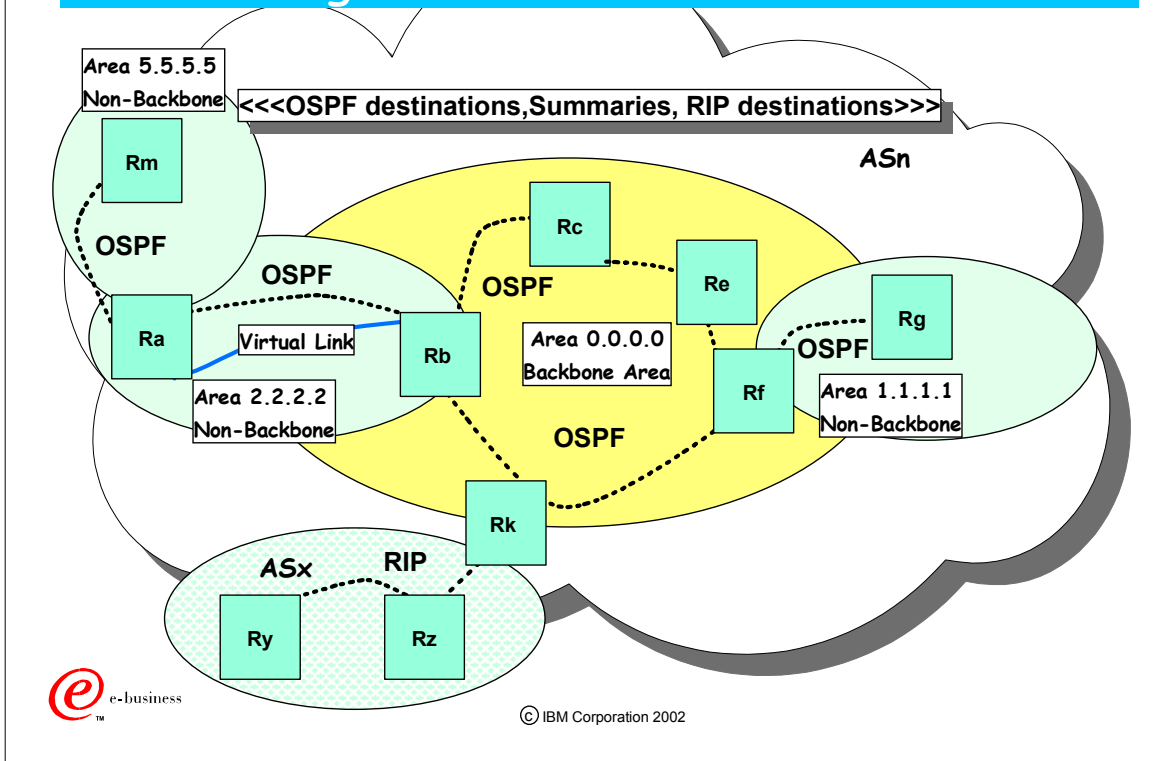
1. This diagram depicts an OSPF Autonomous System (ASn) and a RIP Autonomous System (ASx).
2. There are many details in this diagram that we will discuss separately on the following pages in order to explain the OSPF protocol.

OSPF Areas: Backbone & Non-Backbone



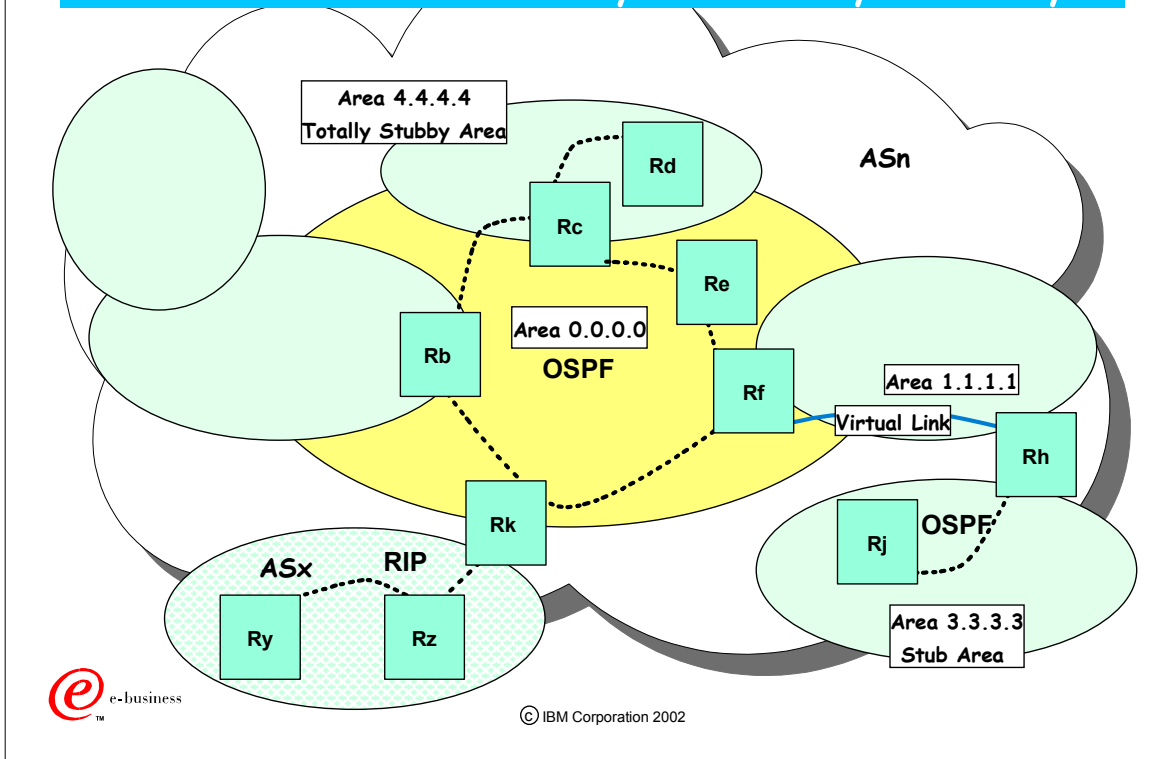
1. A backbone area is a requirement in an OSPF AS. Depending on the size of the routing trees, OSPF can be compute-intensive. It is common to populate only non-backbone areas with application servers.
2. Areas are identified by an ID that is four octets in length. The backbone area is always assigned ID 0.0.0.0. Each Area number must be unique within an AS.
3. All non-backbone areas must be connected to the backbone.
4. In this subset of the OSPF AS we depict a backbone area (0.0.0.0) and three non-backbone areas (1.1.1.1, 2.2.2.2, 5.5.5.5). Note how Areas 1.1.1.1 and 2.2.2.2 are **physically** connected via Area Border Routers Rf and Rb to the backbone. Note, however, that Area 5.5.5.5 is not physically connected to the backbone; it is physically connected to Area 2.2.2.2, a non-backbone area. This appears to violate the rule that an area must be connected to the backbone. Yet this configuration is a valid OSPF network design. Why? Because we have **logically** or **virtually** connected the non-backbone area to the backbone! We have accomplished this feat by defining a **virtual link** between the ABR Ra and the ABR Rb, which is **physically** connected to the backbone area.
5. Although not depicted, virtual links are also used to connect discontinuous areas to make them appear logically contiguous.
6. How big should an area be? This is a vendor-specific decision. It depends on the capacity of the routers to handle the link-state database and the computation of the routing trees. Network stability, number of links, IP network address assignment design all influence the size of the Area.
7. Some customers have areas up to 200 routers in size; others have areas of 50 routers in size.

Advertising: To Backbone & Non-Backbone

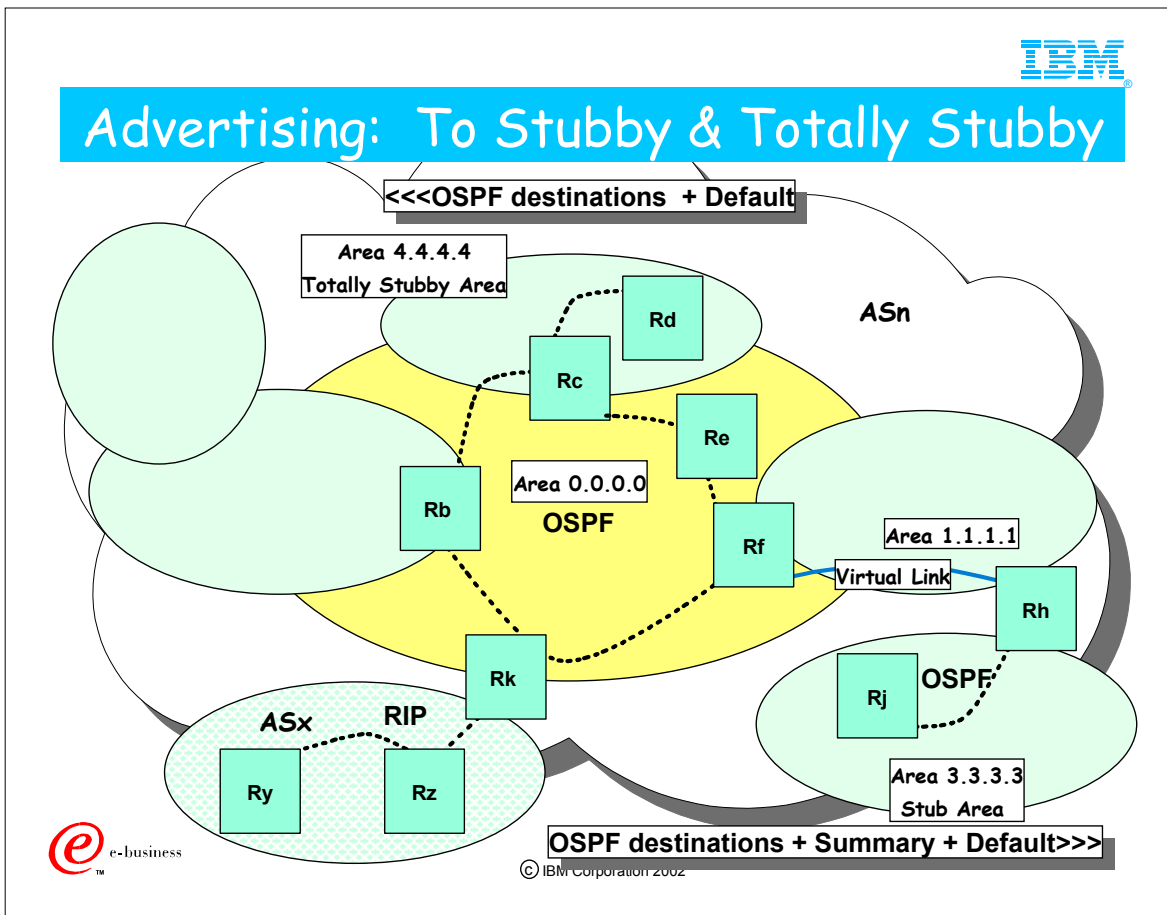


1. We have added back into this picture the RIP AS, "ASx." It is connected to the backbone area via ASBR Rk. The purpose of this visual is now to describe how routing advertisements are shipped throughout the areas so that all destinations -- both OSPF and RIP -- may be reached from anywhere in the network.
2. All routers within an area ("intra-area routers" or "interior routers") maintain a copy of the topology of the area, including information about intra-area OSPF destinations. They also know about inter-area destinations due to advertisements sent to them by the ABRs. They also know about how to reach RIP destinations, either explicitly or through a Default route.
3. The Area Border Router maintains a copy of the database for each area to which it is connected.
 1. An area border router summarizes all OSPF links into other areas and advertises these to adjacent areas.
 2. An ABR can be customized to use address ranges for route summarization. We depict route summarization later in this presentation.
4. ASBRs (in this case, Rk) import (or "redistribute") external destinations from other AS's into OSPF advertisements and originate what are called "External Link State Advertisements" known as Type 5 LSAs. Our ASBR, Rk, can thus originate information about RIP destinations and send this information into the OSPF AS.
 1. External Link State Advertisements (Type 5 LSAs) flow freely across all OSPF areas, except for Stub Areas.
 2. (External Links are RIP, Static, and Direct.)
 3. ASBRs may also be configured to send or "originate" default routes into the area or into the other AS.
5. All the OSPF Areas depicted in this subset diagram have knowledge of routes to the RIP Autonomous System.

OSPF Areas: Stubby & Totally Stubby



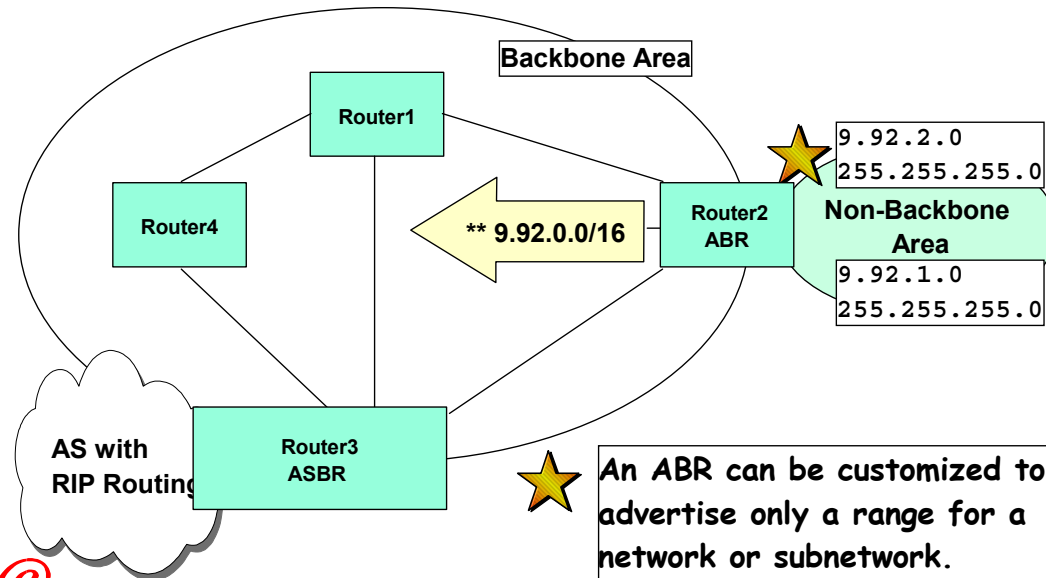
1. This subset of our original diagram now focuses on special non-backbone areas called "Stub Areas" and "Totally Stubby Areas."
2. Stub Areas minimize storage and CPU utilization at nodes that are part of the Stub Area because they maintain less knowledge about the topology of the AS than do other types of non-backbone routers. They maintain knowledge only about intra-area destinations, summaries of inter-area destinations, and a default route in order to reach external destinations in other AS's.
 1. A Stub Area is defined with STUB=YES in the Area definition.
 2. A stub area can be adjacent to the backbone (Area 4.4.4.4) or not adjacent to the backbone (Area 3.3.3.3, which is adjacent to a non-backbone area).
 3. Stub Area 3.3.3.3 is connected via a virtual link to the backbone area.
 1. As you already know, every area must be adjacent to the backbone. If non-adjacent physically, there must be a virtual link to the backbone. In our diagram, Rh becomes part of the Backbone because of the virtual link.
 2. Stub Areas can be the end-points of a virtual link, but they cannot be on the intermediate path of a virtual link.
3. A Totally Stubby Area receives less routing information than a Stub Area; it receives only Default Routes. Designing with a Totally Stubby Area minimizes the compute-intensive operations necessary to build routing trees and minimizes the storage requirements for maintaining the topology database. A Totally Stubby Area is coded inside CS for OS/390 with an Import_Summaries=NO on the AREA statement.
4. CS for OS/390 OSPF cannot be an NSSA (Not So Stubby Area), an area that is not discussed in this brief tutorial. (Such an area sends Type 7 LSAs, not supported in CS for OS/390 or z/OS.) An NSSA is described in IETF RFC 1587 and in the IETF draft document: draft-ietf-ospf-nssa-update-10.txt. CS in OS/390 or z/OS can be a Stub Area or a "Totally Stub Area," as we will see later in this series of presentations.



1. This subset of our original diagram now focuses on the means by which Stub Areas learn of intra-area, inter-area, and external destinations.
 1. A Stub Area receives Type 3 Summary LSAs for inter-area destinations, including the Default Route packaged inside a Type 3 Summary LSA.
 2. Stub Areas can have multiple default routes from different attached Area Borders. (One for each.)
 3. External Link State Advertisements (e.g., RIP, static, Direct routes described in Type 5 LSAs) cannot flow into a Stub Area. The Stub Areas can reach hosts in the RIP AS only by means of a default route made known to them by their Area Border Router.
2. What do these statements mean? Recall the function of the ABRs we just looked at:
 1. An area border router summarizes all OSPF links into other areas and advertises these to adjacent areas.
 2. An ABR can be customized to use address ranges for route summarization. We depict route summarization later in this presentation.
 3. You see in this subset diagram that an Area Border Router can interface with a Stub Area. So, an ABR summarizes inter-area destinations and sends Link State Advertisements into Stub Areas in the form of Type 3 Summary LSAs. It sends Default destinations into the Stub Areas, also in the form of Type 3 Summary LSAs. It does not, however, forward any External Destinations it has learned from ASBRs into the Stub Areas. (The External Destinations are described in Type 5 LSAs.) It does not need to, because the Stub Area is given a Type 3 Summary LSA Default way to get to any part of the network it may need to reach. The Area Border Router says, "I'm not telling you the external links, so here's a default route for you if you need to reach an external route."
 4. So, in summary, all OSPF links and Type 3 Summary LSA, including the LSA for a Default, can flow into a Stub Area. (More about Type 3 Summary LSAs later....)
3. A Totally Stubby Area receives only Default Routes. That is, just as with a Stub Area, it receives no External Routes (Type 5 External LSAs). But unlike a Stub Area, it also does not receive Type 3 Summary LSAs to reach inter-area destinations, with one exception: the Default Route that is packaged inside a Type 3 Summary LSA. It is coded inside CS for OS/390 with an Import_Summaries=NO on the AREA statement.

Route Aggregation

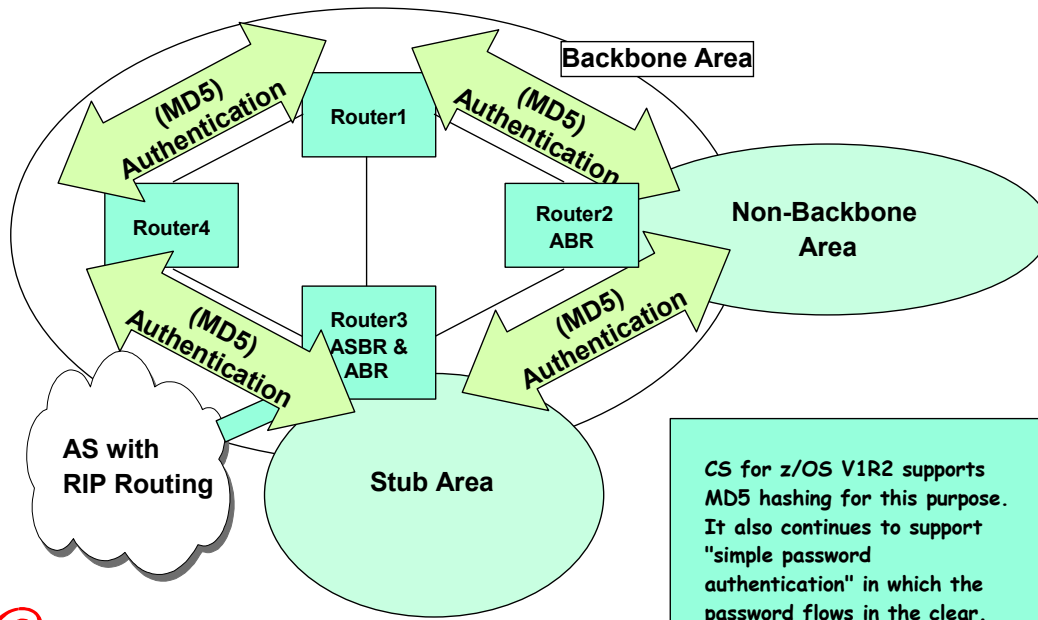
** = Router2 Advertising 9.92.0.0/16



© IBM Corporation 2002

1. In a well-designed network in which areas have been carved out according to the networks or subnetworks they serve, the ABRs can be customized to advertise a range, thus reducing the number of LSAs for inter-area routes.
2. The term "Classless InterDomain Routing" (CIDR) or supernetting is also applied to this type of Route Aggregation.
3. Within the subnet, the actual subnet mask in use may be quite different, for example, 255.255.255.0. The advertised value for the mask (255.255.0.0) specifies the mask that is common to all the subnets within 9.82.0.0 (9.82.1.0 and 9.82.2.0) when the actual mask is 255.255.255.0.
 1. CS in OS/390 allows you to specify a RANGE and "Advertise=YES/NO" in order to provide Route Summarization in the LS Advertisements.
4. Without Route Aggregation (Ranges), the LSA Summaries that are sent by Router 2 about the Non-Backbone Area are:
 1. 9.92.1.0
 2. 9.92.2.0
5. With Route Aggregation (Ranges), the LSA Summary that is sent by Router 2 about the Non-Backbone Area is:
 1. 9.92.0.0
6. Some implementations call Route Aggregation "Route Summarization" or "Summarizing Routes." Do not confuse "route summarization" with LSA Summaries, as LSA Summaries may contain either both of the non-aggregated destinations from our example (9.9.2.20 and 9.92.1.0) or may contain the aggregated destination (9.92.0.0).

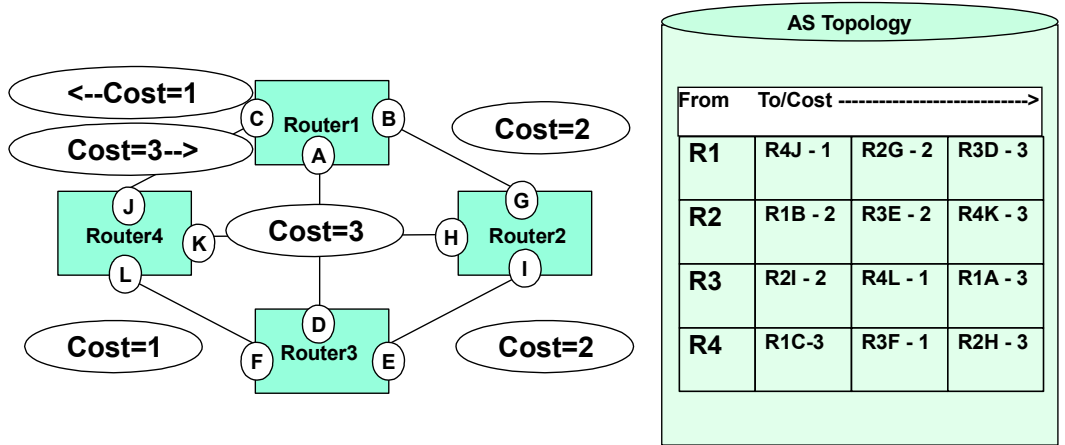
Authenticating Routers in an Area



© IBM Corporation 2002

1. Within an area the link advertisements can be authenticated so that a "rogue" router may not send routing advertisements that should not be accepted.
 1. Only trusted routers are allowed to introduce any changes.
 2. This capability protects more against the accidental introduction of a router into the area than it does against the malicious insertion of an alien router into the area.
2. zOS V1R2 Communications Server adds MD5 support. MD5 is not supported in V2R6/8 or V2R10, although there is a security mechanism with "simple password authentication."
3. IBM provides a utility, "pwtokey," that can be used to generate an MD5 authentication key. However, be aware that some routers -- Cisco being one of them -- may generate keys for MD5 authentication using nonstandard methods, so that care must be taken to ensure that Cisco and OMPROUTE define the same key. The character value entered on the Cisco side of a network connection must be converted to HEXADECIMAL in ASCII format on the IBM side of the connection and must not have been generated with the "pwtokey" utility. Please see examples of coding in the OSPF presentation that covers OMPROUTE coding for OSPF.
4. The key id is a one-byte constant that identifies the key for MD5 authentication. You should really consider it to be part of the key, and it must match on both sides of a connection. Some platforms (including Cisco) support multiple keys. The key id identifies which key is being used. IBM supports only one key, but we provide the key id for compatibility, permitting it to be defined to match what Cisco is using, for example.

Link State Cost

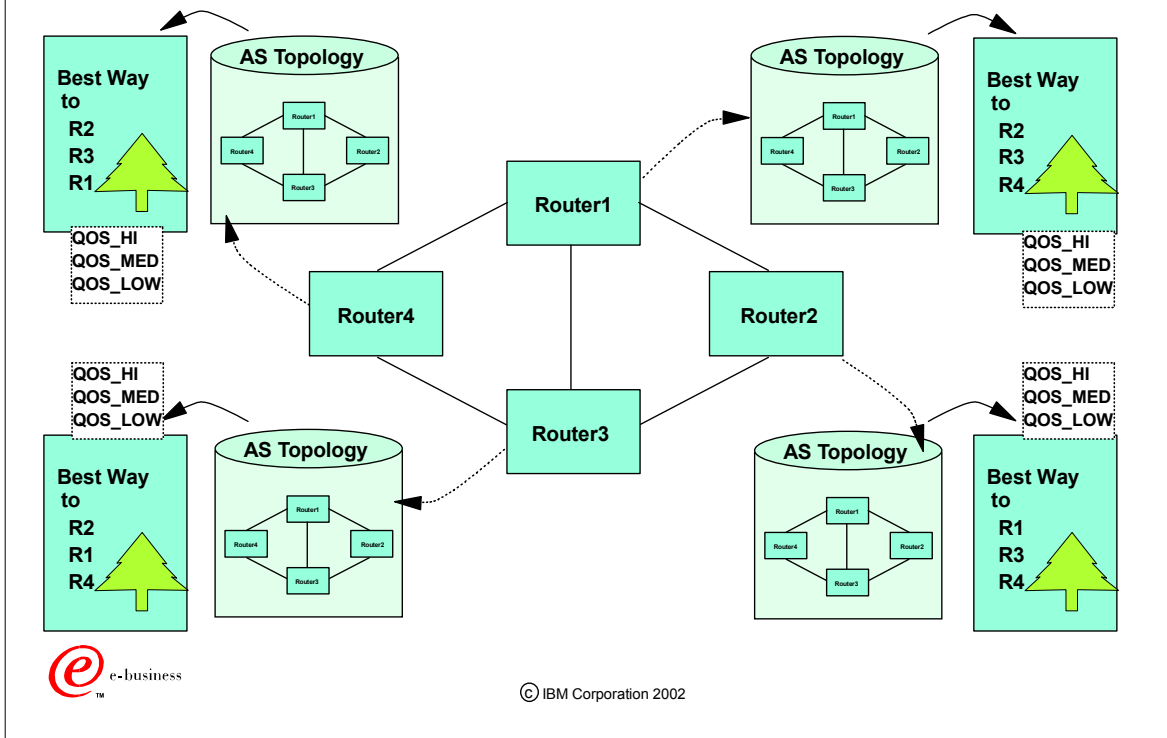


Identical copy of database in each router of the Area
Links are really to IP Adrs/Subnets - not to Routers
(A, B, C, and so on are IP Addresses.)



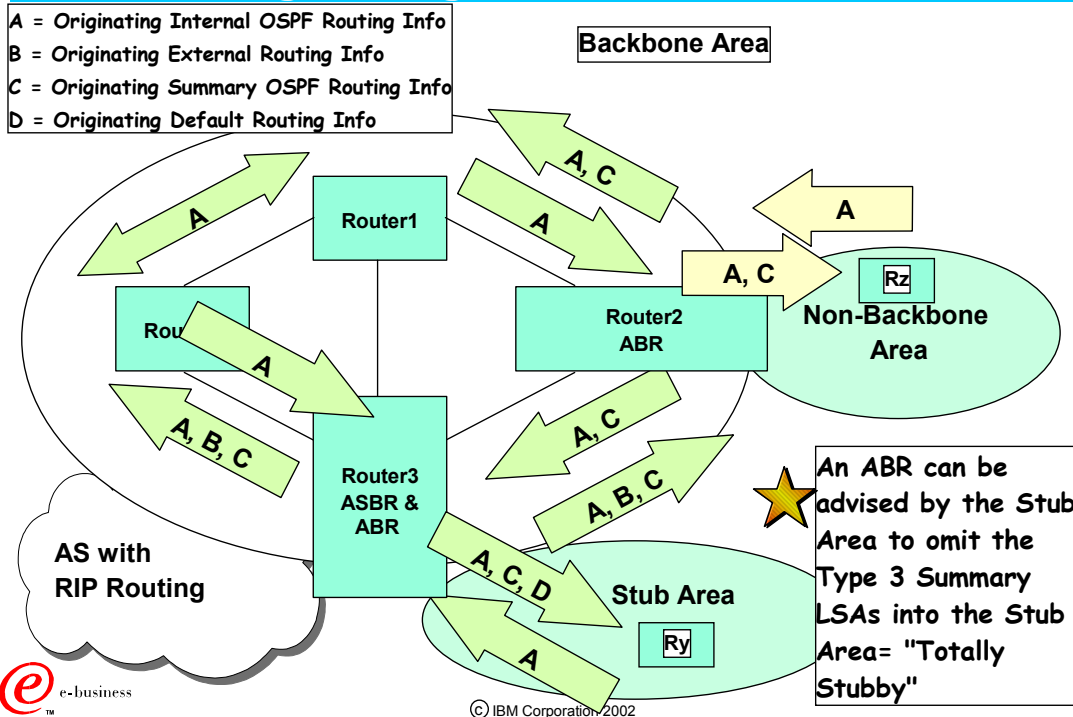
1. Every link is associated with a Cost or metric that is used to compute the route to any given destination.
 1. The lower the cost, the better the route.
 2. The cost may be assigned according to any of several metrics: largest throughput, lowest delay, lowest cost, highest speed, best reliability, and so on.
2. Each router has a map of the topology of the Area. The map contains the cost of each link in the area.
 - ★ The diagram/table has been simplified. Although it implies that the links extend to Routers, in fact they extend to IP addresses or Subnets in a Router. From Router 1 (R1) there is a link to one of the IP Addresses in R4 with a cost of 1. There is a link to one of the IP addresses in R2 with a cost of 2 and one to an IP address in R3 with a cost of 3.
3. The network administrator can assign a cost to every link in the network.
 1. Alternatively, some implementations of OSPF use the link speed to assign a default cost.
 2. Although not a requirement, it is usual to assign the same cost to a link in both directions. Consistency in the network helps to avoid unpredictable routes between two points.
 1. You may assign different costs in either direction if you are intending to influence traffic flow inbound and outbound to Router1.
 2. Coordinate link costs with the other router implementers in your network.
 3. In CS/390, the default link cost is 1.
4. Router 1 can reach Router 3 in three ways
 1. via Router 2 for a combined cost of 4
 2. via Router 4 for a combined cost of 2
 3. directly to Router 3 for a cost of 3

Basics of Link-State Routing Protocol



1. Each router maintains a database describing the topology of the area. The database is identical at each router.
 1. Each record in the Database represents one link in the network.
 2. A cost is associated with the link. This cost is platform-specific
 1. Some platforms assign costs according to throughput or link speed.
 2. Other platforms require that the costs be assigned manually.
 3. A route to a destination is the sum of the costs of all links to the destination. A lower cost represents a better route.
 1. There is virtually no limit on the total cost of a route, although in reality a cost of 65534 is the highest metric that is available to a valid route. (65535 indicates infinity, meaning that the route is not available.) (RIP, on the other hand, has a limit of 15 hops/metric count for valid routes; 16 represents an unreachable destination.)
2. Each router uses the flooding protocol to inform the other routers in the AS of the state of its interfaces and of its reachable neighbors.
3. External routes, that is, routes learned from RIP, learned statically, or otherwise, are maintained separate from the OSPF topology.
 1. OSPF and External routes are communicated throughout the Autonomous Area unless there are Stub Areas.
 1. Stub Areas do not receive External Link advertisements; they reach areas identified in External Link advertisements by sending traffic to a default route through an Area Border Router.
4. Each router computes a "tree" of best routes to each location using itself as the root.
 1. Although OMPROUTE in CS for OS/390 or CS in z/OS currently (as of z/OS V1R2) does not compute a "tree" of best routes for each Quality of Service (QOS), the RFCs for OSPF do define this capability.

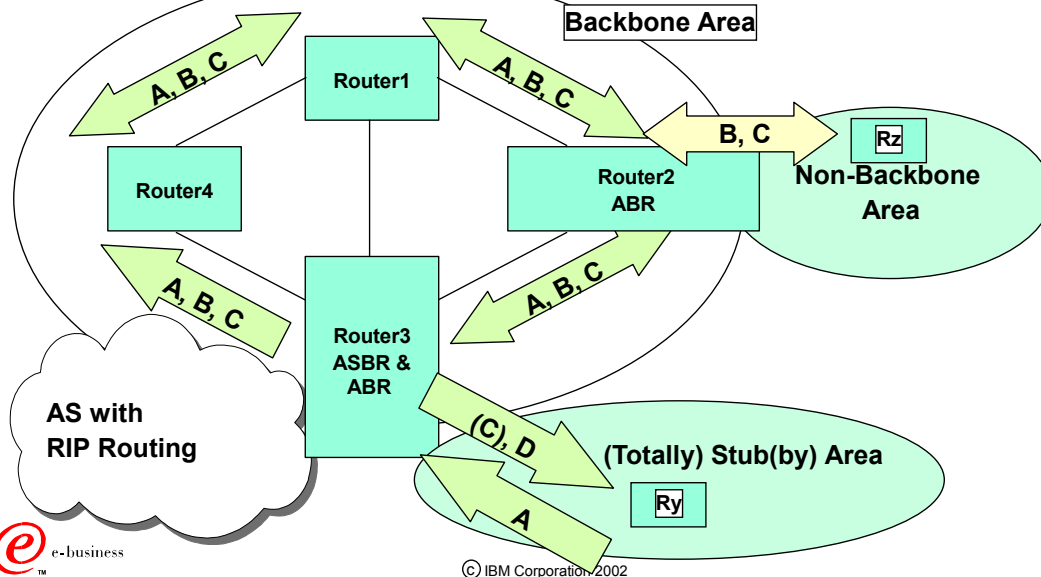
Originating LSAs: Overview



1. The whole idea of using link state routing protocols is to maintain a synchronized copy of the link state database in all nodes of an area.
2. Routers either originate route advertisements or they propagate the route advertisements received from other routers. The latter is called "flooding."
3. A routing advertisement, called a Link State Advertisement (LSA) is sent under any of the following conditions:
 1. A new link is activated or deactivated
 2. The metric of an existing link changes
 3. Every 30 minutes to refresh the databases of the neighbors
4. The ASBR (Router3) has redistributed the RIP Routes into OSPF. The redistributed (imported) routes are communicated to the other routers in the OSPF Area as External Link State Advertisements. OSPF routes are communicated as any of several types of OSPF LSAs.
5. External LSAs are not sent into a Stub Area. The Area Border Router (ABR, also Router 3) informs the Stub Area of a default route to reach any areas that are connected via RIP routes. OSPF route advertisements are also sent into the Stub Area, but they are summarized (as TYPE 3 LSAs) to indicate all the destinations that can be reached by the ABR. If the Stub Area is to be a "Totally Stubby Area," it should specify "IMPORT SUMMARIES=NO" in its AREA statement.
 1. NOTE: In reality the OSPF protocol provides for two types of Summary LSAs: Type 3 Summaries for IP destinations in other areas and Type 4 Summaries about ASBRs in the Area. LSA Type 4 Summaries pertaining to ASBRs are never advertised into Stub Areas.
6. Router 2, an ABR between the Backbone Area and the Non-Backbone Area, sends OSPF Route Summaries and External Routes into the Non-Backbone Area depicted.
7. The Stub Area could also originate a DEFAULT route, which would mean that it would send a "D." However, this would not be a good network design; the Stub Area should receive its default route from the ABR.

Flooding LSAs: Overview

A = Flooding Internal OSPF Routing Info
B = Flooding External Routing Info
C = Flooding Summary OSPF Routing Info
D = Flooding Default Routing Info




© IBM Corporation 2002

1. Once databases have been synchronized with neighbors, Route Advertisements are Flooded throughout the area. The synchronization of databases is also called "bringing up adjacencies."
 2. Each LSA has a timer on it. It is flooded again after 30 minutes in order to refresh the database.
 3. Recall that the Stub Area can be customized so as to receive only the Default Route. Omitting the inter-area route advertisements (Type 3 Summary LSAs) makes the Stub Area a "Totally Stubby Area." The Stub Area "AREA" statement simply needs to indicate "IMPORT SUMMARIES=NO."
1. NOTE: A Stub Area is never sent Type 4 Summary LSAs. (Type 4 LSAs indicate LSAs to reach ASBRs.)

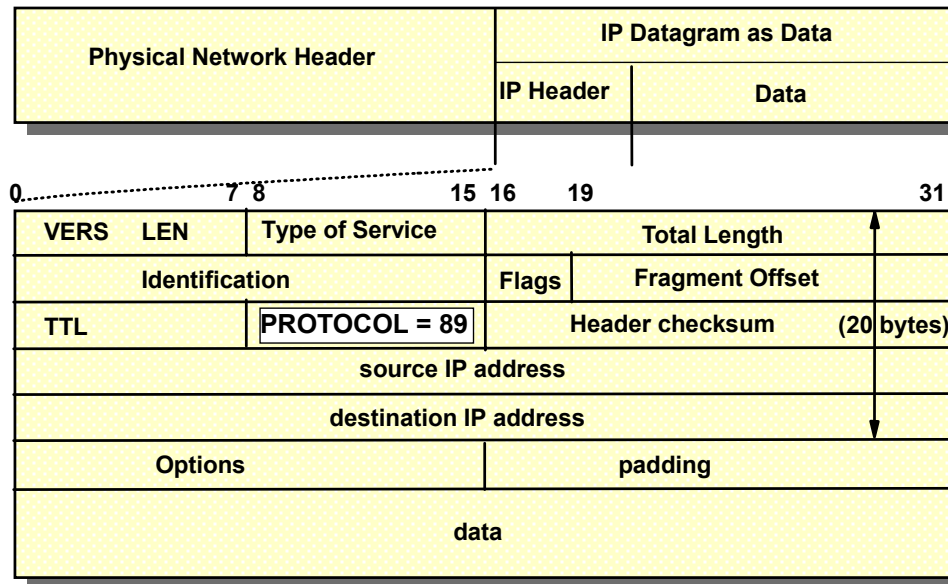
How OSPF Operates: Details

- Discover the Neighbors
- Elect a Designated Router (DR)
- Form Adjacencies
- Synchronize the DataBases
- Compute the Routing Table (Tree)
- Advertise the Link States



1. The OSPF protocol has its own protocol number in IP (number 89). It uses Raw Sockets and so, unlike TCP or UDP, does not have the concept of "Port." There is no specified port for OSPF.
2. The OSPF protocol is designed to align databases at all nodes within an area. This synchronization of databases is called "bringing up adjacencies."
3. The OSPF protocol consists of three subprotocols:
 1. Hello
 2. Exchange
 3. Flooding
4. Hello is used to check that links are operational and to select what is called a designated router. Hello packets are sent at regular intervals that are customized by the network administrator on each OSPF router.
5. The Exchange protocol is used during startup of a router where it receives a copy of the network database from a neighboring router.
6. The Flooding protocol is used to propagate link state changes through the network.

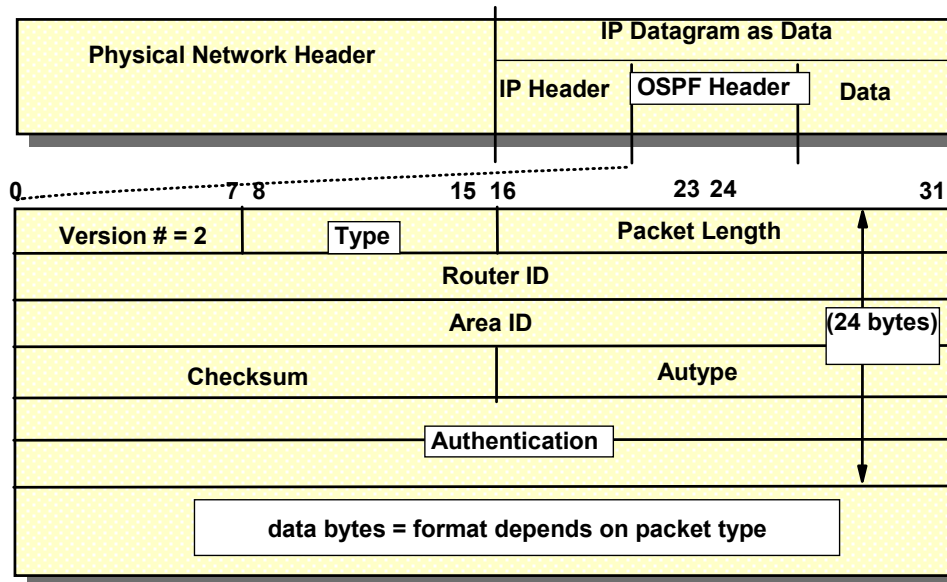
IP Datagram Format (OSPF)



© IBM Corporation 2002

1. This is the IP datagram format.
 1. The IBM manuals do not show the layout of IP datagrams and TCP segments since the layouts are located in the TCP/IP RFCs and also in TCP/IP tutorials available as IBM Redbooks or as other non-IBM publications.
 2. If you have worked with SNA in the past, this fact about TCP/IP documentation and formats will be new to you. SNA is an architecture owned and controlled by IBM. IBM publishes Formats and Protocols manuals as well as Data Areas Manuals to describe internals of SNA.
2. OSPF packets are sent as IP Datagrams with a Protocol Type of 89.

OSPF Packet Format

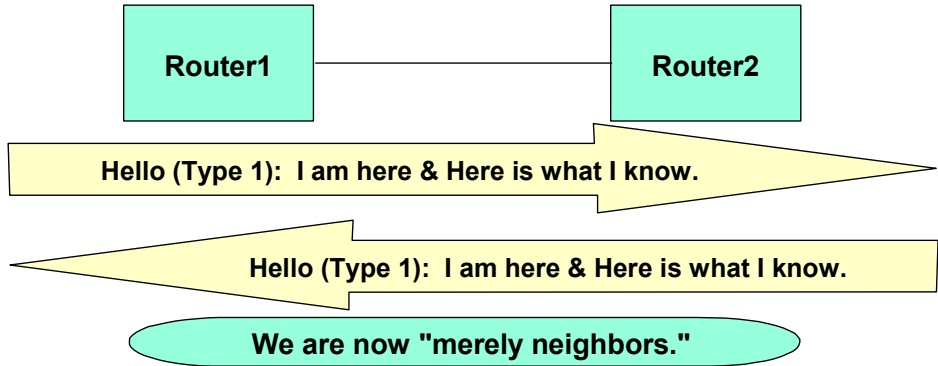


1. OSPF Version # is 2 (RFC 1583).
2. OSPF Packet Types
 1. Type 1 = Hello Packet
 2. Type 2 = DataBase Description Packet
 3. Type 3 = Link State Request Packets
 4. Type 4 = Link State Update Packets
 5. Type 5 = LSA Acknowledgement Packets
3. The packet length in bytes includes the OSPF header.
4. The Router ID is the the ID of the router originating the packet.
5. The Area ID is the area that the packet is being sent into.
6. The Checksum is the standard IP checksum of the entire contents of the packet, excludig the 64-bit authentication field.
7. The AuType identifies either 0 for no authentication or 1 for plain text 64-bit password or 2 for MD5 authentication.
8. Authentication is used by the authentication scheme.

Types of OSPF Packets

- Type 1 = Hello Packets
- Type 2 = DataBase Description Packets
- Type 3 = Link State Request Packets
- Type 4 = Link State Update Packets
- Type 5 = Link State Acknowledgement Packets

Discovering the Neighbors



OSPF Packet Header, Type 1		
Network Mask		
Hello Int.	Options	RTR Priority
Dead Router Interval		
Designated Router		
Backup Designated Router		
RouterIds of All Neighbors on This Medium		




1. The HELLO packet functions as a "keepalive" packet to determine whether a neighbor is still operating; it also sets up the initial negotiation of the connection, negotiating who the Designated Router is on the network.
 1. Connections between neighbors must maintain the same Hello_Interval in order to establish a neighbor relationship. They must also maintain the same Dead_Router_Interval.
 2. The Dead_Router_Interval is recommended to be at least 4x the value of the Hello Interval. If a neighbor has been unavailable for the Dead_Router_Interval, the neighbor relationship is marked as "down."
 3. These intervals have different defaults depending on the platform implementation. Console error messages at OS/390 indicate whether these values are out of synch.
2. The Options field indicates whether the router supports the TOS field or is capable of sending External Routes.
3. The Priority field indicates with a non-zero value that the router is eligible to be a Designated Router.

Discovering Neighbors & DR Eligibility


Non-Broadcast Network Multi-Access (NBMA)

- Neighbors must be pre-defined
 - X.25, Frame Relay, Native ATM
- Unicast to Interface IP Addresses: Hello (Type 1)
- Requires DR election



Broadcast Network

- Neighbors dynamically discovered
 - TR, Ethernet, FDDI, LANE, HiperSockets
- Multicast to 224.0.0.5: Hello (Type 1)
- Requires DR election



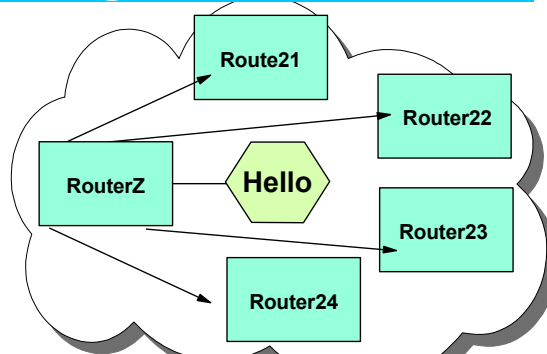
© IBM Corporation 2002

1. Technically, OSPF requires multicast capability on all interfaces over which it receives OSPF information. Otherwise, packets are discarded.
2. However, certain types of interfaces are incapable of multicast. Either they do not have the microcode to support it or the architecture of the medium does not include multicast capability. For such interfaces, OSPF will work with unicast if the network is defined as "non-broadcast=yes" and all DR neighbors and Non-DR neighbors are defined. In such instances, all devices on such a network must be defined as unicast even if some of them are multicast-capable. Since the definition burden for unicast can be onerous, the recommendation is to use multicast wherever possible.
3. A new type of Broadcast network was introduced in z/OS V1R2: HiperSockets, also known as iQDIO. It enables an internal connection within a CEC among z/OS images or among z/OS images together with LINUX images. You must ensure that the HiperSockets network or LAN has at least one Designated Router and preferably a Backup Designated Router.
4. This visual describes the Hello Packet from the perspective of only one router in the network. In fact, all the routers in the NBMA and Broadcast networks send each other Hello Packets.
5. Depending on the type of network in which the neighbors are participating, you may need to define who the eligible Designated Routers are inside your OSPF definitions. In CS/390, these definitions are maintained in the OMROUTE Configuration file.
 1. The Priority field of the HELLO packet indicates with a non-zero value that a router is eligible to become a DR.
6. The HELLO protocol determines who the Designated Router (DR) will be.
 1. Role of the DR:
 1. It is adjacent to all other routers on the network.
 2. It generates and floods the network links advertisements on behalf of the network. In a broadcast network this reduces the amount of router protocol traffic that is generated, as only the DR is responsible for flooding the network with the information.
 3. It is responsible for maintaining the network topology database that is replicated at all other routers on the network and in the same area.
 2. Role of the Backup DR:
 1. It takes over should the DR fail.
7. The router with the higher Router_Priority becomes the DR on a broadcast or non-broadcast multiaccess network. If there is a tie, the router with the higher Router_ID becomes the DR. For point-to-point links do not specify Router_Priority; it is irrelevant. If your OS/390 or z/OS system is to be used primarily for production or test work and not for routing, consider setting Router_Priority to 0 for all CS interfaces so that the CS node is ineligible to become the DR.
 1. There is a wait period in the election of the DR - so, if Router 12 and 13 have communicated and Router 11 and RouterY have communicated, nevertheless, only one of them would become the DR. This wait is the Dead_Router_Interval, during which time the router declares itself eligible for DR (with a non-zero router priority); however, until the expiration of the DR interval, no one router declares itself the DR. This wait is introduced so that as many eligible routers as possible have a chance to come up and announce themselves before a DR is selected.
8. If more than one router declares itself the Designated Router (DR), the one with the higher RouterID becomes the DR.
9. If a DR has been declared and another router with a higher RouterID joins the network, the original DR maintains its role as Designated Router.
10. There must be a Backup DR; the original election of the DR leaves the router with the next higher RouterID as the Backup DR.
11. If the Router ID is not specified, the IP address of one of the OSPF interfaces will be used as Router ID. For CS in OS/390 or z/OS, we recommend that you use the IP address of a STATIC VIPA or a physical interface for the RouterID to avoid the default selection of a Dynamic VIPA which could move.

Discovering Neighbors

Special Types of Point-to-Multipoint Networks



- Neighbors dynamically discovered through special signalling protocols
 - CS/390 MPC, XCF, IUTSAMEH
- Unicast to Each Interface: Hello (Type 1)



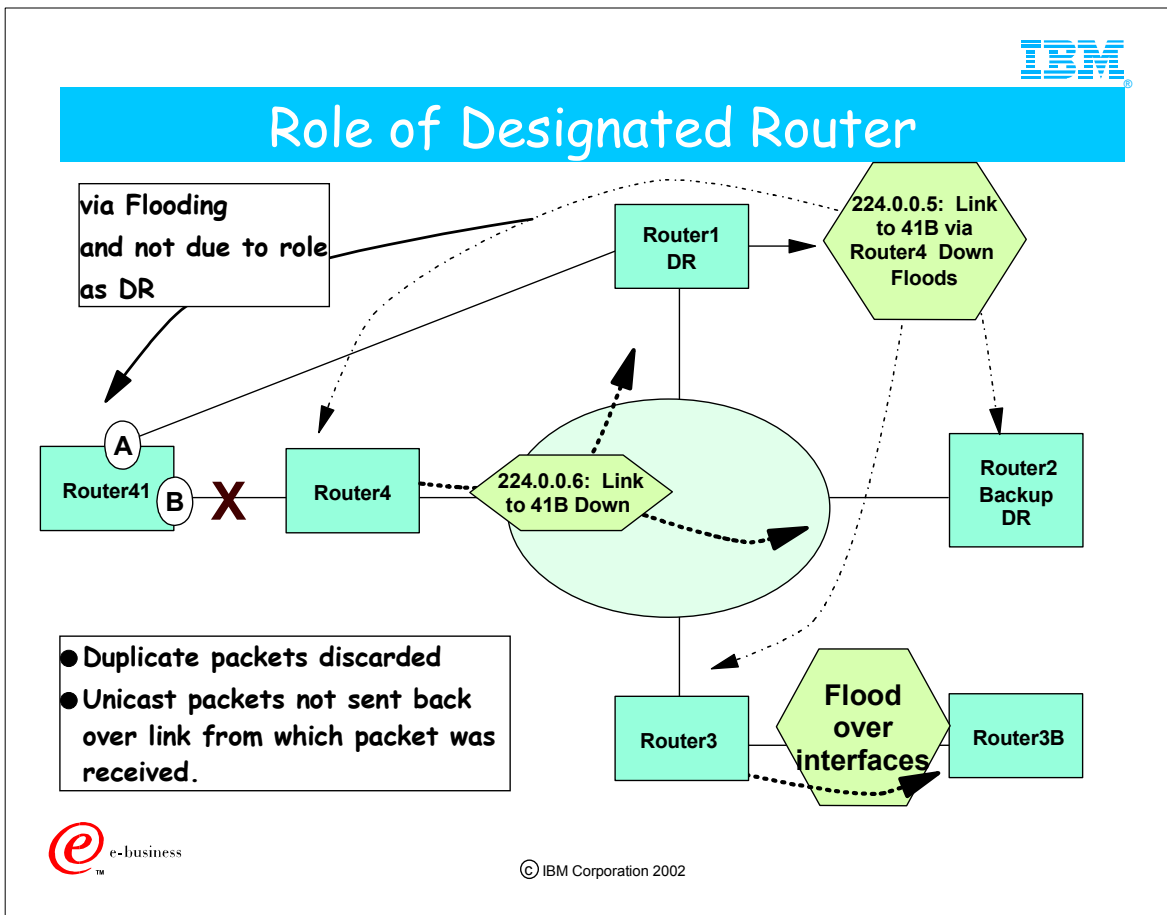
Point-to-Point Networks

- Neighbors dynamically discovered through HELLO protocol (CTCA, CLAW)
- Link-state Updates multicast to 224.0.0.5
- Link-state Update retransmissions multicast to station at end of connection: Hello (Type 1)

"D TCPIP,,OMPR,OSPF,IF"

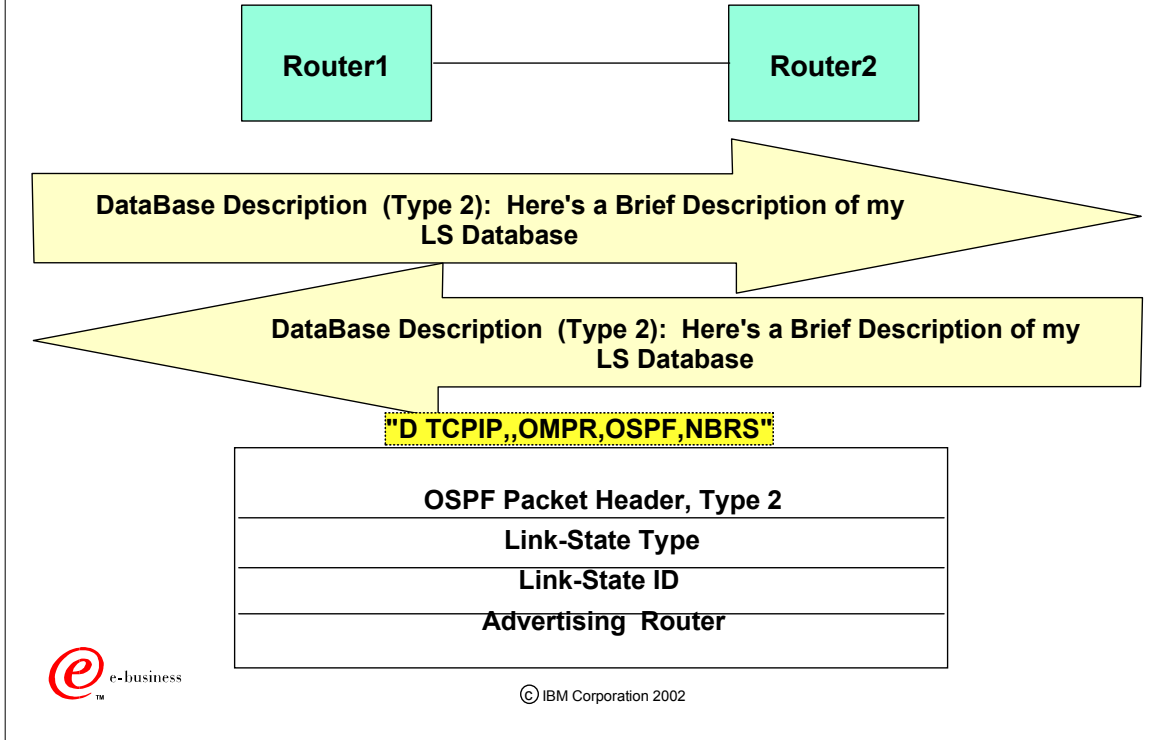


1. S/390 and z/Series with IP in CS/390 can dynamically discover its neighbors on an XCF or an IUTSAMEH network.
 1. OSPF definition is thus simplified, as we do not need to predefine the routers that are DR-eligible in the network in the OMPROUTE Configuration file.
2. This special type of Pt-to-Multipoint network does not require a DR, as we do not need to predefine the neighbors in the network in the OMPROUTE Configuration File.
3. CLAW and CTCA are point-to-point in CS for OS/390
 1. Cisco Routers describe themselves as point-to-multipoint when they are configured with the following interfaces to S/390:
 1. CLAW (NOTE: An APAR released in July of 2001 permits the definition of a CLAW interface as point-to-multipoint with a parameter on the LINK statement. See APAR PQ48766.)
 2. ESCON EMIF MPC+
4. An OMPROUTE command will display the type of network that is available to the connection you have defined:
 1. "D TCPIP,,OMPR,OSPF,IF" displays the interface types.
5. Networks that do not require DR election like Point-to-Point and Point-to-Multipoint networks may be set up as Demand Circuits and may use Hello-Suppression to minimize OSPF traffic: with Demand_Circuit=YES, LSAs are not periodically refreshed over the interface; only LSAs with real changes are advertised and LSAs flooded over this interface never age out.

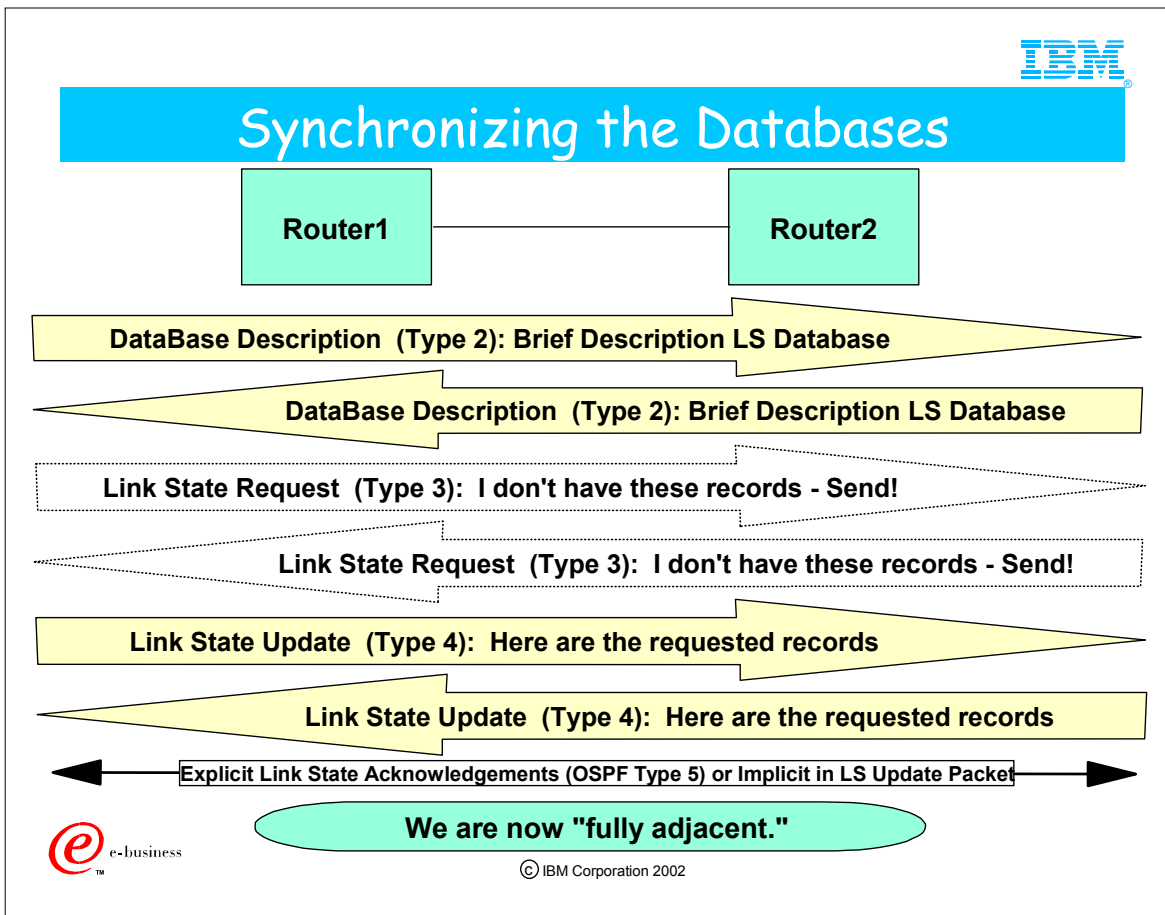


1. All routers synchronize their databases with that of the Designated Router.
 1. When a link-state advertisement is required due to a change, the router multicasts this LSA to the DR and its Backup DR using the special multicast address of 224.0.0.6 (for "All Designated Routers").
 1. If this is a non-broadcast network, the unicast address and not the multicast address is used
 2. The DR then floods this information to all its interfaces and to the network using the multicast address of 224.0.0.5 ("All OSPF Routers").
 1. If this is a non-broadcast network, the unicast address and not the multicast address is used. Furthermore, the advertisement is not sent back to the same interface from which it was received. (That is, Router 4 does not hear back from the DR that the link to 41 is down if it is attached to an NBMA network using unicast.)
2. If duplicate packets are received, as could be the case when there are point-to-point links in the network or alternate connections to the DR, or if there is router redundancy, the sequence number in the packet let the recipient know this is a duplicate. The recipient then discards the duplicate.
3. Routers that hear updates from their DR flood the information over their other interfaces to other attached routers. (See Router3-to-Router3B communication.)
 1. Each LSA must be acknowledged, either implicitly or explicitly with a Link State Acknowledgment packet (OSPF packet Type 5).
 1. An implicit acknowledgment can be perceived due to an update that was intended for the flooding router anyway.
 2. Router 1 floods the information to Router 41A, since this is a pt.-to-pt. connection. Router1 is not acting in the capacity of a DR here.

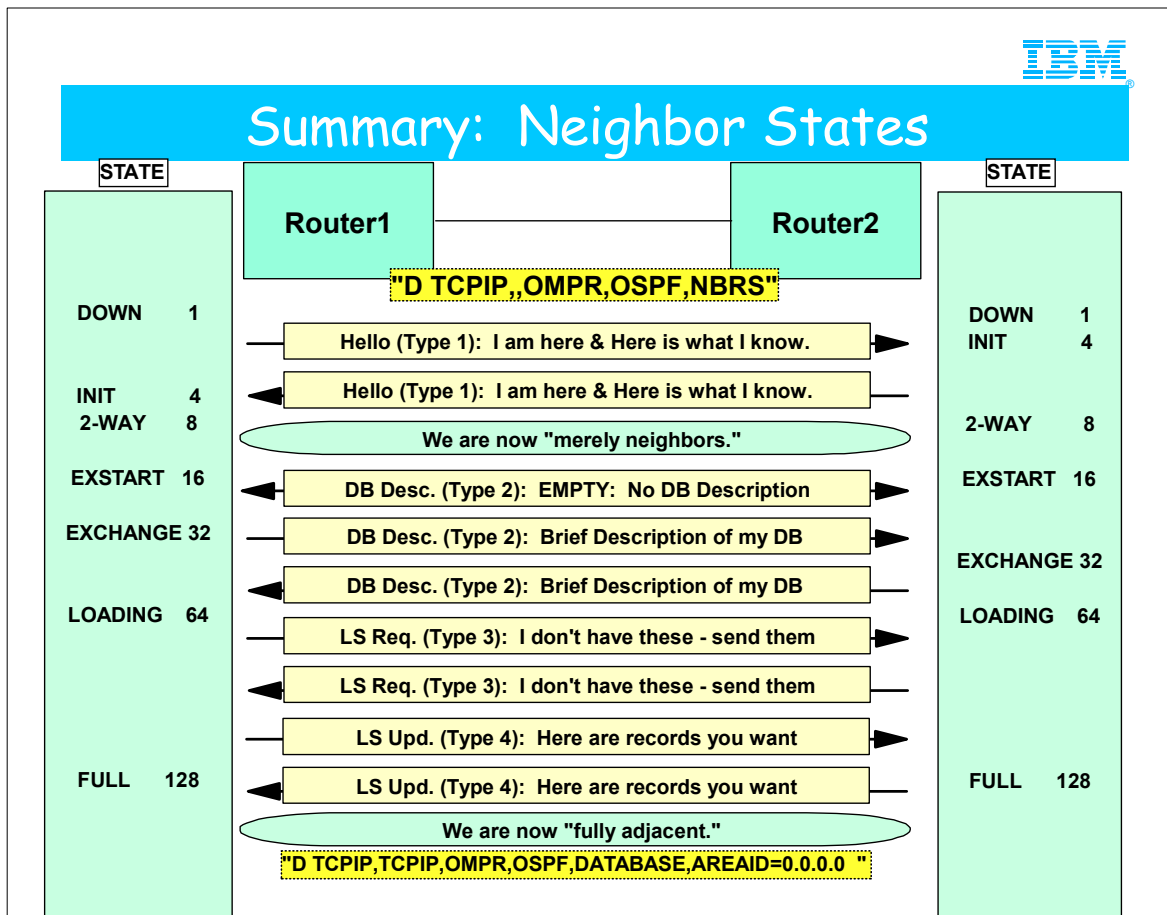
Forming Adjacencies: DataBase Exchange



- Once the HELLO packets have been exchanged there is two-way communication and the routers are "merely neighbors." They then know with whom they must establish adjacencies. (Remember: bringing up an adjacency means that the databases are synchronized.)
 - If the network is multiaccess, all routers are adjacent to the Designated and Backup Designated Routers. The adjacencies are only with the DR or Backup DR.
 - If the network is point-to-point, or special point-to-multipoint like MPC, XCF and IUTSAMEH, the router forms an adjacency with the partner at the other end of the connection.
- The routers next use the Exchange Protocol to exchange a description of their link state database (via OSPF Type 2 packets) with their adjacent partners. If they determine that they are fully synchronized, they are considered "fully adjacent."
- To display the status of adjacencies, use the OMROUTE command...
 - "D TCPIP, OMPR, OSPF, NBRS"

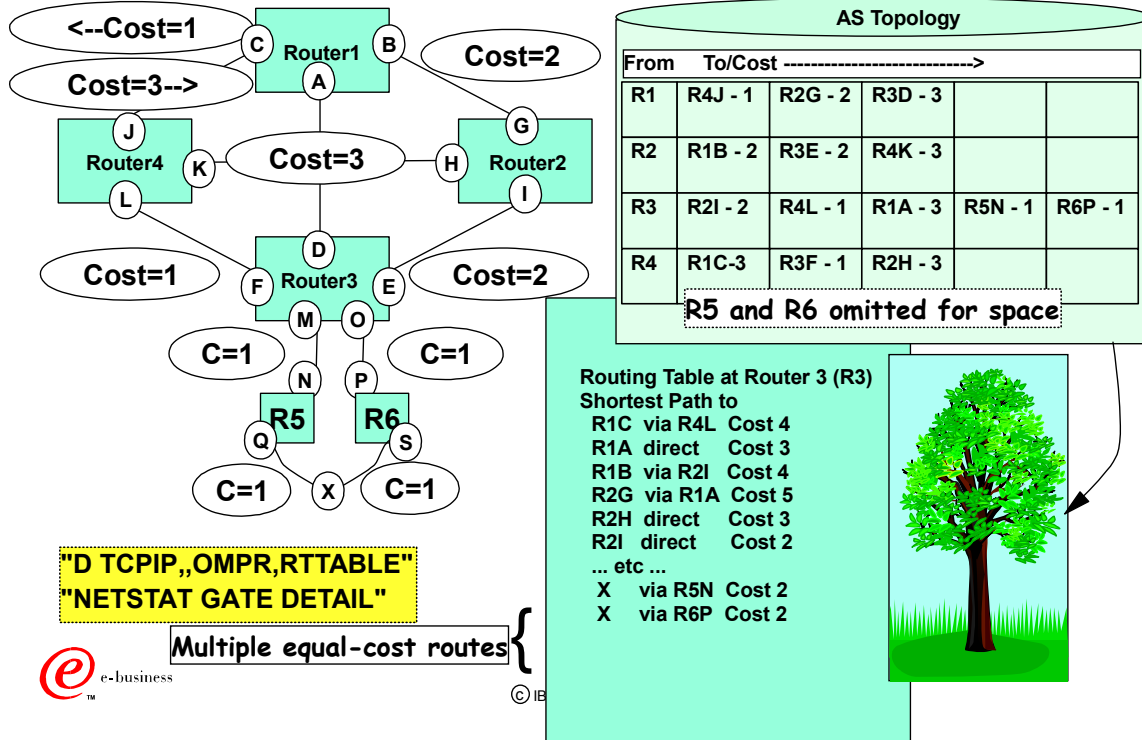


1. If the router databases are not fully synchronized, they request more information from the adjacent router (Link State Request/OSPF packet Type 3); then they exchange Link State Updates (OSPF Type 4 packets) to get in synch so that they may become "fully adjacent."
 1. The Exchange Protocol is more complicated than what is depicted here. One router (the one with the higher RouterID) assumes the master role and the other the slave role. The master sends its database descriptions, one at a time. The slave acknowledges each one and includes in the acknowledgement its database descriptions. If the new description indicates that this record is newer than what the recipient already has in its database, this description is saved as a record "of interest."
 1. The records are compared according to "type," "advertising router," and "link state ID." A sequence number in the record reveals whether the record is newer or older.
 2. Once all descriptions have been received, the neighbors send out database requests for more complete information about the records "of interest," which are then followed with the actual update packets. The Link State Update packets (OSPF Packet type 4) implement the "flooding" procedure.
 3. Each update packet must be acknowledged, either explicitly with a DataBase Acknowledgement packet (OSPF Type 5) or implicitly in the LS Update packets (OSPF Type 4).
2. It may not be until this point that the routers are "fully adjacent." This means that the link state databases are fully synchronized. Recall that whether or not a router interface is "fully adjacent" with every neighbor it knows about depends upon the network type to which the interface is attached. If the interface is part of a network that requires the election of a Designated Router and the router itself is neither the DR nor the Backup DR, it might have many neighbors (neighbor state 8), but it will be fully adjacent only with the Designated Router and the Backup Designated router (neighbor state 128). If the Database synchronization process is quite lengthy, you might consider setting the DB_Exchange_Interval higher than the default length (equivalent to the Dead_Router_Interval). If the database synchronization fails due to insufficient time, the neighbor process will never stabilize beyond Neighbor State of 8 (and may even continually cycle through states 8, 16, and 32 as indicated in Messages EZZ7919I and EZZ7921I, event 12), and the DB exchange and loading process will fail with event code 15: "failure to thrive" in Message EZZ7921I OSPF ADJACENCY FAILURE. DB_Exchange_Interval was introduced with APARs PQ45413, PQ37048, and PQ39760 for V2R7, V2R8 and V2R10 respectively.
3. Again, to display the status of adjacencies, use the OMPROUTE command...
4. "D TCPIP,,OMPR,OSPF,NBR5"



1. This page represents a summary of what we have seen in the protocol exchanges on the previous pages. It depicts the state of the conversation that is transpiring between neighboring routers in the OSPF network.
2. It also indicates the state changes that the routers (or CS nodes) go through during each phase of the protocol exchange.
3. Depending on the platform, the command displays for monitoring the routing in the network might show either the verbal status ("DOWN," INIT," etc.) or a number that can be equated with that state (DOWN=1, INIT=4, and so on).
4. 1 (DOWN): Initially the routers are down and have no contact with each other.
5. 2 (Attempt): In NonBroadcast MultiAccess (NBMA) networks a different state may be indicated even when a router is marked down. "Attempt" indicates that no contact has been made but the Hello packet will continue to send the packets to "attempt" to make contact.
6. 4 (Initialize): The Hello packet has been identified/received but no exchange of information has taken place between neighboring routers. Contact between the neighboring routers has been made, however.
7. 8 (2-Way): The Hello packets have been received and acknowledged by both neighboring routers, as indicated by the presence of the router itself in the neighbor's Hello packet. That is, each router sees itself in the neighbor's Hello packet. The designated router (DR) is selected and thereafter follows the selection of the Backup DR.
8. 16 (ExStart): Neighbor routers form adjacencies between themselves. Neighbor routers' communication is more advanced in this state and the routers decide who is the "master" and who is the "slave" and what is the initial DataBase Sequence Number. The transmission of the link state database can begin.
9. 32 (Exchange): The neighbors send their link state database to their adjacent routers. The link state database records describes the characteristics of the database and each must be acknowledged. Link state request packets requesting the neighbor's recent LSA's status may be sent to the neighboring routers. In this state the neighbors are capable of receiving and sending all types of OSPF routing protocol packets.
10. 64 (Loading): The link state request packets are being transmitted and received from neighboring routers requesting the most recent LSAs.
11. 128 (Full): If the neighboring routers are displayed with this state, it means that they are fully adjacent. The adjacent routers exchange LSAs and appear in their neighbors' router-LSAs.
12. The synchronization of databases can be verified with a command to display the OSPF database; there is a field called CHECKSUM TOTAL, which can be used to compare if 2 routers have synched their DBs.
 1. D TCPIP, TCPIP, OMPR, OSPF, DATABASE, AREAID=0.0.0.0

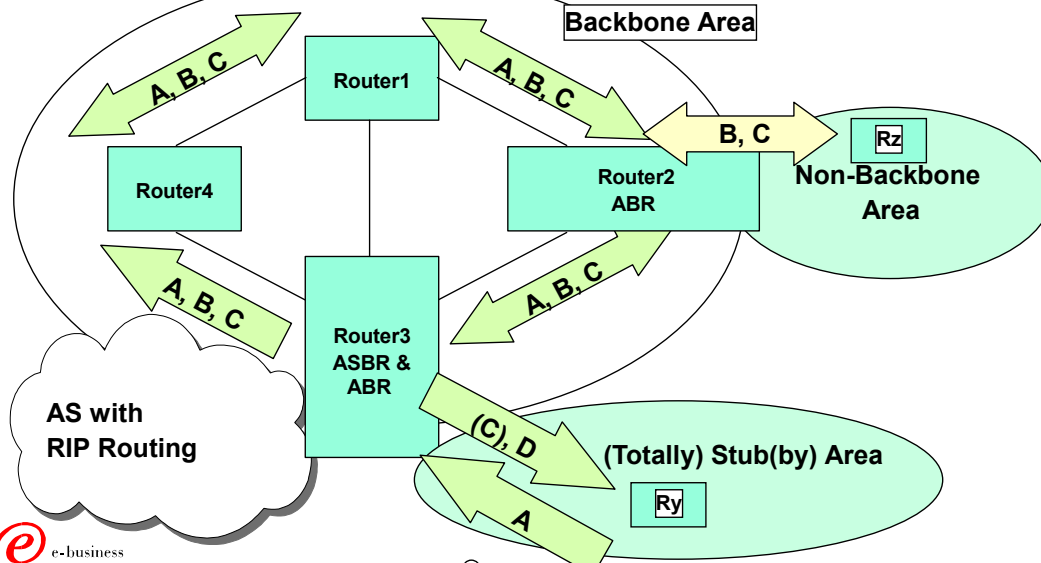
Computing the Routing Table



- All routers in the area have identical copies of the AS topology database.
- Each router computes its own routing table using a "spanning tree" algorithm.
 - This is also called the "Dijkstra" algorithm for computing "Shortest Path Flrst."
 - Every time a link-state update is received the entire tree is recomputed.
 - Christian Huitema quotes a time of 200 milliseconds to compute the routes for a network with 200 nodes. (p. 107 of "Routing in the Internet," Prentice Hall.
- If there are multiple equal-cost routes, the routing tree retains them.
 - If there are multiple routes to the same destination, the routing tree selects only the shortest path for retention.
- D TCPIP,OMPR,RTTABLE shows the cost of the routes in the network. TSO NETSTAT GATE DETAIL also shows the costs of the links to each gateway attached to the router.
- The RFC provides for the OSPF protocol to compute a different routing tree for each type of service (represented by the TOS bits - bits 3, 4, and 5 - of the TOS Byte in the IP Header). However, no known implementation of OSPF provides for multiple trees.

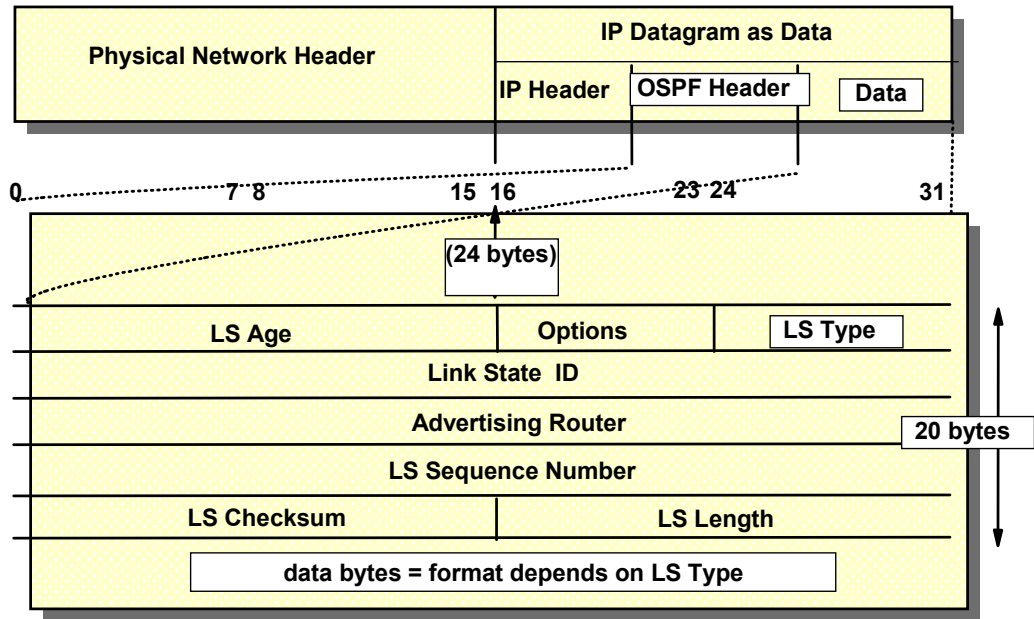
Advertising Link States: Flooding Review

A = Flooding Internal OSPF Routing Info
B = Flooding External Routing Info
C = Flooding Summary OSPF Routing Info
D = Flooding Default Routing Info



1. You have seen this visual before. It is just a reminder that there are several types of LSAs that can be sent into the network. These LSA Types are summarized on the next pages.
2. LSAs can be updated only every 5 seconds. If an LSA with a different sequence number arrives within 5 seconds of a previous LSA for the same route, it is ignored.
3. If the age of the LSA in the DB is less than 1 second, another LSA for the same route cannot be accepted.
4. A router is required to acknowledge LSAs, even if they are duplicates. However, it forwards only one copy of the LSA if it has received duplicates.
5. Every 30 minutes the Designated Routers must exchange their link state databases.

LSA Header Format



1. OSPF in OMPROUTE supports five types of LSAs.
2. An LSA Type 7 is described in the RFCs. It is used to advertise a Not So Stubby Area (NSSA). OMPROUTE does not support this type of LSA.

Types of Link State Advertisements

- Type 1 = Router Links
 - Describes all router's interfaces into an area and their states (Intra-area destinations)
- Type 2 = Network Links
 - Describes DR connection to medium in broadcast or NBMA nets
- Type 3 = Summary Links, created by Area Border Routers
 - Describe inter-area routes
- Type 4 = Summary Links, created by Area Border Routers
 - Describe routes to the AS Boundary Routers
- Type 5 = AS External Links, created by AS Boundary Routers
 - Describe routes to destinations outside the AS



```
"D TCPIP,,OMPR,OSPF,DATABASE,AREA=..."
```


```
"D TCPIP,,OMPR,OSPF,DBSIZE"
```

```
"D TCPIP,,OMPR,OSPF,LSA,LSTYPE=..."
```

© IBM Corporation 2002

1. More about Type 1: Router LSA
 1. Indicates whether originating router has Virtual Links, is an ASBR or an ABR
 2. Indicates for each link whether it is pt-to-pt, connection to a transit network, connection to a Stub area, or Virtual Link.
2. More about Type 2: Network LSA
 1. Originated by DR in broadcast and NBMA networks
 2. Contains list of all routers that are attached (adjacent) to the DR
3. More about Type 3: Summary LSA
 1. Originated by ABR
 2. Contains IP Destination Address inside the area
 3. One advertisement per destination
 4. Network Mask is mask of IP Address Destination
 5. Used to describe Default Route to Stub Area
 1. LS Id Field and Network Mask both set to 0.0.0.0
4. More about Type 4: Summary LSA about ASBR
 1. Originated by A=ABR
 2. Contains IP Address of ASBR
 3. Network Mask is 0.
5. More about Type 5: External LSAs
 1. Originated by ASBRs
 2. Describe Destinations outside the AS
 1. Can be default destination
 1. LS ID Field is 0.0.0.0
 2. Can be network destination
 1. LS ID Field contains network address
 3. Identifies whether to use Type 1 metric for destination
 1. Metric is comparable to a link-state metric
 4. Or... Identifies whether to use Type 2 metric
 1. Metric is greater than link-state metric
6. Notes on Type 7 LSAs: Not So Stubby Area
 1. This optional OSPF feature is described in IETF RFC 1587 and in Internet Draft draft-ietf-ospf-nssa-update-10.tx. NSSAs are similar to the existing OSPF Stub Area. However, unlike an OSPF Stub Area, they may import AS external routes in a limited fashion.
7. The following OMPROUTE commands display information about the LSAs that have reached this router:
 1. D TCPIP,,OMPR,,OSPF,DATABASE,AREA=...
 2. D TCPIP,,OMPR,OSPF,DBSIZE
 3. D TCPIP,,OMPR,OSPF,LSA,LSTYPE=...

Aging the Link State Records



"D TCPIP,,OMPR,OSPF,RTTABLE"

- A new record has an Age of "0."
- The record age is incremented by the Transmission Delay (D value in diagram) that is configured for the interface it is sent over.
- It is also incremented for every second it is maintained.
- Maximum Age of a record is one hour, at which point routers will no longer use the record in their route table calculation. The record will eventually be removed from the database.
- A router must retransmit a record *at least* once every 30 minutes.

1. D TCPIP,,OMPR,RTTABLE displays the age of information about connections.



For More Information ...

URL	Content
http://www.ibm.com/servers/eserver/zseries	IBM Enterprise Servers (z900 & S/390)
http://www.ibm.com/servers/eserver/zseries/networking	z900 Networking
http://www.ibm.com/servers/eserver/zseries/networking/technology.html	Networking White Papers and Information
http://www.ibm.com/software/network	Networking & Communications Software
http://www.ibm.com/software/network/commserver	Communications Server
http://www.ibm.com/software/network/commserver/library	CS White Papers, Product Doc, etc.
http://www.redbooks.ibm.com	ITSO Redbooks
http://www.rfc-editor.org/rfcsearch.html	RFCs
http://www.ibm.com/support/techdocs/	Advanced Technical Support (Flashes, Presentations, White Papers, etc.)

Whitepaper on OSPF with IBM & CISCO:

"OSPF Design and Interoperability Recommendations for Catalyst 6500 and OSA-Express Environments" http://www-1.ibm.com/servers/eserver/zseries/networking/pdf/ospf_design.pdf

Redbook on "Networking with z/OS and Cisco Routers: An Interoperability Guide (SG24-6297)



© IBM Corporation 2002

Bibliography

RFCs 1583 & 2328	OSPF V2 by John Moy: URL http://www.ietf.cnri.reston.va
GG24-3376	TCP/IP Tutorial and Technical Reference
SC31-8513	OS/390 Communications Server: IP Configuration
Huitema, Christian	Routing in the Internet, Prentice Hall
Black, Uyless	IP Routing Protocols: RIP, OSPF, BGP, PNNI & CISCO Routing Protocols (SR23-9498)
Parkhurst, William R.	Cisco Router OSPF Design and Implementation Guide, Parkhurst McGraw-Hill (SR23-8683)
Cisco Systems	CCIE Fundamentals: Network Design and Case Studies, 2nd Edition (SR23-9437)
Doyle, Jeff (Cisco Systems)	CCIE Professional Development: Routing TCP/IP, Volume 1 (SR23-9241)
SG24-5631	SecureWay Communications Server for OS/390 V2R8 TCP/IP Guide to Enhancements
SG24-5227	V2Rn TCP/IP Implementation Guide, Vol. 1 (Configuration and Routing)
SG24-5228	V2Rn TCP/IP Implementation Guide, Vol. 2 (UNIX Apps)
SG24-5229	V2Rn TCP/IP Implementation Guide, Vol. 3 (MVS Apps)
FLASH N3196	TCP/IP Migration Tips & Hints (V2R8) at www.ibm.com/support/techdocs
FLASH N3190	OMPROUTE and VIPA at www.ibm.com/support/techdocs





e-business™

© IBM Corporation 2002



e-business™

© IBM Corporation 2002