



Lucent Technologies
Remote Access Business Unit

OSPF and the Internet



Table of Contents

1.0 OSPF and the Internet	1
Why OSPF?	1
When to Use OSPF	1
Lucent Advantages	1
Easy to Use, Easy to Configure, Focused on PortMaster Applications	1
Streamlined to Maximize System Resources	2
NSSA Support	2
OSPF and RIP Background	2
OSPF vs. RIP	2
RIP Update Process	2
The OSPF Update Process	3
Key OSPF Concepts	4
OSPF Router Relationships	4
Neighbors and Adjacencies	5
OSPF Backbones	5
Stub Areas	5
NSSA Areas	5
Variable-Length Subnet Masks (VLSMs) with OSPF	5
OSPF "Costing"	6
OSPF Packet Authentication	6
OSPF and Lucent: Typical Applications	7
Scenario #1: VLSM, Subnetting a Class C Address across Multiple Sites	7
Scenario #2: VLSM, Subnetting a Class C Address across and ISP POP	7
2.0 Summary	8
3.0 Bibliography	8

1.0 OSPF and the Internet

This paper provides background information and an application guide to the Open Shortest Path First (OSPF) routing protocol. It highlights the features and benefits of OSPF, describes Lucent Technologies OSPF implementation, explains the difference between OSPF and the Routing Information Protocol (RIP), and describes typical OSPF applications. It is meant for managers and network administrators at Internet Service Providers (ISPs), corporations, and other organizations who want to identify the advantages that Lucent Technologies OSPF can provide their networks. All topics in this paper apply to routing on both the Internet and TCP/IP based private inter-networks.

Why OSPF?

OSPF offers all the functionality of RIP, plus:

- Variable-length subnet mask (VLSM) support
- Routing updates without the 30-second "hold-down" period required by RIP
- Bandwidth optimization, including less frequent routing updates and a choice of metrics for defining the best links between routers
- Up to 255 routed segments between routers
- Packet authentication of routing updates with both simple password and MD5 authentication

When to Use OSPF

The following are typical scenarios for using OSPF:

When a single router or communications server must accommodate different sized TCP/IP networks. Increasingly, ISPs need to divide or combine subnets to ensure the most efficient use of TCP/IP addresses. This capability, called variable-length subnet masks (VLSM) or "classless" networking, is supported by OSPF. In contrast, RIP does not allow a network to be segmented or combined with others to create networks of different sizes.

When routing changes need to be propagated quickly. RIP can create too much network downtime by taking too long to update routers with network changes; RIP needs a hold-down period to ensure that information it has generated has been

properly propagated through the network. If a network has many routers, RIP updates can take several minutes to alert the entire network to the failure of a single router. OSPF updates are much faster than RIP updates. (Note also that sites using "one way out" or default gateways usually are much faster than sites using RIP.)

When more than 15 hops between routers are required. More than 15 hops might be a requirement in some larger networks. RIP will only support 15 hops between routers, but OSPF can support up to 255 hops.

When routing advertisements need to be password-protected to prevent network instability or sabotage. OSPF has packet authentication capability; RIP does not.

Lucent Advantages

Lucent Technologies OSPF complies with RFC-1583 with additional support for MD5 Authentication and RFC-1587 (NSSA). OSPF is implemented on all Lucent Technologies remote access servers and routers, including PortMaster® 2 Communications Servers, PortMaster 3 Integrated Access Servers, PortMaster 4 Integrated Access Concentrators, IRX™ Routers, and Office Routers. The following points demonstrate the key advantages of the Lucent implementation.

Easy to Use, Easy to Configure, Focused on PortMaster Applications

For the typical PortMaster application, only a few commands are required to run OSPF. You need not set up the complex redistribution and filtering schemes required for some other vendor OSPF implementations. Figure 1 lists the commands required for typical Lucent OSPF setups.

```
Command> set ospf enable
Command> set ospf priority 1
Command> add ospf area 0.0.0.0
Command> set ether0 ospf on
```

Figure 1. Command strings required for typical Lucent OSPF setup.

Streamlined to Maximize System Resources

When enabled, Lucent Technologies compact OSPF software image requires a minimum of RAM. Additionally, the OSPF code is designed to load into RAM only if OSPF routing is enabled. This feature makes more system resources available to PortMasters not running OSPF and allows you to run OSPF later without needing to change software.

NSSA support

Lucent supports "not so stubby networks" (NSSA), enabling Lucent Technologies OSPF routing tables and link state databases to take up less memory. Because many other remote access vendors do not support NSSA, Lucent Technologies NSSA support makes its communications servers and routers uniquely interoperable with Internet backbone routers that do support NSSA.

OSPF and RIP Backgrounder

Routing protocols define the rules that routers use to communicate with each other. Routing protocols dynamically provide the network topology information necessary to choose paths amongst routers, allowing routers to automatically choose routes, and to alter them when network changes occur. Beyond these basics, routing protocols vary greatly in design, capability, implementation, and impact on network infrastructure.

The most widely implemented routing protocol is the Routing Information Protocol (RIP). RIP was the first common TCP/IP routing protocol and is supported by most routers. RIP became a component of TCP/IP when it was included with Berkeley Standard Distribution (BSD) UNIX in 1982. Even though RIP has many limitations, RIP's simplicity and interoperability have spurred its implementation in TCP/IP networks worldwide.

In today's complex internetworking environments, especially on the Internet, RIP's limitations have become most apparent. RIP does not scale well to larger networks, consuming large amounts of network bandwidth. Also, RIP lacks several key features that can make today's networks much more responsive and flexible.

The OSPF routing protocol was developed to overcome many of the limitations of RIP. Although the current version of OSPF was

first formalized in 1991, OSPF has become more widely deployed only recently. Larger ISPs and corporations alike are beginning to require the broad feature set offered by OSPF. In contrast to RIP, OSPF scales to larger networks. It's faster, generally places much less strain on the network, optimizes throughput, and adapts more easily to existing internetworking needs.

OSPF vs. RIP

The fundamental difference between OSPF and RIP is that they are based on two different algorithms. OSPF is based on the Dijkstra link-state algorithm. RIP is based on the Bellman-Ford distance-vector algorithm. Using OSPF's link-state algorithm, every router maintains a similar network map identifying all links between neighbors. Best paths are calculated from these maps. OSPF also ensures that updates sent to neighboring routers are acknowledged by neighbors, verifying that all routers have consistent network maps. Using RIP's distance-vector algorithm, every router creates a unique routing table identifying the best path from itself to all other routers in the network.

Of the two protocols, OSPF's acknowledgment-oriented routing update process is far more responsive to changes in network topology. Routers can make decisions faster when their network information is known to be consistent with that of other routers.

RIP Update Process

Generally, RIP routers send updates to their neighbors every 30 seconds. These routing updates carry information about the number of hops between routers. Routers revise their routing tables with the network topology status by taking the update information from a neighboring router and adding another hop to the information received from that router. Figure 2 shows a four-router network with a fifth router (router E) added.

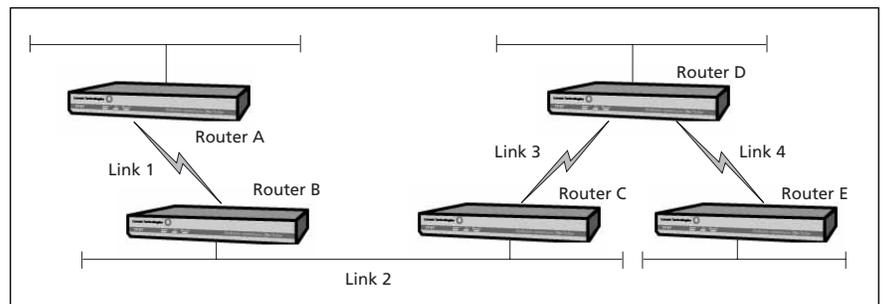


Figure 2. Four-router network with Router E added

Router E is to be added to an existing four router RIP network consisting of routers A, B, C, and D. Each has a unique routing table that identifies the appropriate path to take when forwarding packets. Before router E is added, router A's routing table is as shown in Table 1.

<i>Destination</i>	<i>Gateway</i>	<i>Metric (hop count)</i>	<i>Interface</i>
<i>A</i>	<i>A</i>	<i>0</i>	<i>Local</i>
<i>B</i>	<i>B</i>	<i>1</i>	<i>1</i>
<i>C</i>	<i>B</i>	<i>2</i>	<i>1</i>
<i>D</i>	<i>B</i>	<i>3</i>	<i>1</i>

Table 1. Router A's routing table before Router E is added.

- "Destination" is the TCP/IP address list of the routers to which router A has access.
- "Gateway" is the TCP/IP address list of the router through which traffic flows to reach chosen destinations.
- "Metric" is the number of links (or "hops") between the source and destination router.
- "Interface" is the source router port to be used for the route.

Once router E has been added, router E sends out notification of its location to router D over link 4. Router D updates its routing table with this new information. Within 30 seconds, router D forwards its new routing table in an update to router C over link 3, within 30 seconds, router C forwards its routing table update to router B over link 2, and so on. Ultimately router A's routing table will include another entry showing access to router E through router B, with a metric of 4, through interface 1.

When multiple paths exist between routers, hop counts are used to identify the optimal routing path—the one with the lowest cost (the lowest number of hops). For example, if a direct connection were established between routers A and E, router A's new entry would show access to router E over link 5 with a hop count of 1, replacing any less optimal router A-to-E entries.

The OSPF Update Process

In contrast to RIP, OSPF does not repeatedly broadcast routing tables to others and incrementally update hop counts. With OSPF, each router maintains a complete network map of the local area and sends updates and update acknowledgments when network changes occur or on 30 minute refresh cycles. OSPF sends only the minimum data required to communicate a change. This approach contrasts with RIP, where every router has a unique routing table tailored to its specific place in the network. In an OSPF network, every router within an area contains the same routing table information in the form of a network map. As shown in Figure 2, router E is added to an existing four router OSPF network consisting of routers A, B, C, and D. All possess the same network map showing all routers in the

network and their direct links to other routers. Before E is added, router A's topology database is as shown in Table 2.

<i>From</i>	<i>To</i>	<i>Total Cost</i>
<i>A</i>	<i>Link 1</i>	<i>1</i>
<i>Link 1</i>	<i>B</i>	<i>1</i>
<i>B</i>	<i>Link 2</i>	<i>2</i>
<i>Link 2</i>	<i>C</i>	<i>2</i>
<i>C</i>	<i>Link 3</i>	<i>3</i>
<i>Link 3</i>	<i>D</i>	<i>3</i>

Table 2. Router A's topology database before Router E is added.

Once router E is added, router E sends out notification (called a "link state advertisement") of its location to router D. Router D updates its network map and immediately forwards E's update message to router C, which immediately forwards E's update message to router B, and so on. Ultimately router A's routing table will include another entry showing that router D has access to router E over Link 4 with a cost (to router A) of 4. Indeed, the

same advertisement generated by router E makes its way to router A.

OSPF's update process affords three benefits over RIP's:

1. OSPF routing updates take place less often, every 30 minutes or when network changes occur. Thus, OSPF optimizes network bandwidth by keeping the frequency of update traffic to a minimum.
2. OSPF updates typically propagate throughout the network more rapidly than do RIP updates, enabling OSPF networks to recover more rapidly from broken links.
3. OSPF does not have RIP's 15-hop-count limitation. As a result, OSPF can accommodate many more routed network segments.

Key OSPF Concepts

The key OSPF concepts you need to understand to properly design an OSPF network are as follows:

- OSPF router relationships including autonomous systems, neighbors and adjacencies, backbones, and stub areas
- Variable-length subnet masks with OSPF
- OSPF "costing"
- OSPF packet authentication

OSPF Router Relationships

The concept of the OSPF area is a fundamental part of OSPF network design. OSPF is a CPU-intensive protocol, and unlike RIP networks OSPF networks are not bound by a hop count limitation. Very large OSPF networks can experience routing and update traffic problems that seriously impact network performance. In addition, routers in large OSPF networks require large amounts of memory. To avoid these problems, OSPF networks can be divided into more manageable OSPF "areas."

OSPF areas are made up of "internal routers" and are linked to other areas by "area border routers" (ABRs). Supersets of OSPF areas are called "autonomous systems" (AS), which are linked to other autonomous systems by "autonomous system border routers" (ASBR). OSPF autonomous systems

can be interlinked by an exterior gateway protocol such as the Border Gateway Protocol (BGP).^{*} All OSPF routers must be capable of acting as internal routers, area border routers, or autonomous system border routers. Figure 3 illustrates these concepts.

**For information about the Lucent Remote Access Business Unit's implementation of BGP, refer to our PortMaster Routing Guide. A copy of this guide can be downloaded from our web page at the following URL: <http://www.livingston.com/tech/docs/routing/about.fm.html>.*

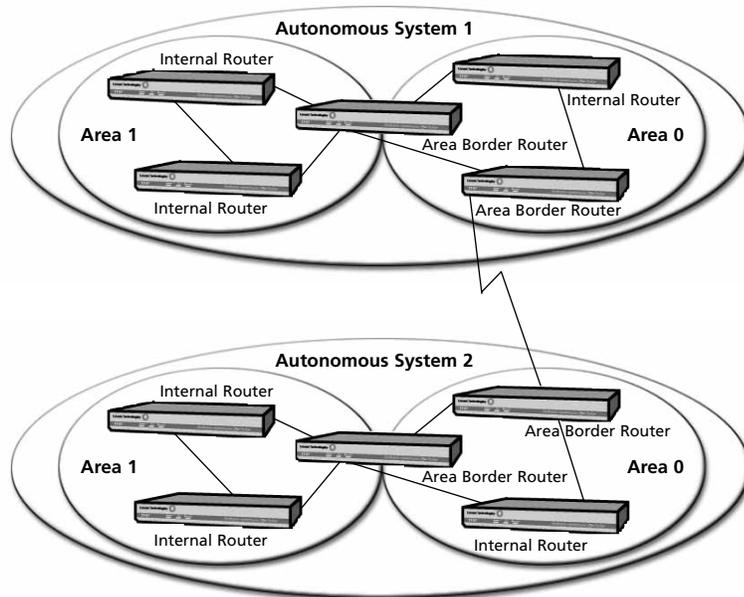


Figure 3. OSPF autonomous systems and routers

By grouping subnets into areas and areas into autonomous systems, network designers can create more efficient and manageable OSPF networks. Routers within an area need only maintain network maps for their respective area. This feature minimizes routing updates from other areas and conserves router memory. The autonomous system concept further conserves system and router resources by minimizing the flow of routing updates and decreasing the resources required to keep track of these updates.

Because traffic patterns and links vary by network, there is no definitive rule for the size and makeup of an OSPF area. Nevertheless, a general rule of thumb is to limit areas to no more than 40 or 50 routers to ensure adequate OSPF network performance.

Neighbors and Adjacencies

Neighbors and adjacencies are relationships established among OSPF routers within an area for intra-area router communications. Neighbors are routers that share a common network segment and area. Neighbors are created by OSPF's "hello" protocol. Small hello packets are frequently sent to verify two-way communication between neighboring routers. These periodic hello packets are a much more bandwidth efficient method for verifying connectivity than are the full network table refreshes performed by RIP.

Adjacencies are created when neighboring routers exchange routing information. To minimize update information on a segment, OSPF creates a designated router (as well as a backup designated router) to act as the central point for routing table updates. All routers in a segment keep up-to-date tables but exchange routing information with only the designated routers. Adjacent routers free up network resources by centralizing the routing table update process, limiting the update information traffic between neighbors. In addition, OSPF can optimize router CPU usage by allowing any router to act as the designated router, allowing routers with more available resources to be chosen to administer this activity.

OSPF Backbones

Any OSPF network containing more than one area requires an area numbered as "0," which is called the "backbone." All areas in an autonomous system must be connected to the backbone. The backbone is not necessarily made up of additional routers or hosts, but instead can be viewed as a logical routing construct created to manage inter-area traffic. In some cases, backbones can consist solely of routers belonging to other areas. To free up backbone resources for routing issues, hosts should be located in areas other than the backbone. Generally, if you are going to design an OSPF autonomous system with only one area, you should use Area 0.

Stub Areas

"Stub areas" are recommended in OSPF areas that are connected to other areas through one or more area border routers (ABRs). Stub areas cannot support autonomous system border routers (ASBRs). A likely stub area location would be an OSPF

remote office with a single point of access to a central office (CO). Routing out of stub areas is based on default routes—fixed, predefined routing paths. Stub areas are beneficial because their routers require less memory and generally create less network overhead.

NSSA Areas

NSSA areas also enhance the use of network resources. NSSA should be used when an area has one or more ASBRs but attaches through ABRs.

Variable-Length Subnet Masks (VLSMs) with OSPF

As TCP/IP network addresses become more scarce, organizations are assigning only the number of TCP/IP addresses required for a given network. Unfortunately, limitations within the RIP protocol have severely restricted the ability of organizations to assign TCP/IP addresses.

For the RIP protocol to route information properly between separate subnetworks or "subnets," every subnet must have the same subnet mask and the networks must be contiguous. This limitation has especially serious consequences for multiport communication devices routing traffic among many networks. Whether a given network connected to the communication device needs 6 or 126 addresses, each attached network must be assigned the same number of IP addresses. Therefore, RIP can be a very wasteful protocol for organizations such as ISPs and corporate central sites that need to assign subnets with different network masks or that communicate among noncontiguous networks.

OSPF is not saddled by this RIP limitation because OSPF updates include network mask information. Armed with this information, OSPF enables a single multiport router to work with different subnet masks and noncontiguous networks. This capability allows much more efficient use of TCP/IP addresses, thereby allowing network designers greater freedom in assigning addresses. The ability to work with different network masks and noncontiguous networks is called "variable-length subnet mask" (VLSM) support.

Figure 4 illustrates the importance of VLSM. If router A is trading RIP updates with routers B and

C, router A is unable to distinguish between router B's and C's networks because router A does not know the network masks of routers B and C. Routers B and C send RIP update information to router A. This update information makes them both appear to be part of network 192.168.3.x (x=0-255). When data is directed from router A to a 192.168.3.x IP address, router A sends the packet to whichever router last provided a RIP update, making that last router appear to be the gateway for all network 192.168.3.x addresses. Hence, RIP cannot provide reliable routing in this network.

1. Intra-area routes stay within a single area.
2. Inter-area routes extend within the autonomous system, crossing area border routers (ABRs).
3. Type 1 External routes are learned from outside the autonomous system and have OSPF-like metrics.
4. Type 2 External routes are learned from outside the autonomous system and have non-OSPF-like metrics.

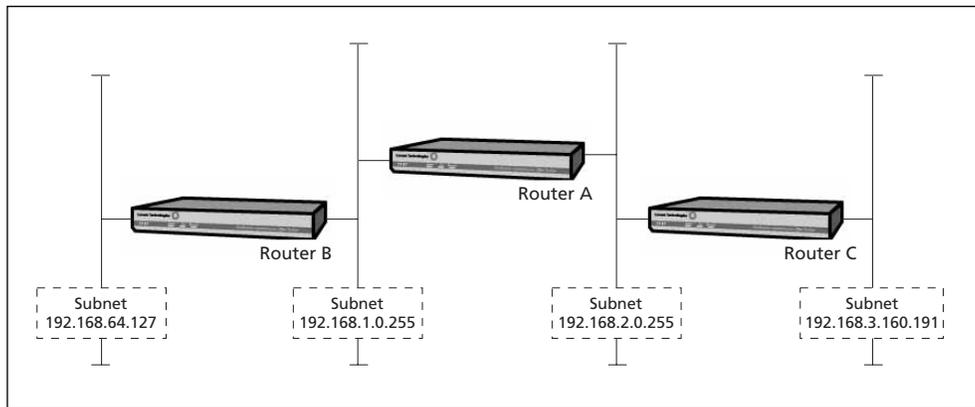


Figure 4. The need for variable-length subnet mask (VLSM) support

In contrast, OSPF does provide reliable routing. In an OSPF network, router A has both the IP address and network mask information required to identify the unique set of addresses associated with router B's and C's networks. Router B's OSPF update states that it is 192.168.2.65 and that its network mask is 255.255.255.192. With this information, router A can forward to router B any IP traffic sent to 192.168.2.64 through .127. Router C's address of 192.168.3.161 and network mask of 255.255.255.224 ensure that router A can properly forward all traffic sent to 192.168.3.160 through.191.

OSPF "Costing"

OSPF uses an hierarchy of routing categories and bandwidth calculations to choose optimal routing paths. Optimal routes are chosen on a least-cost basis. OSPF places routes into four categories, presented as follows in order of their OSPF cost with the lowest cost categories first:

OSPF chooses intra-area routes over inter-area routes, inter-area routes over Type 1 External routes, and so on. If multiple routes from within a given category are available, OSPF generally defaults to the route that offers the greatest bandwidth. Although OSPF allows the customization of routing cost metrics, in practice most OSPF networks base routing decisions on default bandwidth metrics.

OSPF Packet Authentication

All OSPF packets include authentication information. OSPF network routers can be protected against unauthorized routing information through the assignment of networkwide passwords. This protection can be useful, for example, in a case where two independent OSPF networks share the same cable. Passwords can keep networks more stable by protecting against unintentional or spurious routing updates and against intentional router sabotage.

OSPF and Lucent: Typical Applications

The following two scenarios are examples of popular applications for using OSPF with Lucent products.

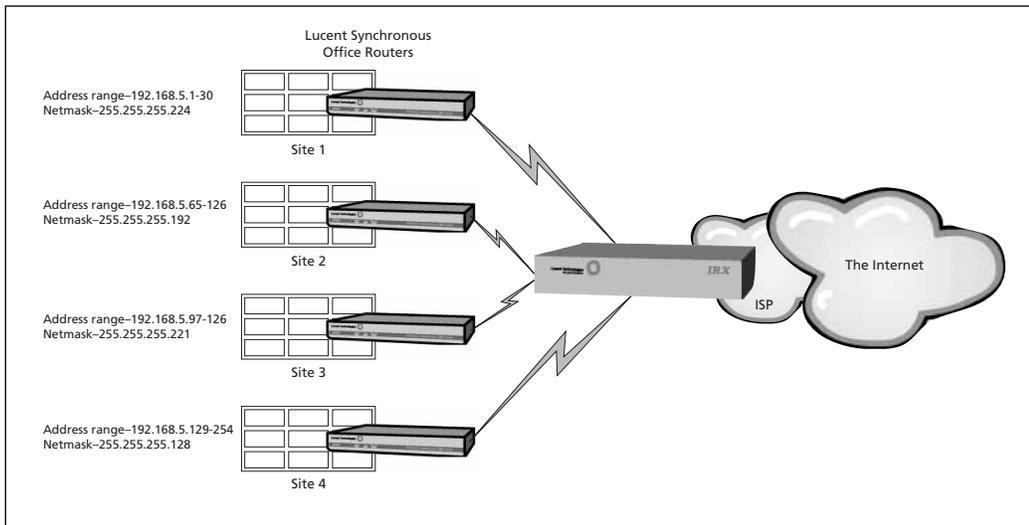


Figure 5. VLSM subnetting across multiple sites

Scenario #1: VLSM, Subnetting a Class C Address across Multiple Sites

For many ISPs and corporations, VLSM is the sole reason for implementing OSPF. As the number of available TCP/IP addresses have become more scarce, more ISPs and corporations are looking for ways to better utilize IP addresses. These organizations often want to spread a class C network across multiple networks or sites, rather than wasting full class C-size address ranges on a single network or site.

Figure 5 demonstrates Internet service based on an OSPF network scheme using VLSM. In this example, the ISP has four separate customer sites of various sizes utilizing a single class C network address, 192.168.2.0, for their network access. Each customer site network can best be configured as a default route stub area because their routers are linking to only one site.

Scenario #2: VLSM, Subnetting a Class C Address across an ISP POP

Another VLSM application applies to the networks at the ISP's own points-of-presence (POPs), their network operations centers. In this case, the ISP needs to subnet a single class C network across multiple PortMaster Communication Servers and Office Routers as demonstrated in Figure 6.

Using RIP in this configuration can create routing problems. When a dialin customer with a 192.168.2.75 address accesses the Internet through PortMaster #2, the IRX will not

reliably get back to that specific PortMaster. Without VLSM, the router considers all IP addresses beginning with 192.168.2 to be on the same network. It can only attempt to contact

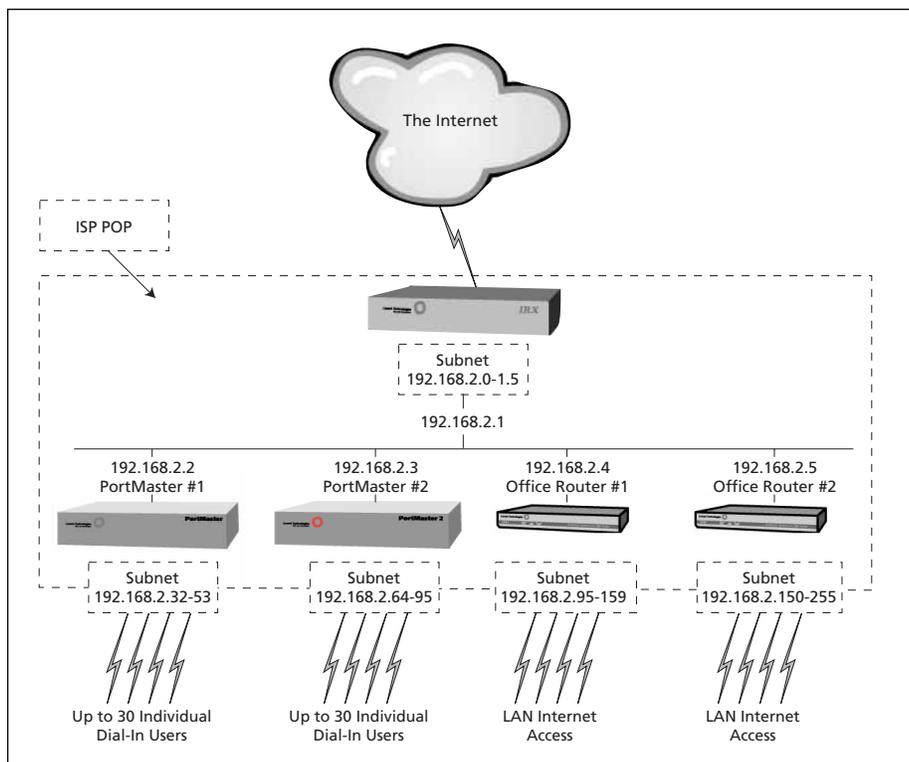


Figure 6. Subnetting across an ISP's POP

192.168.2.75 through any one of the two PortMasters or two Office Routers shown in Figure 6. If it contacts any except PortMaster #2, communication will be rejected.

With OSPF's VLSM support, however, the IRX can reliably forward data. The IRX's routing table states that the 192.168.2.64 network can be reached via 192.168.2.3. The network mask 255.255.255.224 provides the necessary address range information, revealing that this subnet supports 30 IP addresses beginning with 192.168.2.64. Because 192.168.2.75 is within the 192.168.2.64-95 subnet, the IRX can correctly forward the data to the PortMaster at the 192.168.2.3 address.

2.0 Summary

Many network managers are migrating from RIP to OSPF for reasons like the following:

- RIP doesn't scale well, but OSPF can effectively support much larger networks.
- RIP updates can bog down a larger network. When an OSPF autonomous system has been correctly divided into multiple areas, OSPF updates create much less overall burden on network performance.
- RIP is incapable of recognizing classless, subnetted segments with network masks, but OSPF supports variable-length subnet masks that recognize subnets of any size.

Although OSPF is not for every network and involves more network planning and setup than RIP does, OSPF provides the performance and the flexibility required by many of today's ISPs and enterprise-wide networks. For these advanced networks, OSPF is a reliable and proven routing protocol choice.

3.0 Bibliography

Douglas Comer, *Internetworking with TCP/IP, Volume 1, Principles, Protocols, and Architecture*, 3rd Edition, Prentice Hall, 1995.

C. Hedrick "Routing Information Protocol," RFC-1058, June 1988.

Christian Huitema, *Routing in the Internet*, Prentice Hall PTR, 1995.

J. Moy, "OSPF Version 2", RFC-1583, March 1994.

Radia Perlman, *Interconnections: Bridges and Routers*, Addison-Wesley Professional Computing Series, 1992.

Lucent Technologies
Remote Access Business Unit

4464 Willow Road

Pleasanton, California 94588-8519

Tel: (925) 730-2500 Fax: (925) 730-2510

website: <http://www.lucent.com/dns/portmaster>

Copyright © 1999 Lucent Technologies, Inc.

Printed in the USA 04/99 LT-DAT0717

Lucent Technologies
Bell Labs Innovations

