Network Working Group J. Moy, Editor
Request for Comments: 1246 Proteon, Inc.
July 1991

Experience with the OSPF protocol

Status of this Memo

This memo provides information for the Internet community. It does not specify any Internet standard. Distribution of this memo is unlimited

Abstract

This is the second of two reports on the OSPF protocol. These reports are required by the IAB/IESG in order for an Internet routing protocol to advance to Draft Standard Status. OSPF is a TCP/IP routing protocol, designed to be used internal to an Autonomous System (in other words, OSPF is an Interior Gateway Protocol).

Version 1 of the OSPF protocol was published in RFC 1131. Since then OSPF version 2 has been developed. Version 2 has been documented in RFC 1247. The changes between version 1 and version 2 of the OSPF protocol are explained in Appendix F of RFC 1247. It is OSPF Version 2 that is the subject of this report.

This report documents experience with OSPF V2. This includes reports on interoperability testing, field experience, simulations and the current state of OSPF implementations. It also presents a summary of the OSPF Management Information Base (MIB), and a summary of OSPF authentication mechanism.

Please send comments to ospf@trantor.umd.edu.

Table of Contents

1.0 Introduction

This document addresses, for OSPF V2, the requirements set forth by the IAB/IESG for an Internet routing protocol to advance to Draft Standard state. This requirements are briefly summarized below. The remaining sections of this report document how OSPF V2 satisfies these requirements:

- The specification for the routing protocol must be well written such that independent, interoperable implementations can be developed solely based on the specification. For example, it should be possible to develop an interoperable implementation without consulting the original developers of the routing protocol.
- A Management Information Base (MIB) must be written for the protocol. The MIB must be in the standardization process, but does not need to be at the same level of standardization as the routing protocol.
- The security architecture of the protocol must be set forth explicitly. The security architecture must include mechanisms for authenticating routing messages and may include other forms of protection.
- Two or more interoperable implementations must exist. At least two must be written independently.
- There must be evidence that all features of the protocol have been tested, running between at least two implementations. This must include that all of the security features have been demonstrated to operate, and that the mechanisms defined in the protocol actually provide the intended protection.
- There must be significant operational experience. This must include running in a moderate number routers configured in a moderately complex topology, and must be part of the operational Internet. All significant features of the protocol must be exercised. In the case of an Interior Gateway Protocol (IGP), both interior and exterior routes must be carried (unless another mechanism is provided for the exterior routes). In the case of a Exterior Gateway Protocol (EGP), it must carry the full complement of exterior routes.

This report is a compilation of information obtained from many people. The reader is referred to specific people when more information on a subject is available. People references are gathered into Section 10.0, in a format similar to that used in [4].

1.1 Acknowledgments

The OSPF protocol has been developed by the OSPF Working Group of the Internet Engineering Task Force. Many people have contributed to this report. They are listed in Section 10.0 of this report.

2.0 Documentation

Version 1 of the OSPF protocol is documented in RFC 1131 [1]. OSPF Version 2, which supersedes Version 1, has been documented in RFC 1247 [2]. The differences between OSPF Version 1 and Version 2 are relatively minor, and are listed in Appendix F of RFC 1247 [2]. All information presented in this report concerns OSPF V2 unless explicitly mentioned otherwise.

The OSPF protocol was developed by the OSPF Working Group of the Internet Engineering Task Force. This Working Group has a mailing list, *ospf@trantor.umd.edu*, where discussions of protocol features and operation are held. The OSPF Working Group also meets during the quarterly Internet Engineering Task Force conferences. Reports of these meeting are published in the IETF's Proceedings. In addition, two reports on the OSPF protocol have been presented to the IETF plenary (see "Everything You Ever Wanted to Know about OSPFIGP" in [5] and "OSPF Update" in [6]).

The OSPF protocol began undergoing field trials in Spring of 1990. A mailing list, *ospf-tests@seka.cso.uiuc.edu*, was formed to discuss how the field trials were proceeding. This mailing list is maintained by Ross Veach of the University of Illinois [rrv]. Archives of this list are also available. There has been quite a bit of discussion on the list concerning OSPF/RIP/EGP interaction.

A OSPF V2 Management Information Base has also been developed and published in [3]. For more information, see Section 3.0 of this report.

There is a free implementation of OSPF available from the University of Maryland. This implementation was written by Rob Coltun [rcoltun]. Contact Rob for details.

3.0 MIB

An OSPF Management Information Base has been published in RFC 1248 [3]. The MIB was written by Rob Coltun [rcoltun] and Fred Baker [fbaker]. The OSPF MIB appears on the mgmt subtree as SMI standard-mib 13.

The OSPF MIB was originally developed by Rob Coltun of the University of Maryland, under contract to Advanced Computer Communications. A subset of his proposal was implemented to facilitate their development, and represents operational experience of a sort.

The MIB consists of a general variables group and ten tables:

TABLE 1. OSPF MIB Organization

Group Name	Description
ospfGeneralGroup	General Global Variables
ospfAreaTable	Area Descriptions
ospfStubAreaTable	Default Metrics, by Type of Service
ospfLsdbTable	Link State Database
ospfAreaRangeTable	Address Range Specifications
ospfHostTable	Directly connected Hosts
ospfIfTable	OSPF Interface Variables
ospfIfMetricTable	Interface Metrics, by Type of Service
ospfVirtIfTable	Virtual Links
ospfNbrTable	(Non-virtual) OSPF Neighbors
ospfVirtNbrTable	Virtual OSPF Neighbors

As MIBs go, the OSPF MIB is quite large; 105 objects. The following are some statistics describing the distribution of the MIB's variables:

- 11 define the above Group and Tables
- 10 define the Entry in a Table
- 7 are Counters
- 6 are Gauges
- 68 objects mandated by the OSPF Version 2 Specification

Section D.2 of the OSPF V2 specification [2] lists a set of required statistics that an implementation must maintain. These statistics have been incorporated into the OSPF MIB. The MIB's thirteen Counters and Gauges enable evaluation of the OSPF protocol's performance in an operational environment. Most of the remainder of the MIB's variables parameterize the many features that OSPF provides the network administrator.

For more information on the MIB contact Fred Baker [fbaker].

4.0 Security architecture

In OSPF, all protocol packet exchanges are authenticated. The OSPF packet header (which is common to all OSPF packets) contains a 16-bit Authentication type field, and 64-bits of Authentication data. Each particular OSPF area must run a single authentication scheme, as indicated by the Authentication type field. However, authentication keys can be configured by the system administrator on a per-network basis.

When an OSPF packet is received from a network, the OSPF router first verifies that it indicates the correct Authentication type. The router then authenticates the packet, running a verification algorithm using the configured authentication key, the 64-bits of Authentication data and the rest of the OSPF packet data as input. The precise algorithm used is dictated by the Authentication type. Packets failing the authentication algorithm are dropped, and the authentication failure is noted in a MIB-accessible variable (see [3]).

There are currently few Authentication types in use. The current assignments are:

TABLE 2. Current OSPF Authentication types.

Type code	Algorithm
0	No authentication performed.
1	Simple (clear) password.
2-255	Reserved for assignment by the IANA (iana@isi.edu)
> 255	Available for local (per-AS) definition.

For more information on OSPF's authentication procedures, see Sections 8.1, 8.2, and Appendix E of [2].

5.0 Implementations

The are multiple, interoperable implementations of OSPF currently available. This section gives a brief overview of the five implementations that have participated in at least one round of interoperability testing[1]. Other implementations do exist, but because of commercial realities (e.g., the product is not yet announced) they unfortunately cannot be listed here.

The five implementations that have participated in the OSPF interoperability testing are (listed in alphabetical order):

- **3com**. This implementation was wholly developed by 3com. It has participated in all three rounds of interoperability testing. It is also the only implementation of OSPF's TOS routing.. The 3com implementation consists of approximately 9000 lines of C code, including comments but excluding user interface and MIB code. Consult Dino Farinacci [dino] for more details.
- **ACC.** This implementation is based on the University of Maryland code. It participated in the last two rounds of interoperability testing. It also contains the only implementation of (a precursor to) the OSPF MIB (see Section 3.0 for details), which it uses for monitoring and configuration. The ACC implementation consists of approximately 24,000 lines of C code, including its OSPF MIB code. Consult Fred Baker [fbaker] for more details.
- **Proteon**. This implementation was wholly developed by Proteon. It has participated in all three rounds of interoperability testing. It is also the only implementation that has a significant amount of field experience (see Section 6.0 for details). The Proteon implementation consists of approximately 9500 lines of C code, including comments but excluding user interface code. Consult John Moy [jmoy] for more details.
- **Wellfleet**. This implementation has participated in all three rounds of interoperability testing. Consult Jonathan Hsu [jhsu] for more details.
- **University of Maryland.** This implementation was developed wholly by Rob Coltun at the University of Maryland. It has formed the basis for a number of commercial OSPF implementations, and also participated in the latest round of interoperability testing. The University of Maryland implementation consists of approximately 10,000 lines of C code. Consult Rob Coltun [rcoltun] for more details.

Note that, as required by the IAB/IESG for Draft Standard status, there are multiple interoperable independent implementations, namely those from 3com, Proteon and the University of Maryland.

---

1. For a detailed discussion of OSPF interoperability testing, see Section 7.0 of this report.

6.0 Operational experience

This section discusses operational experience with the OSPF protocol. Version 1 of the OSPF protocol began to be deployed in the Internet in Spring of 1990. The results of this original deployment were reported to the mailing list *ospf-tests@seka.cso.uiuc.edu*[2]. No protocol bugs were found in this first deployment, although several additional features were found to be desirable. These new features were added to the protocol in OSPF Version 2.

The OSPF protocol is now deployed in a number of places in the Internet. In this section we focus on three highly visible systems, namely the NASA Sciences Internet, BARRNet and OARnet. The dimensions of these three OSPF systems is summarized in the following table:

TABLE 3. Three operational OSPF deployments

Name	Version 1 date	Version 2 date	# routers	#externals
NSI	4/13/90	1/1/91	15	496
BARRNet	4/90	11/90	14	1816
OARnet	10/15/90	not yet	13	135-140

All the above deployments are using the Proteon OSPF implementation. There is one other deployment worth mentioning in this context. 3com has started to deploy OSPF on their corporate network. They have 8 of their routers running OSPF (the 3com implementation), and are planning on cutting over the remaining routers (20 in all). Currently they have two operational routers running OSPF and RIP simultaneously. One converts OSPF data to RIP data, and the other RIP data to OSPF data. For more details, contact Dino Farinacci [dino].

6.1 NSI

The NASA Science Internet (NSI) is a multiprotocol network, currently supporting both DECnet and TCP/IP protocols. NSI's mission is to provide reliable high-speed communications to the NASA science community. The NASA Science Internet connects with other national networks including the National Science Foundation's NSFNET, the Department of Energy's ESnet and the Department of Defense's MILNET. NSI also has international connections to Japan, Australia, New Zealand, Chile and several European countries.

For more information on NSI, contact Jeffrey Burgan [jeff] or Milo Medin [medin].

2. Archives of this mailing list are available from Ross Veach [rrv].

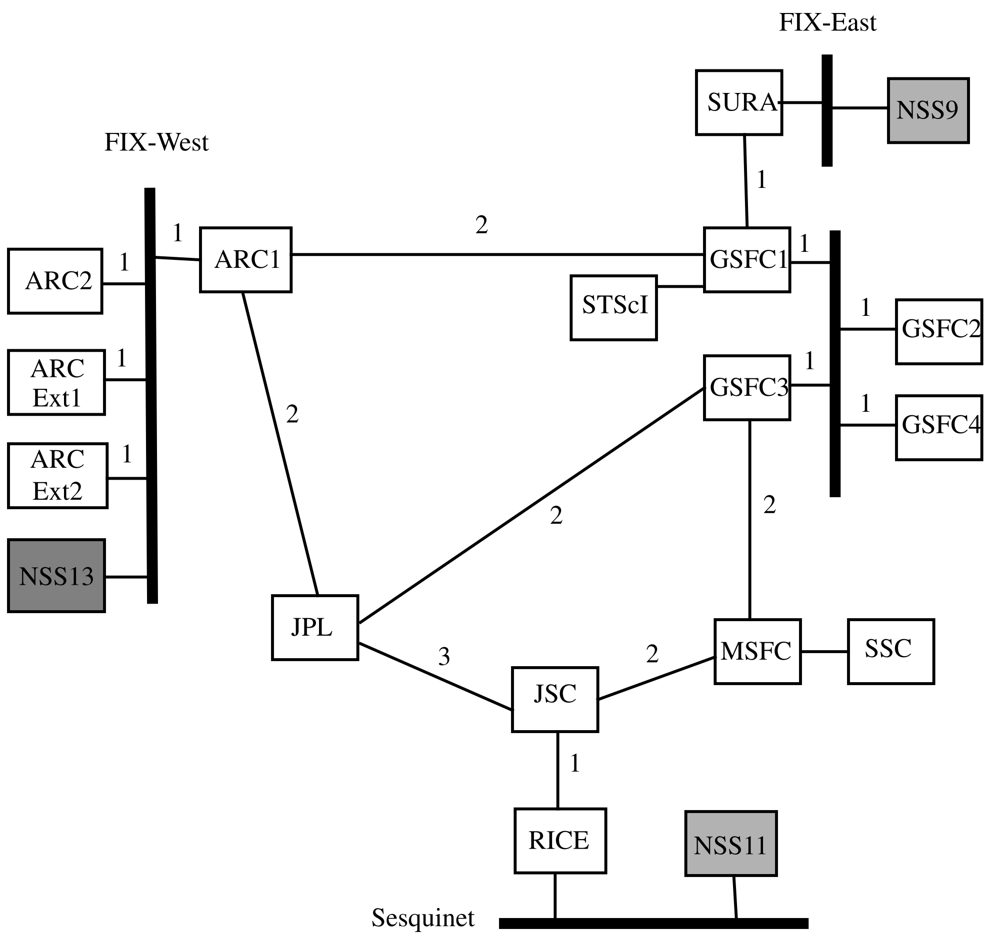

Figure 1: The NASA Science Internet OSPF System

6.1.1 NSI's OSPF system

NSI was one of the initial deployment sites for OSPF Version 1, having deployed the protocol in April 1990. NSI has been running OSPF V2 since 1/1/91. They currently have 15 routers in their OSPF system. This system is pictured in Figure 1. It consists of a nationwide collection of serial lines, with ethernets at hub sites. The numbers associated to interfaces/links in Figure1 are the associated OSPF costs. Note that certain links have been weighted so that they are less preferable than others.

Many of NSI's OSPF routers are speaking either RIP and/or EGP as well as OSPF. Routes from these other routing protocols are selectively imported into their OSPF system as externals. The current number of imported externals is 496.

All NSI externals are imported as OSPF type 2 metrics. In addition, NSI uses the OSPF external route tag to manage the readvertisement of external routes. For example, a route learned at one edge of the NSI system via EGP can be tagged with the number of the AS from which it was learned. Then, as the OSPF external LSA describing this route is flooded through the OSPF system, this tagging information is distributed to all the other AS boundary routers. A router on the other edge of the NSI can then say that it wants to readvertise (via EGP) routes learned from one particular AS but not routes learned from another AS. This allows NSI to implement transit policies at the granularity of Autonomous Systems, instead of network numbers, which greatly reduces the network's configuration burden.

NSI has also experimented with OSPF stub areas, in order to support routers having a small amount of memory.

6.1.2 NSI - Deployment analysis

NSI ran a couple of experiments after OSPF's deployment to test OSPF's convergence time in the face of network failures, and to compare the level of routing traffic in OSPF with the level of routing traffic in RIP. These experiments were included in NSI status reports to the OSPF plenary.

The first experiment consisted of running a continuous ICMP ping, and then bringing down one of the links in the ping packet's path. They then timed how long it took OSPF to find an alternate path, by noticing when the pings resumed. The result of this experiment is contained in Milo Medin's "NASA Sciences Internet Report" in [8]. It shows that the interrupted ping resumed in three seconds.

The second experiment consisted in analyzing the amount of routing protocol traffic that flow over an NSI link. One of the NSI links was installed, but did not have any active users yet. For this reason, all traffic that flowed over the link was routing protocol traffic. The link was instrumented to continuously measure the amount of bandwidth consumed, first in the case where RIP was running, and then in the case of where OSPF was running. The result is shown graphically in Jeffrey Burgan's "NASA Sciences Internet" report in [9]. It shows that OSPF consumes many times less network bandwidth than RIP.

6.2 BARRNet

BARRNet is the NSFNet regional network in Northern California. At the present time, it serves approximately 80 member sites in an area stretching from Sacramento in the north-east to Monterey in the in the south-west. Sites are connected to the network at speeds from 9.6Kbps to full T1 using Proteon and cisco routers as well as a Xylogics terminal server. The membership is composed of a mix of university, government, and commercial organizations. BARRNet has interconnections to the NSFNet (peering with both T1 and T3 backbones at Stanford University), ESNet (peering at LLNL, with additional multi-homed sites at LBL, SLAC, and NASA Ames), and DDN national networks (peering at the FIX network at NASA Ames), and to the statewide networks of the University of California (peering at U.C. Berkeley) and the California State University system (peering at San Francisco State and Sacramento State).

Topologically, the network consists of fourteen OSPF-speaking Proteon routers, which as a “core”, with six of these redundantly connected into a ring. All “core” sites are interconnected via full T1 circuits. Other member sites attach as “stub” connections to the “core” sites. The bulk of these are connected in a “star” configuration at Stanford University, with lesser numbers at other “core” sites.

Contact Vince Fuller [vaf] for more information on BARRNet.

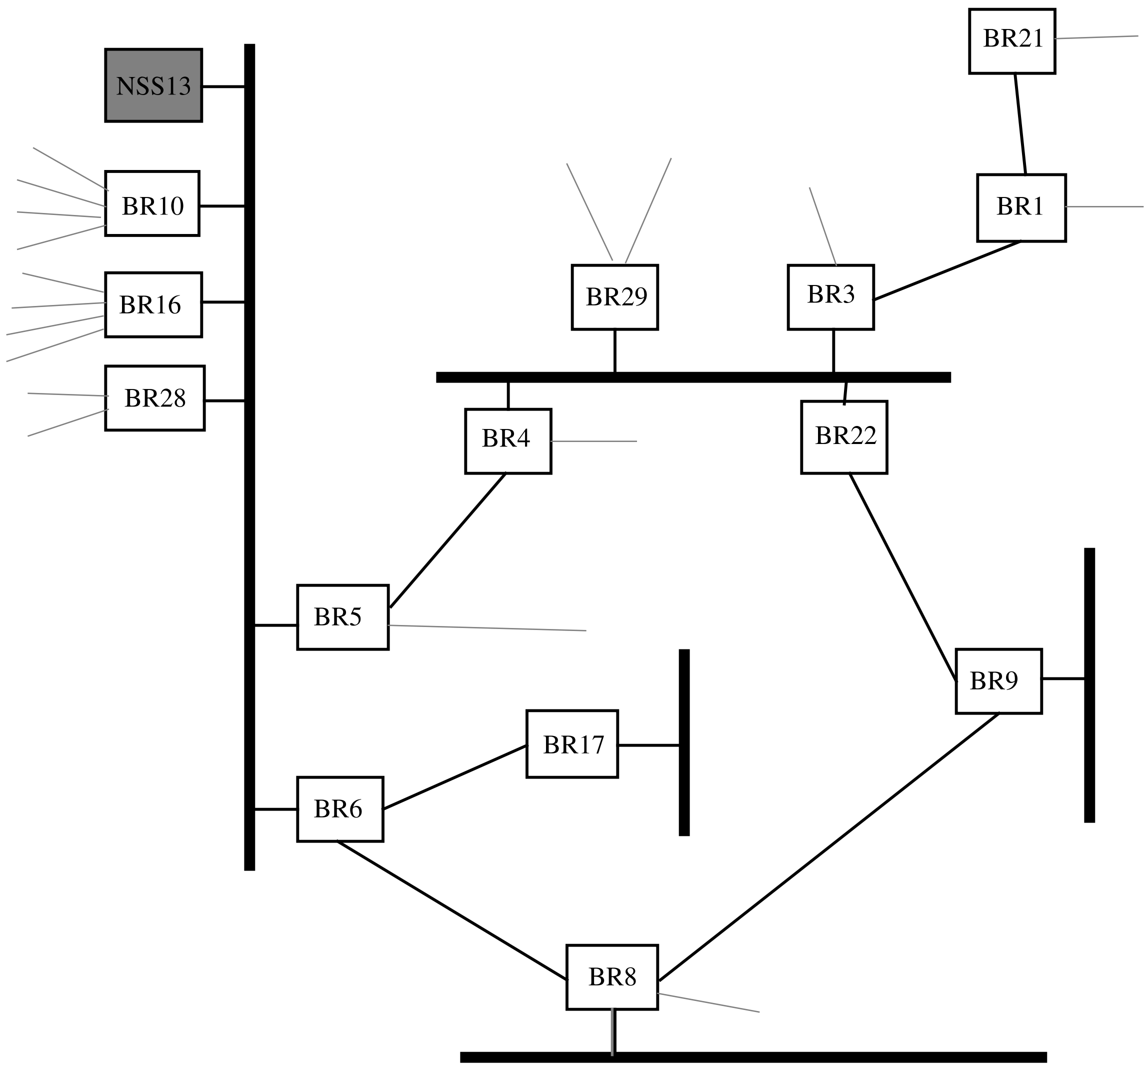

Figure 2; The BARRNet OSPF system

6.2.1 BARRNet's OSPF system

BARRNet was also one of the initial deployment sites for OSPF Version 1, having deployed the protocol in April 1990. BARRNet has been running OSPF V2 since November 1990. They currently have 14 routers in their OSPF system. The BARRNet OSPF system is pictured in Figure 2. It consists of a collection of T1 serial lines, with ethernets at hub sites.

Most of BARRNet's OSPF routers are speaking either RIP and/or EGP as well as OSPF. Routes from these other routing protocols are selectively imported into their OSPF system as externals. A large number of external routes are imported; the current number is1816. The bulk of these are the T1 NSFNet routes, followed by several hundred NSN routes, around 60-80 BARRNet routes from the non-OSPF system, and several dozen from ESNet.

All external routes are imported into the BARRNet system as external type 1 metrics. In addition, BARRnet, like NSI, uses the OSPF's external route tagging feature to help manage the readvertisement of external routes via EGP.

BARRnet is also using four stub OSPF areas in order to collapse subnet information. These stub areas all consist of a single LAN. They do not contain any OSPF routers in their interiors.

6.2.2 BARRNet - Deployment analysis

Initial deployment of OSPF Version 1 in BARRNet pointed to the need for two new protocol features that were added to OSPF V2, namely:

- Addition of the **forwarding address** to OSPF external LSAs. This eliminated the extra hops that were being taken in BARRNet when only routers BR5 and BR6 were exchanging EGP information with the NSS (see Figure 2). Without the forwarding address feature, that meant that NSFNet traffic handled by routers BR10, BR16 and BR28 was taking an extra hop to get to the NSS.
- Addition of **stub areas**. This was an attempt to get OSPF running on some of the BARRNet routers that had insufficient memory to deal with all of BARRNet's external routes.

6.3 OARnet

OARnet, the Ohio Academic Resources Network, is the regional network for the state of Ohio. It serves the entire higher education community, providing Ohio schools access to colleagues worldwide. The Ohio Supercomputer Center and the NSF Supercomputer Centers are reached through OARnet. Libraries, databases, national and international laboratories and research centers are accessible to faculty, helping make Ohio schools competitive.

OARnet was established in 1987 to provide state-wide access to the CRAY at the Ohio Supercomputer Center in Columbus, Ohio. Since then it has evolved into a network supporting all aspects of higher education within Ohio. A primary goal of OARnet is to facilitate collaborative projects and sharing of resources between institutions, including those outside the state. OARnet connections are available to Ohio academic institutions and corporations engaged in research, product development, or instruction. Colleges, universities, and industries currently use OARnet connections to communicate within the state and with colleagues around the country.

OARnet uses the Internet (TCP/IP) and DECNET protocols. OARnet participants using TP/IP protocols are connected to the worldwide Internet, which includes all the major networks open to non-classified research. OARnet is also connected to NSFNet, the national research and education network sponsored by the National Science Foundation. It has gateways to BITNET, CSNET, CICNet (a network connecting the Big Ten universities), and the NASA Science Internet.

For more information on OARnet, contact Kannan Varadhan [kannan].

6.3.1 OARnet's OSPF system

OARnet has been running OSPF Version 1 since October 15, 1990. They currently have 14 routers in their OSPF system. The OARnet OSPF system is pictured in Figure 3.

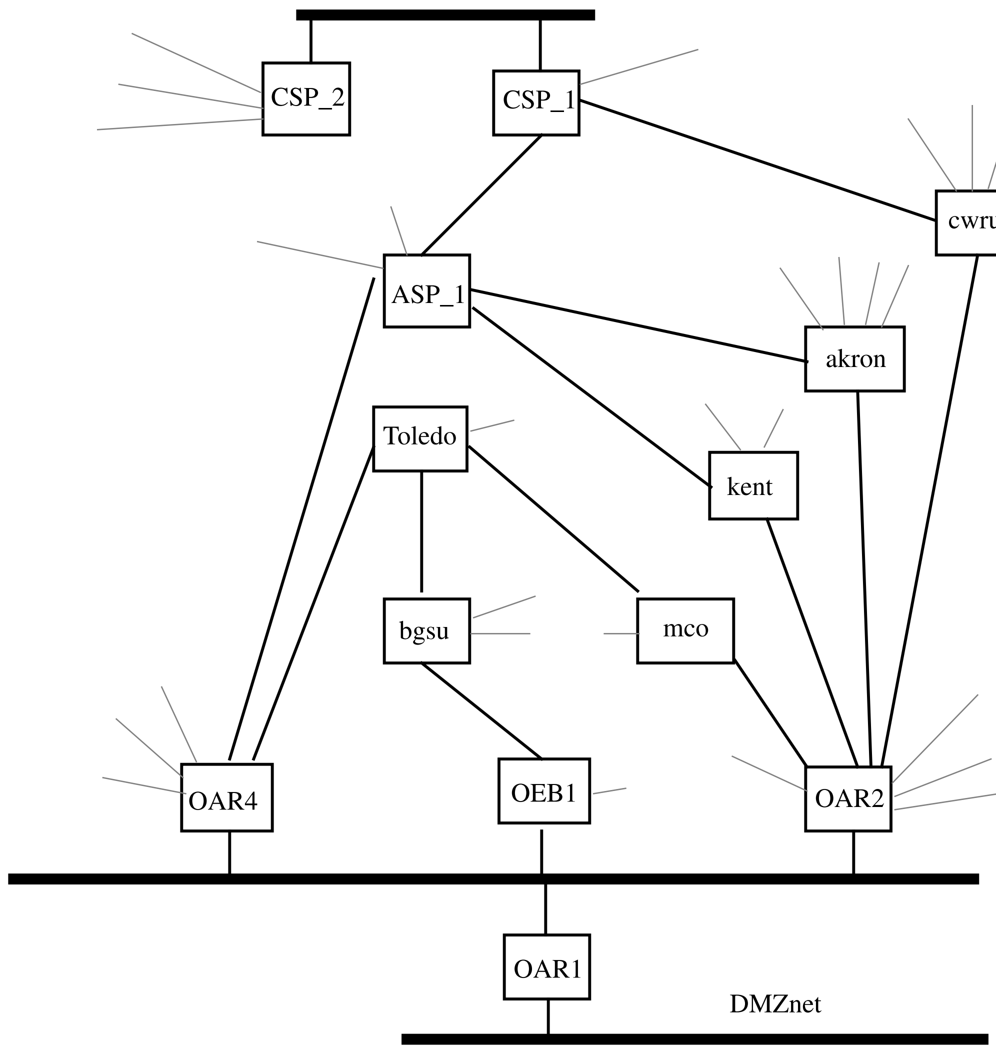

Figure 3; The OARnet OSPF system

There are 29 sites connected directly to the OARnet backbone. All 13 of OARnet's OSPF routers act as ASBRs. There are 40 OSPF internal routes on OARnet's network, and they import about 120 routes from RIP. OARnet runs EGP on the DMZnet at Columbus, which connects them to CICNet. The router connecting OARnet to DMZnet (OAR1 in Figure 3) runs EGP on the DMZnet side, and OSPF and RIP on the OARnet backbone. No EGP routes are imported into the OSPF system. The OAR1 router is configured to generate a default when EGP routes are available. The OAR1 router is the keystone for routing on OARnet's network, in that it acts as an intermediary for all of OARnet's RIP centric routers.

OARnet uses the Event Logging System on its Proteon routers to generate traps for “interesting” events related to routing. They have these traps sent to an SNMP management station, where the logs are collected for later perusal.

6.3.2 OARnet - Deployment analysis

OARnet is monitoring their OSPF system via collection of traps on their SNMP management station. The following is a report on their observations. It has been edited slightly to conform better with the other text and maps presented in this report. For more information, contact Kannan Varadhan [kannan]:

3 of our 10 DS1 circuits are on digital microwave, and these tend to flap occasionally. Our observations indicate that the routers bring up links, and adjacencies, on average, in about 2 seconds. Routes fallback to alternate backup paths instantly. Whole blocks of routes cut over the instant the adjacencies are formed.

In contrast to this, our RIP routes would take about 3-6 minutes to cutover, and, on occasion, would not cut back to the preferred paths. This was our prime motivation in switching to OSPF.

We attempted to duplicate Milo Medin’s ping test to dramatically illustrate the performance of RIP over OSPF. To do this, we selected a host on the farthest point from our workstation, and ran a continuous ping to it. We would then bring down a primary DS1 circuit, and watch the time it took to switch to the fallback route. Following this, we would bring the circuit back up, and study the time it took to re-sync to the new path. With RIP, we were unable to fully complete the experiment, because the farthest point was exactly equal to the new (and preferred) primary path, and therefore, RIP would never choose it on it’s own, until the path it was currently using failed. With OSPF, it took about 2 seconds to synchronize over a new, much slower 56kb path, and less than a second when the DS1 circuit came back up.

Here are some more observations of the OARnet OSPF system’s behavior:

- 131.187.36.0 is the 56kb line to Kent State University. Kent also has a DS1 circuit leading into ASP, the Akron Pop. Likewise, UAkron.edu has a similar configuration. A roundabout backup path exists when traffic heads up to Cleveland over a couple of DS1 circuits, and then down a 56kb backup path used by another school in the Cleveland area.

  Some statistical information:

  1. 09:55:17: SPF.37: new route to Net 131.187.36.5, type SPF cost 32
  2. 09:55:18: SPF.37: new route to Net 131.187.36.6, type SPF cost 22
  3. 09:55:20: SPF.21: State Change, nbr 131.187.27.6, new state <Full>, event 9
  4. 09:55:21: SPF.37: new route to Net 131.187.36.5, type SPF cost 31
  5. 09:55:22: SPF.37: new route to Net 131.187.36.6, type SPF cost 21

6. 09:55:28: SPF.21: State Change, nbr 131.187.21.5, new state <Full>, event 9
7. 09:55:29: SPF.21: State Change, nbr 131.187.51.6, new state <Full>, event 9
8. 09:55:31: SPF.37: new route to Net 131.187.36.5, type SPF cost 22
9. 09:55:33: SPF.37: new route to Net 131.187.36.5, type SPF cost 11

The Akron router restarts, and has to re-sync with all the lines. This restart is confirmed when one looks at the traps from gwCSP1. Traps from gwASP1 still do not get through to us, because traps are sent via UDP, and gwASP1's routing tables are not fully configured yet.

Events 1 and 2 are route changes routing traffic via Cleveland, across 2 DS1 circuits and a 56kb line.

When the DS1 circuit to UAkron came up, routes instantly cut over to use this as a better least cost path. This is shown in events 3, 4 and 5.

In a few seconds, the line to Columbus is the next one up. This is event 6. Event 8 relates to this cutover, and is the best path yet. When the DS1 circuit to Kent is up, the link is used instantly.

We are able to make such a definitive conclusion of these traps on the basis of the topological information that we have about the network and the means used to monitor them.

- To illustrate the time required to fully synchronize a database, we piece together a few adjacency forming traces...

  Please bear in mind that these time stamps are the time stamps on the management station, and are not to be taken as the absolute truth. Things we haven't taken into account are transit times of messages, ordering of events (SNMP traps are sent using UDP), loss of event reports (recall that an entire synchronization sequence of gwASP1 on the ASP-CSP link is missing), etc.

  The trace below corresponds to the Akron router, gwASP1 bring up the link in the previous section. This is as observed on the other end of the line, gwCSP1.

  REPORT DATE: 02/26/91ROUTER: gwcsp1

  09:55:06: SPF.15: State Change, ifc 131.187.22.6, new state <Point-To-Point>, event 1

  09:55:06: GW.xxx: Link Up Trap: 09:55:07: SPF.37: new route to Net 131.187.22.5, type SPF cost 1

  09:55:07: SPF.21: State Change, nbr 131.187.22.5, new state <Init>, event 1

  09:55:09: SPF.37: new route to Net 131.187.27.5, type SPF cost 22

  09:55:11: SPF.21: State Change, nbr 131.187.22.5, new state <ExStart>, event 14

  09:55:11: SPF.21: State Change, nbr 131.187.22.5, new state <2-Way>, event 3

  09:55:12: SPF.21: State Change, nbr 131.187.22.5, new state <Exchange>, event 5

  09:55:12: SPF.21: State Change, nbr 131.187.22.5, new state <Full>, event 9

  09:55:12: SPF.21: State Change, nbr 131.187.22.5, new state <Loading>, event 6

  Below, is another trace of the same router restart sequence, where the router is proceeding to bring up other DS1 circuits. Bringing up the first adjacency took about 5 seconds. Subsequent adjacencies take the router less than a second as seen below.

REPORT DATE: 02/26/91ROUTER: gwasp1

09:55:20: SPF.15: State Change, ifc 131.187.27.5, new state <Point-To-Point>, event 1

09:55:20: GW.xxx: Link Up Trap: 09:55:20: SPF.21: State Change, nbr 131.187.27.6, new state <Init>, event 1

09:55:20: SPF.21: State Change, nbr 131.187.27.6, new state <ExStart>, event 14

09:55:20: SPF.21: State Change, nbr 131.187.27.6, new state <Exchange>, event 5

09:55:20: SPF.21: State Change, nbr 131.187.27.6, new state <Full>, event 9

09:55:21: SPF.21: State Change, nbr 131.187.27.6, new state <Loading>, event 6

09:55:24: SPF.21: State Change, nbr 131.187.51.6, new state <Init>, event 1

09:55:24: SPF.21: State Change, nbr 131.187.21.5, new state <Init>, event 1

09:55:25: SPF.37: new route to Net 131.187.21.6, type SPF cost 13

09:55:25: SPF.37: new route to Net 131.187.51.5, type SPF cost 22

09:55:28: SPF.21: State Change, nbr 131.187.21.5, new state <ExStart>, event 14

09:55:28: SPF.21: State Change, nbr 131.187.21.5, new state <2-Way>, event 3

09:55:28: SPF.21: State Change, nbr 131.187.21.5, new state <Exchange>, event 5

09:55:28: SPF.21: State Change, nbr 131.187.21.5, new state <Full>, event 9

09:55:28: SPF.21: State Change, nbr 131.187.21.5, new state <Loading>, event 6

09:55:29: SPF.37: new route to Net 131.187.51.6, type SPF cost 1

09:55:29: SPF.37: new route to Net 131.187.21.5, type SPF cost 1

09:55:29: SPF.21: State Change, nbr 131.187.51.6, new state <Exchange>, event 5

09:55:29: SPF.21: State Change, nbr 131.187.51.6, new state <ExStart>, event 14

09:55:29: SPF.21: State Change, nbr 131.187.51.6, new state <2-Way>, event 3

09:55:29: SPF.21: State Change, nbr 131.187.51.6, new state <Full>, event 9

09:55:29: SPF.21: State Change, nbr 131.187.51.6, new state <Loading>, event 6

A transient fault on a DS1 circuit, causes the line to flap. All routers quickly reroute around the flap, and the router itself takes about 2 seconds to bring up the adjacency once more.

REPORT DATE: 02/26/91ROUTER: gwasp1

14:33:43: GW.xxx: Link Up Trap:

14:34:19: SPF.15: State Change, ifc 131.187.22.5, new state <Down>, event 7

14:34:19: GW.xxx: Link Failure Trap:

14:34:19: SPF.47: Net 131.187.22.6 now unreachable

14:34:36: SPF.15: State Change, ifc 131.187.22.5, new state <Point-To-Point>, event 1

14:34:36: GW.xxx: Link Up Trap:

14:34:37: SPF.37: new route to Net 131.187.22.6, type SPF cost 1

14:34:45: SPF.21: State Change, nbr 131.187.22.6, new state <2-Way>, event 3

14:34:45: SPF.21: State Change, nbr 131.187.22.6, new state <Init>, event 1

```
14:34:46: SPF.21: State Change, nbr 131.187.22.6, new state <ExStart>, event 14
14:34:46: SPF.21: State Change, nbr 131.187.22.6, new state <Exchange>, event 5
14:34:47: SPF.21: State Change, nbr 131.187.22.6, new state <Full>, event 9
14:34:47: SPF.21: State Change, nbr 131.187.22.6, new state <Loading>, event 6
```

- On the amount of time it takes for a router to restart, and become fully synchronized. Taking the logs in the previous instance, we notice that the CSP-ASP link comes up at 9:55:06. The last link is observed to be up at 9:55:29, which is less than a minute.
- On the RIP equivalent of the tests, it took us 3 minutes to cutover to the slower speed fallback route, and we lost countless many packets. The routes never cutover to the higher speed paths when available, and we waited well over 30 minutes watching this, wondering why. Unfortunately, at this point, we seem to have lost the RIP statistics.

  On the OSPF version, we have...

```
{nisca danw 51}
ping 131.187.25.6 PING 131.187.25.6 (131.187.25.6):
56 data bytes 64 bytes from 131.187.25.6: icmp seq=0 ttl(255-ttl)=54(201). time=20 ms
                    [...]
64 bytes from 131.187.25.6: icmp seq=10 ttl(255-ttl)=54(201). time=20 ms
                    ||T1 down
64 bytes from 131.187.25.6: icmp seq=14 ttl(255-ttl)=54(201). time=180 ms
64 bytes from 131.187.25.6: icmp seq=15 ttl(255-ttl)=54(201). time=60 ms
                    [...]
64 bytes from 131.187.25.6: icmp seq=38 ttl(255-ttl)=8(247). time=1300 ms
64 bytes from 131.187.25.6: icmp seq=39 ttl(255-ttl)=54(201). time=820 ms
                    ||Tl Up
64 bytes from 131.187.25.6: icmp seq=40 ttl(255-ttl)=54(201). time=20 ms
64 bytes from 131.187.25.6: icmp seq=41 ttl(255-ttl)=54(201). time=20 ms
131.187.25.6 PING Statistics
51 packets transmitted, 48 packets received, 5% packet loss
round-trip (ms) min/avg/max = 20/277/1300
```

6.4 Features exercised during operational deployment

In operational environments, all basic mechanisms of the OSPF protocol have been exercised. These mechanisms include:

- **Designated Router election**. There have been operational deployments have as many as 8 OSPF routers attached to a single broadcast network.

- **Database synchronization.** This includes OSPF's **adjacency bringup** and **reliable flooding** procedures. Large operational OSPF link state databases (e.g., BARRNet) have provided a thorough test of these mechanisms.
- **Flushing advertisements.** The procedure for flushing old or unreachable advertisements (the **MaxAge** procedure) has been tested operationally. It is interesting to note that flushing of advertisements does occur more during interoperability testing (because of the constant restarting of routers) that it does operationally. For example, in a week of BARRNet statistics, 9650 advertisements were flushed, while 688,278 new advertisements were flooded.
- **Import of external routes**. All options of external LSAs have been tested operationally: **type 1 metrics**, **type 2 metrics**, **forwarding addresses** and the **external route tag**.
- **Authentication.** The OSPF authentication procedure has been tested operationally.
- **Equal-cost multipath**. Operational deployments have included topologies with equal-cost, redundant paths.
- **Stub areas.** These have been deployed both in BARRNet and NSI.

6.5 Limitations of operational deployments

The following things have not been tested in an operational environment:

- **Multi-vendor deployments**. So far all deployments have used a single implementation. However, extensive interoperability testing of OSPF has been done (see Section 7.0 of this report).
- **Regular OSPF areas.** These have however been tested in all three rounds of the OSPF interoperability testing.
- **Virtual links.** These have however been tested in OSPF's interoperability testing.
- **Non-broadcast networks**. However, OSPF interoperability testing has been performed over X.25 networks.
- **TOS routing.** However, this has been tested in OSPF's interoperability testing.

6.6 Conclusions

All basic features of the OSPF protocol have been exercised. Very large OSPF link state databases (e.g., BARRNet's OSPF system) have been deployed, providing a thorough test of OSPF's database synchronization mechanisms. No OSPF protocol problems have been found in operational deployments.

Most of the hassles in operation deployments has to do with the OSPF/RIP interchange. Many of these issues have been ironed out on the ospf-tests mailing list (see Section 2.0). However, the interaction between OSPF, RIP, and EGP continues to be an active area of research.

7.0 Interoperability Testing

There have been three separate OSPF V2 interoperability testing sessions. Five separate implementations have participated in at least one session: implementations from the companies 3com, ACC, Proteon and Wellfleet, and the publicly available implementation from the University of Maryland.

Each of the testing sessions is described in a succeeding section. For each session, the participants are identified, and the testing topologies are described along with the particular protocol features that were exercised. Any protocol problems that were encountered during the testing are also described. In addition, for the second and third rounds testing reports were sent to the ospf mailing lists. These reports are reproduced in this document.

There is quite a bit of commonality in the features that have been tested from session to session. There are several reasons for this commonality. First, in each testing session an attempt has been made to increase the size of the OSPF system under test. For example, the number of external routes imported has doubled each session. Secondly, the interoperability sessions have been debugging sessions as well as protocol sessions. Many things tested in the third round were to verify that implementations had successfully fixed problems found in earlier sessions. A brief overview of the testing session is presented in the following table:

TABLE 4. OSPF interoperability testing at a glance.

Site	Date	# Routers	# Externals	Implementations
Proteon	9/25/90-9/29/90	6	20-30	3com, Proteon, Wellfleet
SURAnet	12/17/90-12/21/90	10	96	3com, ACC, Proteon, Wellfleet
3com	2/4/91-2/8/91	16	400	3com, ACC, Proteon, Wellfleet, UMD

For more information on the interoperability testing, the following people can be contacted: Fred Baker [fbaker], Rob Coltun [rcoltun], Dino Farinacci [dino], Jonathan Hsu [jhsu], John Moy [jmoy], and William Streilein [bstreile].

7.1 Testing methodology

In the interoperability tests, the routers have been interconnected using ethernet, serial lines (PPP and proprietary), X.25 and 802.5 token ring. Monitoring of the routers has been done through connecting terminals (either directly or via telnet) to the router consoles. Each implementation has a different user interface, which makes monitoring somewhat difficult. As explained earlier in this document, there is now an OSPF MIB, which in the future will enable a common monitoring interface to all implementations.

In general, each implementation has an error logging capability, and this is often how problems are discovered. LAN protocol analyzers are also used to capture OSPF protocol packet exchanges that are causing problems. These packet traces are available for analysis either during of after the testing sessions.

Verification of routing was done through visual inspection of implementations' routing table and link state databases (via the console interface), and through network debugging tools such as "ping" and "traceroute".

7.2 First round (Proteon, 9/25/90 - 9/29/90)

The first round of OSPF protocol testing took place at Proteon Inc.'s Westborough facility, the week of September 25, 1990. Three implementations participated, from the vendors 3com, Proteon and Wellfleet.

There were two 3com routers, two Wellfleet routers and two Proteon routers available for testing. These routers were interconnected with ethernets and serial lines. External routes were imported from the Proteon company internet. In addition, during off hours we were able to connect the routers under test to the Proteon company internet, forming one fairly large OSPF system.

The testing at Proteon proceeded as follows:

- All routers were connected to a single ethernet. Then, as routers were taken up and down, the Designated Router election algorithm and the Database Description process were tested. Also OSPF's reliable flooding algorithm was tested in this configuration.
- Twenty to thirty external routes were imported into the OSPF system by a Proteon router (which was simultaneously running RIP). It was then verified that these external routes were installed into the router's routing tables.
- One of the 3com routers was configured to originate an OSPF default route. This tested OSPF default route processing, and also tested the behavior of the system when multiple routers were importing external routes.
- The OSPF system was split into areas. Both regular OSPF areas (non-stub) and stub areas were tested.
- The six routers under test were connected to the Proteon company internet (which was also running OSPF), forming an OSPF system of eighteen routers. This configuration was short-lived, due to a disagreement between the 3com and Proteon routers concerning how to represent an OSPF default route.

Unfortunately, incomplete records were kept of this testing, so that no maps of the testing configurations can be reproduced for this document.

7.2.1 Problems found in the First round testing

A couple of OSPF protocol/specification problems were uncovered in the first round of testing. First, it was noticed that there was a window in the Database Description process where concurrently flooded MaxAge advertisements could prevent database synchronization from completing. This required a change in the specification's handling of MaxAge LSAs.

Secondly, it was found that the OSPF specification did not specify how the Network Mask field should be set in external LSAs that were advertising the DefaultDestination. This was a minor problem, but caused difficulties because of assumptions made in one implementation on how the mask should be set.

7.3 Second round (SURAnet, 12/17/90 - 12/21/90)

The second round of OSPF protocol testing took place at SURAnet's College Park facility, the week of December 12, 1990. Four implementations participated, from the vendors 3com, ACC, Proteon and Wellfleet.

There were two 3com routers, two ACC routers, two Wellfleet routers and four Proteon routers available for testing. These routers were interconnected with ethernets, serial lines and token rings. External routes were imported from SURAnet by one of the Proteon routers.

The testing at SURAnet proceeded as follows. Initially nine routers were configured as a single backbone area, with six of the routers connected to a single ethernet. This is pictured in Figure 4. In this configuration, the **Designated Router transition** and **database synchronization** process were tested. Ninety-six external routes were imported from SURAnet to stress the flooding algorithm. By restarting the router that was importing the routes, the **flushing of advertisements** from the routing domain was tested. Additionally, **variable-length subnets** and OSPF's optional **TOS routing** capability were tested in this configuration.

Next the routers were configured into four separate **OSPF areas**, with each area directly connected to the OSPF backbone (which was a single ethernet). There were no virtual links in this configuration. Inter-area routing was tested, including having AS boundary routers internal to a non-backbone area. Also tested was the case where a single router was both an area border router and an AS boundary router.

For more details of the testing, see the "Official report of the Second Round Testing" listed below.

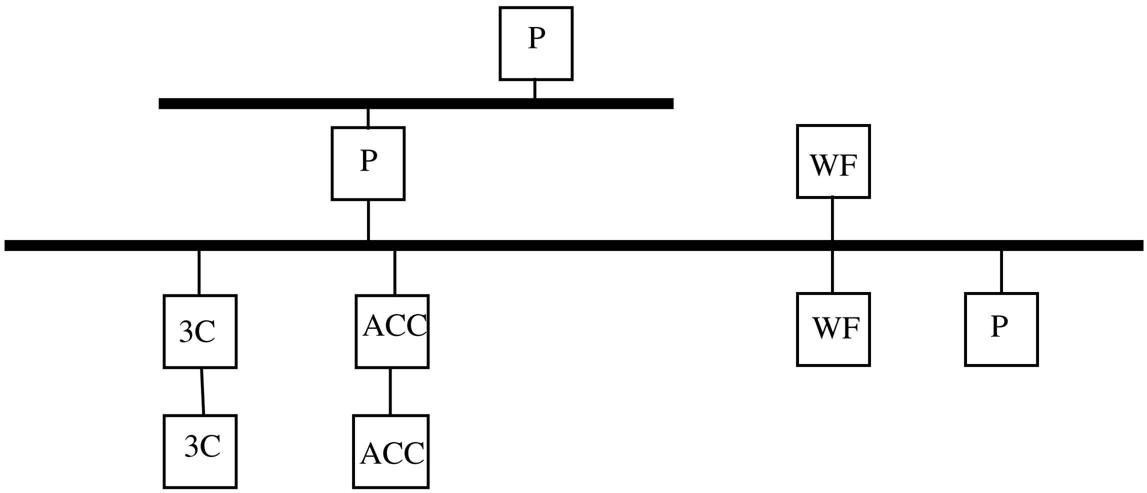

Figure 4: Initial configuration for the Second round testing

7.3.1 Official report of the Second round testing

The following report was sent to the ospf, ospf-tests, and router-requirements mailing lists after the second round of interoperability tests:

The second round of OSPF multi-vendor testing was done in College Park, Maryland the week of 12/17/90. The facilities were provided by SURAnet. Four major router vendors were present, Advanced Computer Communications (ACC), Proteon, 3Com, and Wellfleet. A press conference and presentation was provided for 3 different data communication magazine representatives.

Each vendor provided at least 2 routers. Each vendor had a device connected to a common Ethernet. This Ethernet was configured as the OSPF backbone area. The rest of the routers were attached to the various backbone routers via Ethernet, Token Ring, proprietary serial line, PPP serial line, and X.25 type media. The following test scenarios were performed and completed in the following order:

- Intra-area routing. All routers were configured to reside in the backbone area. Designated Router election was performed various number of times so each vendor could be designated router for a period of time. Packet data was captured on a Sniffer for later observation.
- Variable Length Subnet Masks. Some of the serial lines in the configuration were configured to be on the same IP network but with different subnet masks. It was observed that all routers stored routes for the different length subnets.

- Import SURAnet routes. The Proteon router was listening for RIP routes generated by the SURAnet routers. These routes were imported into our OSPF test system. 96 external link state advertisements were generated as a result. Many scaling type implementation problems surfaced for each vendor during this exercise.
- Type of Service generation. While the test setup was still configured as a single area, the 3Com router generated Type of Service link state advertisements. It was observed how the other vendor implementations reacted to it. Some problems were found.
- Inter-area routing. Multiple areas were configured. Common non-backbone areas were shared by Proteon and Wellfleet and by ACC and 3Com. It was observed that the correct Intra-area and Inter-area routes appeared in each router's routing table. At this point in the test setup, the Proteon router regenerated the 96 SURAnet routes into the configuration. It was observed that the routes were learned and propagated over area boundaries. Some problems occur at this point. More emphasis on this scenario will occur at the next round of testing.
- OSPF over X.25. A point-to-point link was connected between the Proteon router and the 3Com router. The X.25 packet level was configured to run over the link. OSPF was enabled over the link to verify that multi-vendor OSPF over X.25 was performed. Both of these routers were in the same area.
- MaxAge advertisements. Link state advertisements were flushed from the routing domain using the MaxAge procedure. We verified that all routers removed the advertisements from their databases, after they were properly acknowledged by the flooding procedure. Some problems were found in this test, although not nearly as many as in the first round of testing.

Each vendor agreed that this sort of testing was extremely valuable and that it should occur again. 3Com has offered for the third round of testing to occur in Santa Clara sometime in February. 3Com will encourage other OSPF implementations to join in the testing. Items that will be tested are:

- Intra-area routing with loops (as well as equal-cost multipath).
- Inter-area testing including Stub and Transit area support, with both Intra-area and Inter-area loops.
- Virtual link testing in the looped Inter-area configuration.
- RIP/OSPF route interchange including testing forwarding address capability in external link state advertisements.
- EGP/OSPF router interchange including the use of the route tag field in external link state advertisements.
- More than two routers connected to an X.25 network. We would like to test OSPF's non-broadcast multi-access capabilities by attaching more than two vendor's routers to an X.25 packet switch.

- Several vendors running OSPF and RIP simultaneously. This will further test the OSPF/RIP interactions.
- Test processing of links with cost LSInfinity. These links should be treated as unreachable.

Furthermore, we hope that in future rounds of testing an OSPF MIB would allow us to also use a network management station to gather test data.

In summary, the stability of implementations were better this time more so than the first round of testing. No problems with the protocol design were encountered. The exchange of ideas and the cooperation among implementors that occurred during this test effort, continues the spirit that OSPF is truly an open protocol.

7.3.2 Problems found in the Second round testing

No problems were found in the OSPF protocol during the second round of testing.

7.4 Third round (3com, 2/4/91 - 2/8/91)

The third round of OSPF protocol testing took place at 3com's Santa Clara facility, the week of February 4, 1991. Five implementations participated, from the vendors 3com, ACC, Proteon and Wellfleet and the publicly available University of Maryland implementation (running on a SUN workstation).

There were five 3com routers, four ACC routers, three Wellfleet routers, three Proteon routers and the UMD Sun workstation available for testing (giving a total of 16 routers available). These routers were interconnected with ethernets, serial lines and X.25. External routes were imported from BARRNet by one of the 3com routers.

The initial testing configuration is shown in Figure 5. Three areas were configured, along with a non-contiguous backbone. The backbone was then joined by configuring two virtual links. In this configuration the following OSPF functionality was tested: **inter-area routing** and **virtual links**.

The system was then reconfigured so that twelve of the routers were connected to a single ethernet. This configuration is pictured in Figure 6. By bringing routers up and down, this configuration tested **Designated Router election**, **database synchronization** and **reliable flooding**. To see how this functionality, and also the implementations, scale, 400 external routes were imported from BARRNet.

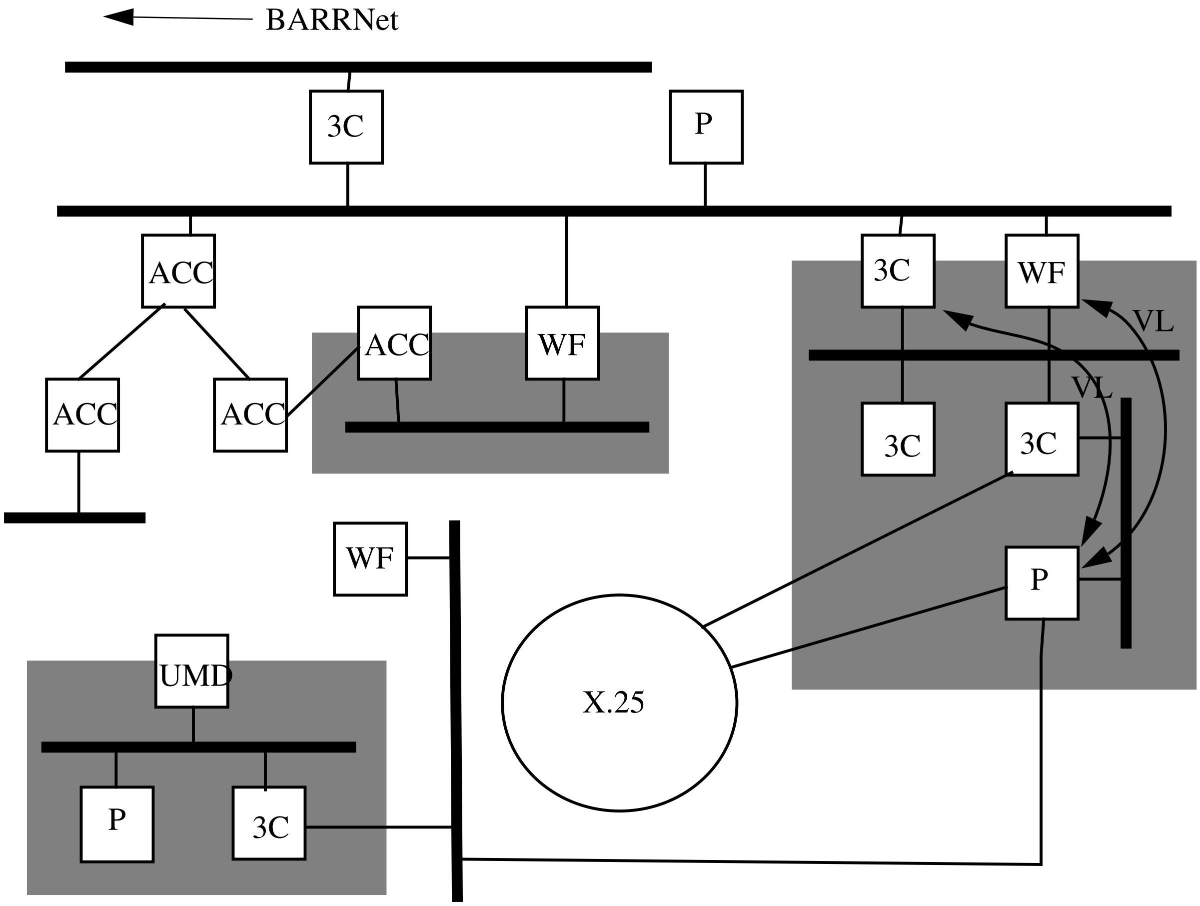

Figure 5: Third round testing: Initial configuration

7.4.1 Official report of the Third round testing

The following report was sent to the ospf, ospf-tests, and router-requirements mailing lists after the third round of interoperability tests:

The third round of OSPF interoperability testing was held at 3com Corporation in Santa Clara the week of February 4-8. Four router vendors came to the testing: 3com, ACC, Proteon and Wellfleet. In addition, Rob Coltun brought the University of Maryland implementation, which was run on a Sun Workstation.

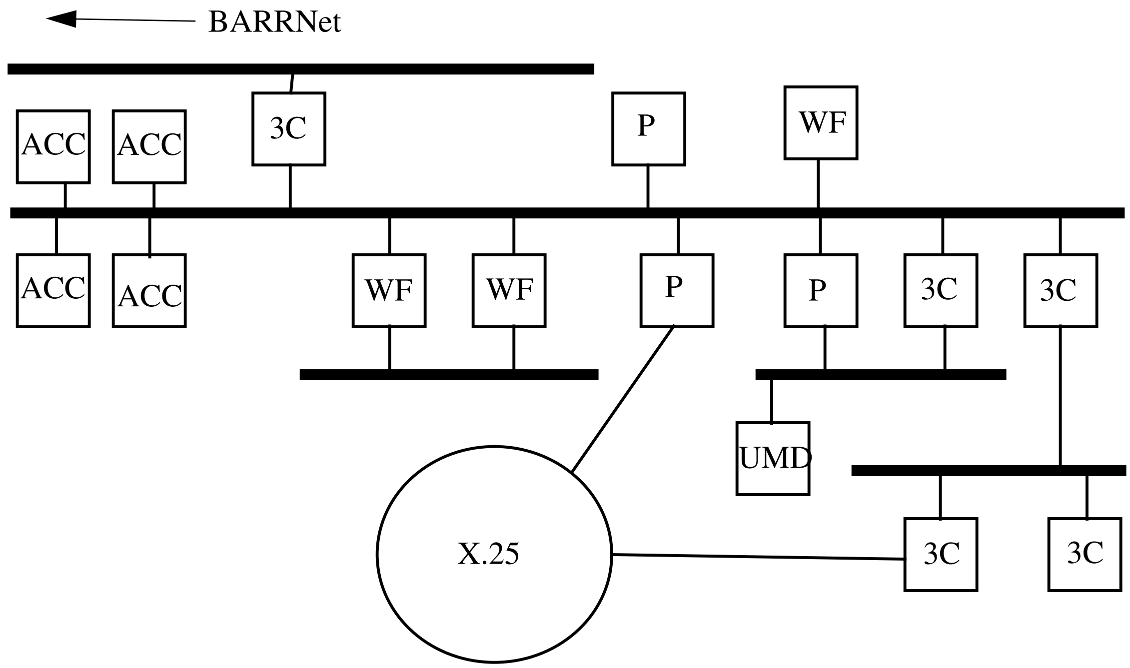

Figure 6: Third round testing: second configuration

Testing was performed over ethernet, point-to-point networks (using PPP) and X.25. In all we had 16 routers available: five 3com routers, four ACC routers, three Proteon routers, three Wellfleet routers and Rob's SUN. We also were able to import external routes from BARRNet.

Specific tests performed included the following:

- Initially we configured the routers into three separate areas and a physically disconnected backbone. The backbone was then reconnected through configuration of several virtual links. These tests verified the generation and processing of summary link advertisements, as well as the operation of virtual links.
- We connected multiple routers to an X.25 packet switch, testing OSPF's non-broadcast network capability. OSPF was successfully run over the X.25 network, using routers that were both DR eligible and DR ineligible. Some problems were encountered, but they all involved running IP over X.25 (i.e., they were not X.25 specific).
- We also connected a 3com router, Proteon router, and Rob's SUN to an ethernet, and then treated the ethernet as a non-broadcast network. This allowed us to connect Rob's SUN into the rest of the routing domain without installing the IP multicast modifications to the SUN kernel. It also further tested the OSPF's non-broadcast network capability.

- We then reconfigured the OSPF system so that all but three of the routers were connected to a single ethernet. This tested the Designated Router functionality (12 routers were synchronizing with the DR). We then also tested the DR election algorithm, by selectively restarting the DR, or sometimes both the DR and the Backup DR. This also tested OSPF's Database Description process.
- In this configuration, we then imported 400 external routes from BARRNet (one of the 3com routers ran both OSPF and EGP). Some problems were encountered in implementations' buffer allocation strategies, and problems were also found in the way implementations avoid IP fragmentation. But overall, this system was fairly stable.

The following problems we found in the OSPF specification:

- The specification should say that the "Network mask" field should not be verified in OSPF Hellos received over virtual links.
- The specification should say that on multi-access networks neighbors are identified by their IP address, and on point-to-point networks and virtual links by their OSPF Router ID. This eliminates confusion when, for example, a router is restarted and comes up with the same IP address but a different Router ID.

Thanks to 3com for providing the testing facility, cables. terminals, X.25 switch. etc. Also thanks to Vince Fuller of BARRNet for helping us import the BARRNet routes.

7.4.2 Problems found in the Third round testing

A couple of specification/protocol problems were found in the third round of interoperability testing. First, it was noticed that the specification did not specify the setting of the Network Mask field in Hellos sent over virtual links. This caused some initial difficulty in bringing up virtual links between routers belonging to different vendors. Secondly, it was noticed that the specification was not strict enough in defining how OSPF neighbors are identified on multi-access networks. This caused difficulties in one implementation when another vendor's router was restarted with the same IP address but a different OSPF Router ID. This is discussed more fully in the above "Official Report of the Third Round Testing".

7.5 Overall: Features tested

All significant protocol features and mechanisms have been tested in the three rounds of interoperability testing. In particular, the following basic pieces of the protocol have been tested:

- **Designated Router election.** With as many as thirteen routers attached to a single LAN, the election of Backup Designated Router and Designated router was verified by bringing routers up and down, singly and in pairs.

- **Adjacency bringup**. The Database Description process was verified, with databases having over 400 LSAs. Adjacency bringup was also verified in times when flooding was taking place simultaneously.
- **Reliable flooding**. It was verified that OSPF's flooding algorithm maintains database synchronization, both in the presence of loops in the topology, and with large databases (over 400 LSAs).
- **Flushing advertisements from routing domain**. OSPF's procedure for flushing old or unreachable LSAs from the routing domain was verified, both in the presence of topology loops and with many LSAs being flushed at once. This is also referred to as OSPF's MaxAge procedure.
- **OSPF routing hierarchy.** The OSPF four level routing hierarchy: intra-area, inter-area. type 1 externals and type 2 externals was tested.
- **Import of external routing information**. The importing of external routes has been tested, with as many as 400 imported at once. Also, the varying options in external LSAs has been tested: type 1 or type 2 metrics and forwarding addresses.escribe all options. In addition, test setups were utilized having AS boundary routers both internal to non-backbone areas and also being simultaneously area border routers.
- **Running protocol over various network types.** In the test setups, OSPF has been run over ethernet, point-to-point serial lines (both PPP and proprietary)**,** 802.5 token ring and X.25.
- **Non-broadcast, multi-access networks.** OSPF has been tested over X.25. Some testing was also done treating ethernet as a non-broadcast network. Two separate situations were tested: when all routers attached to the non-broadcast network were DR-eligible, and when only some of them were.
- **Authentication.** OSPF's authentication procedure was tested for the two current authentication types.
- **Equal-cost multipath.** Much of the testing was done in configurations with redundant paths, and equal-cost multipath was verified through examination of implementations' routing tables.
- **Variable-length subnet masks.** It was verified that implementations paid attention to the network mask field present in OSPF LSAs.

Testing was also performed on the following pieces of OSPF's Area functionality:

- **Extent of advertisements.** It was verified that all advertisements except external LSAs were flooded throughout a single area only.
- **Inter-area routing.** The generation and processing of summary link LSAs was tested. Also tested were configurations having multiple area border routers attaching to a single area.
- **Virtual links.**

The following OSPF options were also tested:

- **TOS routing.** The interplay between TOS-capable and non-TOS-capable routers was tested, by configuring TOS-specific metrics in the only implementation (3com) supporting TOS routing.
- **Stub areas.** OSPF's stub area functionality was verified.

7.6 Testing conclusions

The interoperability testing has proven to be very valuable. Many bugs were found (and fixed) in the implementations. Some protocol problems were found (and fixed), and gray areas of the specification were cleared up. Implementations have also been "bullet-proofed" in order to deal with the unexpected behavior of other implementations. All participants in the testing now understand the maxim "be conservative in what you generate, and liberal in what you accept" (if they didn't already).

7.7 Future work

The one thing that has gone mostly untested at the interoperability sessions is the interaction between OSPF and other routing protocols (such as RIP and EGP). Each interoperability session generally had a router running multiple routing protocols in order to import external routing information into the OSPF system. However, simultaneously running multiple routing protocols between different vendors' routers has not been tested.

Each vendor has developed a slightly different architecture for the exchange of routing information between differing routing protocols. As the OSPF field testing has also shown, this exchange of routing information is an area of ongoing work and a candidate for future standardization.

8.0 Simulation

The OSPF protocol has been simulated by the Distributed Systems Research Group at the University of Maryland Baltimore County (UMBC). The two principal investigators for the OSPF simulation project are Dr. Deepinder P. Sidhu of UMBC and Rob Coltun. They have been aided by three graduate students: S. Abdallah, T. Fu and R. Nair. This section attempts to summarize their simulation setup and results. For more information, contact the Distributed Systems Research Group at the following address:

Dr. Deepinder P. Sidhu
Department of Computer Science
University of Maryland Baltimore County
Baltimore, MD 21228
email: sidhu@umbc3.umbc.edu

A demo was given of their OSPF simulation at the March 4-8, 1991 IETF in St. Louis. Details of the demo should be available in the IETF proceedings.

8.1 Simulator setup

The Distributed System Research Group uses a significantly enhanced version of the MIT network simulator. The simulator is event driven, and contains support for both point-to-point networks and ethernet links. It can simulate characteristics of both packet switches and hosts, and can simulate internet behavior under various types of data traffic load (e.g., Poisson, normal, exponential and uniform distributions). This latter ability could be used, for example, to simulate how a routing protocol works in a congested internet. Specific network topologies can be input into the simulator, or pseudo-random network topologies can be generated. Packet loss rates can also be simulated.

To simulate OSPF, Rob Coltun's OSPF implementation was incorporated into the simulator, essentially unchanged.

The output of the simulator can be displayed in a graphical manner (it uses X windows). Any variable in the implementation under test can be monitored. In addition, statistical reports can later be produced from logging files produced during the simulation runs.

8.2 Simulation results

The OSPF simulation has been run using the following topologies:

- The two sample topologies in the OSPF specification (Figure 2 and Figure 6 in [2]). The first of these topologies shows an Autonomous System without areas, the second the same AS with three areas and a virtual link configured.
- The 19-node hub topology from [7].
- A large network of over 50 nodes, all attached to a single ethernet.
- A large network of over 50 nodes, containing both ethernets and serial lines, pseudo randomly generated.

In these topologies, the correctness of the OSPF database synchronization was verified. This was done through examination of the nodes' OSPF link state databases and the nodes' routing tables. The implementation of multiple OSPF areas was also tested. Also, database convergence time was analyzed as a function of the network components' link speeds.

Also, some formal analysis of the OSPF protocol was undertaken. The neighbor and interface state machines were analyzed. In addition, the Designated Router election algorithm was verified for certain sets of initial conditions.

9.0 Reference Documents

The following documents have been referenced by this report:

[1] Moy, J., “The OSPF Specification”, RFC 1131, October 1989.

[2] Moy, J., “OSPF Version 2”, RFC 1247, July 1991.

[3] Coltun, R. and Baker, F., “OSPF Version 2 Management Information Base”, RFC 1248, July 1991.

[4] Reynolds, J. and Postel, J., “Assigned Numbers’, RFC1060, March 1990.

[5] Corporation for National Research Initiatives, “Proceedings of the Thirteenth Internet Engineering Task Force”, Cocoa Beach, Florida, April 11-14, 1989.

[6] Corporation for National Research Initiatives, “Proceedings of the Sixteenth Internet Engineering Task Force”, Florida State University, February 6-9, 1990.

[7] Gardner, M., et al., “Type-of-service routing: modeling and simulation,” Report 6364, BBN Communications Corporation, January 1987.

[8] Corporation for National Research Initiatives, “Proceedings of the Seventeenth Internet Engineering Task Force”, Pittsburgh Supercomputing Center, May 1-4, 1990.

[9] Corporation for National Research Initiatives, “Proceedings of the Eighteenth Internet Engineering Task Force”, University of British Columbia, July 30-August 3, 1990.

10.0 People

The following people have contributed information to this report and can be contacted for further information:

TABLE 5. People references in this report

Reference Tag	Name	Affiliation	email
bstreile	William Streilein	Wellfleet Communications	bstreile@wellfleet.com
dino	Dino Farinacci	3com Corporation	dino@buckeye.esd.3com.com
fbaker	Fred Baker	Advanced Computer Communications	fbaker@acc.com
jeff	Jeffrey Burgan	Sterling Software	jeff@nsipo.nasa.gov
jhsu	Jonathan Hsu	Wellfleet Communications	jhsu@wellfleet.com
jmoy	John Moy	Proteon, Inc.	jmoy@proteon.com
kannan	Kannan Varadhan	OARnet	kannan@oar.net
medin	Milo Medin	Sterling Software	medin@nsipo.nasa.gov
rcoltun	Rob Coltun	University of Maryland	rcoltun@umd5.umd.edu
rrv	Ross Veach	University of Illinois	rrv@seka.cso.uiuc.edu
vaf	Vince Fuller	BARRNet	vaf@valinor.stanford.edu

Security Considerations

The OSPF protocol's security architecture is described in Section 4.0.

Author's Address

John Moy
Proteon Inc.
2 Technology Drive
Westborough, MA 01581

Phone: (508) 898-2800
Email: jmoy@proteon.com