OSPF Design and Interoperability Recommendations for Catalyst 6500 and OSA-Express Environments

®

Table of Contents

1. Introduction

1.1 ABC Corporation

ABC Corporation is a large financial institution operating in the North Eastern corridor of the USA. They have a large IBM® S/390® Data Center installation. With the movement towards IP connectivity and the availability of the Catalyst 6500 and IBM Open Systems Adapter (OSA) Express technologies, ABC Corp. approached Cisco Systems, Inc. ("Cisco") and IBM with a request for assistance. This request was to provide design and interoperability recommendations for an IP based infrastructure, encompassing Catalyst 6500 and IBM ^ ® OSA-Express environments.

This white paper documents the results of this initiative. It should be noted that the implementation of any component illustrated in this paper neither constitutes nor represents the respective component's only possible implementation.

It is beyond the scope of this paper to exhaustively document each component's full spectrum of supported features, implementation configurations and options, and recommended deployment scenarios. Therefore, it is the reader's responsibility to fully understand the technologies described within this document and to determine this paper's applicability to their particular environment prior to applying or implementing any recommendations cited within this document.

The information contained in this presentation has not been submitted to any formal IBM or Cisco test and is distributed "AS IS". While each item may have been reviewed by IBM and Cisco for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. The use of this information or the implementation of any techniques described herein is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. Customers attempting to adapt these techniques to their own environments do so at their own risk. Neither IBM, nor Cisco makes any claims of this document's applicability to any specific environment, either similar or different than those described within this paper.

1.2 Scope and Rationale

The scope of the document seeks to:

1. Introduce an OSPF design framework for Catalyst 6500 / OSA-Express environments.
2. Describe the design principles and best practices implemented to achieve this design
3. Describe a variety of failure scenarios and how the network reacts to these failures

This document focuses on the connectivity between the Catalyst 6500 and OSA-Express components. For testing purposes an additional Cisco 7500 router was used as a 'WAN Edge' router. The 'WAN Edge' is not a primary focus of the document. Both Cisco and IBM did however make numerous recommendations for this layer, but they are not documented here.

The first area of interest is the actual hardware and software components used in the tests.

1.3 Hardware and Software Inventory

The following is a list of the Cisco and IBM equipment used.

1.3.1 Cisco Systems – Equipment Inventory

Two Catalysts 6500 (with MSFC-1 / PFC combination) and a Cisco 7507 router were deployed in these tests. The following naming convention was thus adopted for these components; '6500-top', '6500-bottom', 'msfc-top', 'msfc-bottom' and '7507-access' respectively. Figure 1. illustrates the layout of the equipment.

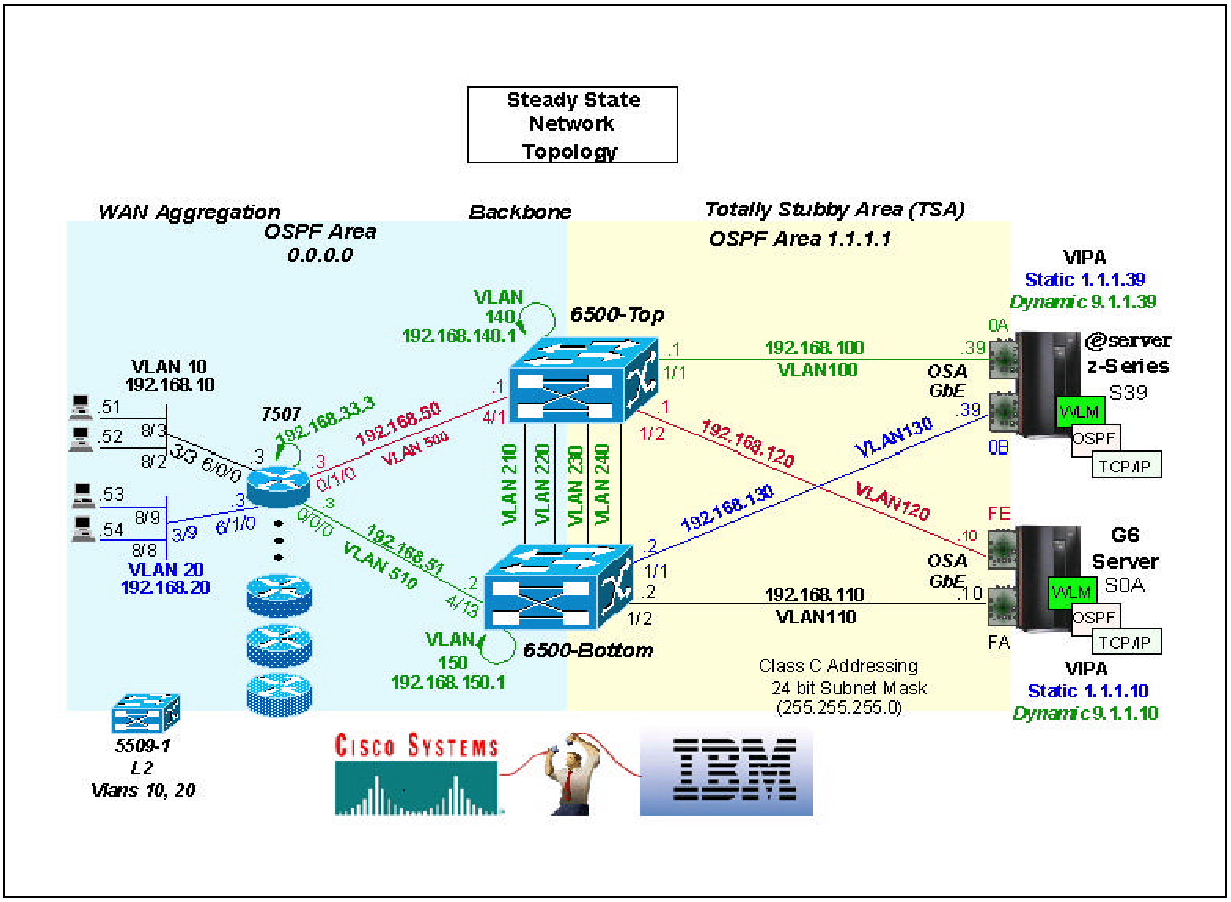

Figure 1 - Steady State Network

The 'show version' command highlights in detail the individual hardware components plus software versions utilized :

6500-top

```
 6500-top> (enable) sh ver

WS-C6506 Software, Version NmpSW: 5.5(3)
Copyright (c) 1995-2000 by Cisco Systems, Inc.
NMP S/W compiled on Sep  8 2000, 17:39:10

System Bootstrap Version: 5.3(1)

Hardware Version: 2.0  Model: WS-C6506  Serial #: TBA03370628
```

```
Mod Port Model              Serial #    Versions
--- ---- ------------------ ----------- ------------------------------------
1   2    WS-X6K-SUP1A-2GE   SAD04260K3X Hw : 3.2
                                        Fw : 5.3(1)
                                        Fw1: 5.1(1)CSX
                                        Sw : 5.5(3)
                                        Sw1: 5.5(3)
         WS-F6K-PFC         SAD042702DY Hw : 1.1
4   48   WS-X6248-RJ-45     SAD03358793 Hw : 1.1
                                        Fw : 4.2(0.24)VAI78
                                        Sw : 5.5(3)
5   48   WS-X6248-RJ-45     SAD03294181 Hw : 1.1
                                        Fw : 4.2(0.24)VAI78
                                        Sw : 5.5(3)
6   8    WS-X6408-GBIC      SAD03361671 Hw : 2.3
                                        Fw : 4.2(0.24)VAI78
                                        Sw : 5.5(3)
15  1    WS-F6K-MSFC        SAD04260LDX Hw : 1.4
                                        Fw : 12.1(2)E,
                                        Sw : 12.1(2)E,

       DRAM                    FLASH                   NVRAM
Module Total   Used    Free    Total   Used    Free    Total Used  Free
------ ------- ------- ------- ------- ------- ------- ----- ----- -----
1       65408K  38177K  27231K  16384K    769K  15615K  512K  250K  262K

Uptime is 0 day, 2 hours, 9 minutes
```

msfc-top

```
msfc-top#sh ver
Cisco Internetwork Operating System Software
IOS (tm) MSFC Software (C6MSFC-JS-M), Version 12.1(2)E, EARLY DEPLOYMENT
RELEASE SOFTWARE (fc1)
Copyright (c) 1986-2000 by Cisco Systems, Inc.
Compiled Mon 19-Jun-00 18:38 by linda
Image text-base: 0x60008900, data-base: 0x611B4000

ROM: System Bootstrap, Version 12.0(3)XE, RELEASE SOFTWARE
BOOTFLASH: MSFC Software (C6MSFC-BOOT-M), Version 12.1(2)E, EARLY DEPLOYMENT
RELEASE SOFTWARE (fc1)

msfc-top uptime is 2 hours, 7 minutes
System returned to ROM by power-on
Running default software

cisco Cat6k-MSFC (R5000) processor with 114688K/16384K bytes of memory.
Processor board ID SAD04260LDX
R5000 CPU at 200Mhz, Implementation 35, Rev 2.1, 512KB L2 Cache
Last reset from power-on
Bridging software.
X.25 software, Version 3.0.0.
SuperLAT software (copyright 1990 by Meridian Technology Corp.).
TN3270 Emulation software.
8 Virtual Ethernet/IEEE 802.3  interface(s)
123K bytes of non-volatile configuration memory.
4096K bytes of packet SRAM memory.

16384K bytes of Flash internal SIMM (Sector size 256K).
Configuration register is 0x2102
```

6500-bottom

```
 6500-bottom> (enable) sho ver
 WS-C6506 Software, Version NmpSW: 5.5(3)
 Copyright (c) 1995-2000 by Cisco Systems, Inc.
 NMP S/W compiled on Sep  8 2000, 17:39:10

System Bootstrap Version: 5.3(1)

Hardware Version: 2.0  Model: WS-C6506  Serial #: TBA03380449

Mod Port Model               Serial #    Versions
--- ---- ------------------- ----------- -------------------------------------
1   2    WS-X6K-SUP1A-2GE    SAD043308YZ Hw : 3.2
                                         Fw : 5.3(1)
                                         Fw1: 5.1(1)CSX
                                         Sw : 5.5(3)
                                         Sw1: 5.5(3)
         WS-F6K-PFC          SAD04310HXA Hw : 1.1
3   4    WS-X6302-MSM        SAD03350729 Hw : 3.0
                                         Fw : 12.0(1a)WX5(6g),
                                         Sw : 12.0(1a)WX5(6g),
4   48   WS-X6248-RJ-45      SAD03363564 Hw : 1.1
                                         Fw : 4.2(0.24)VAI78
                                         Sw : 5.5(3)
5   48   WS-X6248-RJ-45      SAD03311010 Hw : 1.1
                                         Fw : 4.2(0.24)VAI78
                                         Sw : 5.5(3)
6   8    WS-X6408-GBIC       SAD03361741 Hw : 2.3
                                         Fw : 4.2(0.24)VAI78
                                         Sw : 5.5(3)
15  1    WS-F6K-MSFC         SAD043309JY Hw : 1.4
                                         Fw : 12.1(2)E,
                                         Sw : 12.1(2)E,

       DRAM                    FLASH                   NVRAM
Module Total   Used    Free    Total   Used    Free    Total Used  Free
------ ------- ------- ------- ------- ------- ------- ----- ----- -----
1       65408K  38170K  27238K  16384K  13080K   3304K  512K  279K  233K

Uptime is 0 day, 2 hours, 8 minutes
6500-bottom> (enable)
```

msfc-bottom

```
 msfc-bottom#sho ver
 Cisco Internetwork Operating System Software
 IOS (tm) MSFC Software (C6MSFC-JS-M), Version 12.1(2)E, EARLY DEPLOYMENT
 RELEASE SOFTWARE (fc1)
 Copyright (c) 1986-2000 by Cisco Systems, Inc.
 Compiled Mon 19-Jun-00 18:38 by linda
 Image text-base: 0x60008900, data-base: 0x611B4000

ROM: System Bootstrap, Version 12.0(3)XE, RELEASE SOFTWARE
BOOTFLASH: MSFC Software (C6MSFC-BOOT-M), Version 12.1(2)E, EARLY DEPLOYMENT
RELEASE SOFTWARE (fc1)

msfc-bottom uptime is 2 hours, 8 minutes
System returned to ROM by power-on
Running default software

cisco Cat6k-MSFC (R5000) processor with 114688K/16384K bytes of memory.
```

```
Processor board ID SAD043309JY
R5000 CPU at 200Mhz, Implementation 35, Rev 2.1, 512KB L2 Cache
Last reset from power-on
Bridging software.
X.25 software, Version 3.0.0.
SuperLAT software (copyright 1990 by Meridian Technology Corp.).
TN3270 Emulation software.
8 Virtual Ethernet/IEEE 802.3  interface(s)
123K bytes of non-volatile configuration memory.
4096K bytes of packet SRAM memory.

16384K bytes of Flash internal SIMM (Sector size 256K).
Configuration register is 0x102
```

7507-ACCESS

```
 7507-ACCESS#sh ver
 Cisco Internetwork Operating System Software
IOS (tm) RSP Software (RSP-JSV-M), Version 12.0(7)T,  RELEASE SOFTWARE (fc2)
Copyright (c) 1986-1999 by Cisco Systems, Inc.
Compiled Mon 06-Dec-99 19:40 by phanguye
Image text-base: 0x60010908, data-base: 0x61356000

ROM: System Bootstrap, Version 11.1(8)CA1, EARLY DEPLOYMENT RELEASE SOFTWARE
(fc1)
BOOTFLASH: GS Software (RSP-BOOT-M), Version 11.1(22)CA, EARLY DEPLOYMENT
RELEASE SOFTWARE (fc1)

7507-ACCESS uptime is 23 hours, 25 minutes
System returned to ROM by reload
System image file is "slot1:rsp-jsv-mz_120-7_T"

cisco RSP4 (R5000) processor with 262144K/2072K bytes of memory.
R5000 CPU at 200Mhz, Implementation 35, Rev 2.1, 512KB L2 Cache
Last reset from power-on
G.703/E1 software, Version 1.0.
G.703/JT2 software, Version 1.0.
X.25 software, Version 3.0.0.
SuperLAT software (copyright 1990 by Meridian Technology Corp.).
Bridging software.
TN3270 Emulation software.
Chassis Interface.
2 CIP2 controllers (6 IBM Channels).
2 VIP2 R5K controllers (4 FastEthernet).
4 FastEthernet/IEEE 802.3 interface(s)
6 IBM channel interface(s)
123K bytes of non-volatile configuration memory.

20480K bytes of Flash PCMCIA card at slot 0 (Sector size 128K).
20480K bytes of Flash PCMCIA card at slot 1 (Sector size 128K).
8192K bytes of Flash internal SIMM (Sector size 256K).

Slave in slot 3 is running Cisco Internetwork Operating System Software
IOS (tm) RSP Software (RSP-DW-M), Version 12.0(7)T,  RELEASE SOFTWARE (fc2)
Copyright (c) 1986-1999 by Cisco Systems, Inc.
Compiled Mon 06-Dec-99 19:43 by phanguye
Slave: Loaded from system
Slave: cisco RSP4 (R5000) processor with 262144K bytes of memory.
```

This completes the Cisco equipment inventory. Attention now turns to the IBM equipment used :

1.3.2 IBM Corporation – Equipment Inventory

IBM utilized two S/390s for these tests; one z-900 server and one model 9672 Generation 6 (G6) server respectively. A logical partition (LPAR) was used on each S/390 and both were configured to participate in the same Parallel Sysplex®. Both IBM Servers utilized OS/390® V2 R10.

Two Gigabit Ethernet OSA-Express interfaces were installed in each S/390 Server. This facilitated redundancy and load balancing. To enhance the S/390's high availability characteristics, both Static and Dynamic Virtual IP Addresses (VIPA) were defined in each S/390 server's TCP/IP profile.

It should be noted that while the OSA-Express features were dedicated to specific S/390 servers, the customer's (ABC Corp.) primary focus was not in exploring a shared OSA-Express implementation.

Therefore, deploying dedicated OSA -Express features within this paper neither implies nor indicates that an OSA-Express sharing restriction exists. In fact, other than a few minor configuration additions to each S/390 server's TCP/IP and OSPF profiles, there are no known restrictions that would have prevented the OSA-Express features from being successfully shared within the scope of this paper.

The use of VIPA and Dynamic VIPA are key features in this environment. These features are now considered in more detail to aid one's understanding :

1.4 VIPA and Dynamic VIPA

Traditionally, an IP address is associated with each end of a physical link, or each point of access to a shared-medium LAN. IP addresses are unique across the entire visible network. The majority of IP hosts have a single point of physical attachment to the network, but some hosts, particularly large server hosts, have more than one physical link into the network.

Within the IP routed network, failure of any intermediate link or physical adapter disrupts end user service if there is not an alternate path through the routing network. Routers can route IP traffic around failures of intermediate links in such a way that the failures are not visible to the end applications or IP hosts.

The Virtual IP Address (VIPA) removes the adapter as a single point of failure by providing an IP address that is associated with a S/390 server's TCP/IP stack without associating it with a specific physical network attachment, such as an OSA-Express feature. Therefore, since a VIPA has no single physical network attachment associated with it, it is always active and never experiences a physical failure as long as the TCP/IP stack is still active.

To the routed network, a VIPA appears to be a host address (/32 mask) that is indirectly attached to a S/390 Server. When a packet with a VIPA destination reaches the S/390's TCP/IP stack, the IP layer recognizes the address and passes it to the protocol layer in the stack.

Similar to the failure of the physical interface, the failure of a VIPA can be extended to include the failure of either the S/390 server's TCP/IP stack or the failure of a S/390 server. For these failure scenarios, a VIPA address needs to 'move' to a backup S/390 server's TCP/IP stack and the routes to the VIPA address need to be re-advertised so that clients can transparently connect to the backup TCP/IP stack. This process is known as VIPA Takeover Support.

VIPA Takeover has been improved with the introduction of Dynamic Virtual IP Address (DVIPA) in OS/390 V2R8 and Distributed Dynamic Virtual IP Address (Distributed DVIPA) in OS/390 V2R10 . The DVIPA function improves VIPA Takeover by providing a system programmer with the ability to plan for system outages by identifying backup S/390 servers that will dynamically take over responsibility for a VIPA without either operator intervention or external automation. Furthermore, as soon as the failed TCP/IP stack has recovered, the "taken over" VIPA address will immediately be "taken back" automatically and transparently by the recovered TCP/IP stack.

Section 4.5 illustrates this process in further detail.

Having introduced the hardware and software components of this project, attention now turns to the 'Steady State' Network Topology that was implemented.

2. Network Topology and Routing Tables

This section introduces the network topology and the steady-state routing tables for all devices in the network. Figure 1 illustrates the network topology. This diagram reflects the logical design used.

Please note that complete configurations are provided in Appendix A. One's understanding of the routing tables will be reinforced via the discussions in Section 3 and Section 4.

Some key points to note at this stage are :

- Layer 3 oriented OSPF based network
- Use of multiple OSPF Areas (Area 0 and Area 1)
- Catalyst 6500 MSFCs act as Area Border Routers between Area 0 and Area 1
- OSA-Express environment (Area 1) defined as a Totally Stubby Area
- Catalyst 6500 MSFCs acting as Area Border Routers (ABR) providing load balancing / redundant exit points from Area 1.
- Use of VIPA and Dynamic VIPA on the S/390s.
- Minimal Spanning Tree deployment

Steady state routing tables were derived from the following devices :

- **msfc-top**
- **msfc-bottom**
- **S/390 Server S39**
- **S/390 Server S0A**
- **7500-access**

Please reference Figure 1 to consolidate understanding of the routing tables below. Important routes are highlighted for convenience.

2.1 Routing Table - 'msfc-top'

```
msfc-top#sh ip ro

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

     1.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
O       1.1.1.0/24 [110/4] via 192.168.100.39, 01:13:10, Vlan100
                   [110/4] via 192.168.120.10, 01:13:10, Vlan120
O       1.1.1.10/32 [110/4] via 192.168.120.10, 01:13:10, Vlan120
O       1.1.1.39/32 [110/4] via 192.168.100.39, 01:13:10, Vlan100
C    192.168.120.0/24 is directly connected, Vlan120
C    192.168.210.0/24 is directly connected, Vlan210
C    192.168.240.0/24 is directly connected, Vlan240
     192.168.150.0/32 is subnetted, 1 subnets
O       192.168.150.1 [110/2] via 192.168.210.2, 00:03:44, Vlan210
                      [110/2] via 192.168.220.2, 00:03:44, Vlan220
```

```
O    192.168.110.0/24 [110/2] via 192.168.240.2, 01:13:10, Vlan240
                      [110/2] via 192.168.230.2, 01:13:10, Vlan230
C    192.168.230.0/24 is directly connected, Vlan230
O    192.168.10.0/24 [110/2] via 192.168.50.3, 00:03:46, Vlan500
O    192.168.130.0/24 [110/2] via 192.168.240.2, 01:13:12, Vlan240
                      [110/2] via 192.168.230.2, 01:13:12, Vlan230
     192.168.140.0/32 is subnetted, 1 subnets
C       192.168.140.1 is directly connected, Loopback0
C    127.0.0.0/8 is directly connected, EOBC0/6
     9.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
O       9.1.1.10/32 [110/4] via 192.168.120.10, 01:13:12, Vlan120
O       9.1.1.0/24 [110/4] via 192.168.100.39, 01:13:12, Vlan100
                   [110/4] via 192.168.120.10, 01:13:12, Vlan120
O       9.1.1.39/32 [110/4] via 192.168.100.39, 01:13:12, Vlan100
O    192.168.51.0/24 [110/2] via 192.168.220.2, 00:03:46, Vlan220
                     [110/2] via 192.168.210.2, 00:03:46, Vlan210
                     [110/2] via 192.168.50.3, 00:03:46, Vlan500
C    192.168.220.0/24 is directly connected, Vlan220
C    192.168.50.0/24 is directly connected, Vlan500
C    192.168.1.0/24 is directly connected, Vlan1
C    192.168.100.0/24 is directly connected, Vlan100
     192.168.33.0/32 is subnetted, 1 subnets
O       192.168.33.3 [110/2] via 192.168.50.3, 00:03:46, Vlan500
```

2.2 Routing Table – ‘msfc-bottom’

```
 msfc-bottom#sho ip ro

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

     1.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
O       1.1.1.0/24 [110/4] via 192.168.130.39, 01:18:41, Vlan130
                   [110/4] via 192.168.110.10, 01:18:41, Vlan110
O       1.1.1.10/32 [110/4] via 192.168.110.10, 01:18:41, Vlan110
O       1.1.1.39/32 [110/4] via 192.168.130.39, 01:18:41, Vlan130
O    192.168.120.0/24 [110/2] via 192.168.240.1, 01:18:41, Vlan240
                      [110/2] via 192.168.230.1, 01:18:41, Vlan230
C    192.168.210.0/24 is directly connected, Vlan210
C    192.168.240.0/24 is directly connected, Vlan240
C    192.168.150.0/24 is directly connected, Loopback1
C    192.168.110.0/24 is directly connected, Vlan110
C    192.168.230.0/24 is directly connected, Vlan230
O    192.168.10.0/24 [110/2] via 192.168.51.3, 00:09:19, Vlan510
C    192.168.130.0/24 is directly connected, Vlan130
     192.168.140.0/32 is subnetted, 1 subnets
O       192.168.140.1 [110/2] via 192.168.220.1, 00:09:21, Vlan220
                      [110/2] via 192.168.210.1, 00:09:21, Vlan210
C    127.0.0.0/8 is directly connected, EOBC0/6
     9.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
O       9.1.1.10/32 [110/4] via 192.168.110.10, 01:18:43, Vlan110
O       9.1.1.0/24 [110/4] via 192.168.130.39, 01:18:43, Vlan130
                   [110/4] via 192.168.110.10, 01:18:43, Vlan110
O       9.1.1.39/32 [110/4] via 192.168.130.39, 01:18:43, Vlan130
C    192.168.51.0/24 is directly connected, Vlan510
C    192.168.220.0/24 is directly connected, Vlan220
```

```
O    192.168.50.0/24 [110/2] via 192.168.51.3, 00:09:21, Vlan510
                     [110/2] via 192.168.220.1, 00:09:21, Vlan220
                     [110/2] via 192.168.210.1, 00:09:21, Vlan210
C    192.168.1.0/24 is directly connected, Vlan1
O    192.168.100.0/24 [110/2] via 192.168.240.1, 01:18:43, Vlan240
                      [110/2] via 192.168.230.1, 01:18:43, Vlan230
     192.168.33.0/32 is subnetted, 1 subnets
O       192.168.33.3 [110/2] via 192.168.51.3, 00:09:21, Vlan510
```

2.3 Routing Table – ‘7507-access’

```
 7507-ACCESS#  sho ip ro

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

     1.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
O IA    1.1.1.0/24 [110/5] via 192.168.51.2, 00:02:39, FastEthernet0/0/0
                   [110/5] via 192.168.50.1, 00:02:39, FastEthernet0/1/0
O IA    1.1.1.10/32 [110/5] via 192.168.50.1, 00:02:39, FastEthernet0/1/0
                    [110/5] via 192.168.51.2, 00:02:39, FastEthernet0/0/0
O IA    1.1.1.39/32 [110/5] via 192.168.50.1, 00:02:39, FastEthernet0/1/0
                    [110/5] via 192.168.51.2, 00:02:39, FastEthernet0/0/0
O IA 192.168.120.0/24 [110/2] via 192.168.50.1, 00:02:39, FastEthernet0/1/0
O    192.168.210.0/24 [110/2] via 192.168.51.2, 00:02:39, FastEthernet0/0/0
                      [110/2] via 192.168.50.1, 00:02:39, FastEthernet0/1/0
O IA 192.168.240.0/24 [110/2] via 192.168.51.2, 00:02:39, FastEthernet0/0/0
                      [110/2] via 192.168.50.1, 00:02:39, FastEthernet0/1/0
     192.168.150.0/32 is subnetted, 1 subnets
O       192.168.150.1 [110/2] via 192.168.51.2, 00:02:44, FastEthernet0/0/0
O IA 192.168.110.0/24 [110/2] via 192.168.51.2, 00:02:44, FastEthernet0/0/0
O IA 192.168.230.0/24 [110/2] via 192.168.51.2, 00:02:44, FastEthernet0/0/0
                      [110/2] via 192.168.50.1, 00:02:44, FastEthernet0/1/0
C    192.168.10.0/24 is directly connected, FastEthernet6/0/0
O IA 192.168.130.0/24 [110/2] via 192.168.51.2, 00:02:44, FastEthernet0/0/0
     192.168.140.0/32 is subnetted, 1 subnets
O       192.168.140.1 [110/2] via 192.168.50.1, 00:02:44, FastEthernet0/1/0
     9.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
O IA    9.1.1.10/32 [110/5] via 192.168.50.1, 00:02:44, FastEthernet0/1/0
                    [110/5] via 192.168.51.2, 00:02:44, FastEthernet0/0/0
O IA    9.1.1.0/24 [110/5] via 192.168.51.2, 00:02:44, FastEthernet0/0/0
                   [110/5] via 192.168.50.1, 00:02:44, FastEthernet0/1/0
O IA    9.1.1.39/32 [110/5] via 192.168.50.1, 00:02:44, FastEthernet0/1/0
                    [110/5] via 192.168.51.2, 00:02:44, FastEthernet0/0/0
C    192.168.51.0/24 is directly connected, FastEthernet0/0/0
O    192.168.220.0/24 [110/2] via 192.168.51.2, 00:02:44, FastEthernet0/0/0
                      [110/2] via 192.168.50.1, 00:02:44, FastEthernet0/1/0
C    192.168.50.0/24 is directly connected, FastEthernet0/1/0
O IA 192.168.100.0/24 [110/2] via 192.168.50.1, 00:02:44, FastEthernet0/1/0
C    192.168.33.0/24 is directly connected, Loopback1
7507-ACCESS#
```

2.4 Routing Table – 0S/390 Server ’S39’

```
D TCPIP,,N,ROUTE
EZZ2500I NETSTAT CS V2R10 TCPIP 127
DESTINATION      GATEWAY          FLAGS  REFCNT  INTERFACE
DEFAULT          192.168.100.1    UG     000000  GIG0A
DEFAULT          192.168.130.2    UG     000000  GIG0B
1.1.1.0          192.168.130.2    UG     000000  GIG0B
1.1.1.0          192.168.100.1    UG     000000  GIG0A
1.1.1.10         192.168.100.1    UGH    000000  GIG0A
1.1.1.10         192.168.130.2    UGH    000000  GIG0B
1.1.1.39         0.0.0.0          UH     000000  LNKVIPA1
9.1.1.0          192.168.130.2    UG     000000  GIG0B
9.1.1.0          192.168.100.1    UG     000000  GIG0A
9.1.1.10         192.168.130.2    UGH    000000  GIG0B
9.1.1.10         192.168.100.1    UGH    000000  GIG0A
9.1.1.39         0.0.0.0          UH     000000  VIPL09010127
127.0.0.1        0.0.0.0          UH     000004  LOOPBACK
192.168.100.0    0.0.0.0          U      000000  GIG0A
192.168.100.39   0.0.0.0          UH     000001  GIG0A
192.168.110.0    192.168.130.2    UG     000000  GIG0B
192.168.120.0    192.168.100.1    UG     000000  GIG0A
192.168.130.0    0.0.0.0          U      000000  GIG0B
192.168.130.39   0.0.0.0          UH     000000  GIG0B
192.168.230.0    192.168.130.2    UG     000000  GIG0B
192.168.230.0    192.168.100.1    UG     000000  GIG0A
192.168.240.0    192.168.100.1    UG     000000  GIG0A
192.168.240.0    192.168.130.2    UG     000000  GIG0B
23 OF 23 RECORDS DISPLAYED
```

2.5 Routing Table – 0S/390 Server ‘S0A’

```
 D TCPIP,,N,ROUTE
 EZZ2500I NETSTAT CS V2R10 TCPIP 485
 DESTINATION      GATEWAY          FLAGS  REFCNT  INTERFACE
 DEFAULT          192.168.120.1    UG     000000  GIGFE
 DEFAULT          192.168.110.2    UG     000000  GIGFA
 1.1.1.0          192.168.120.1    UG     000000  GIGFE
 1.1.1.0          192.168.110.2    UG     000000  GIGFA
 1.1.1.10         0.0.0.0          UH     000000  LNKVIPA1
 1.1.1.39         192.168.110.2    UGH    000000  GIGFA
 1.1.1.39         192.168.120.1    UGH    000000  GIGFE
 9.1.1.0          192.168.110.2    UG     000000  GIGFA
 9.1.1.0          192.168.120.1    UG     000000  GIGFE
 9.1.1.10         0.0.0.0          UH     000000  VIPL0901010A
 9.1.1.39         192.168.110.2    UGH    000000  GIGFA
 9.1.1.39         192.168.120.1    UGH    000000  GIGFE
 127.0.0.1        0.0.0.0          UH     000004  LOOPBACK
 192.168.100.0    192.168.120.1    UG     000000  GIGFE
 192.168.110.0    0.0.0.0          U      000000  GIGFA
 192.168.110.10   0.0.0.0          UH     000001  GIGFA
 192.168.120.0    0.0.0.0          U      000000  GIGFE
 192.168.120.10   0.0.0.0          UH     000000  GIGFE
 192.168.130.0    192.168.110.2    UG     000000  GIGFA
 192.168.230.0    192.168.110.2    UG     000000  GIGFA
 192.168.230.0    192.168.120.1    UG     000000  GIGFE
 192.168.240.0    192.168.120.1    UG     000000  GIGFE
 192.168.240.0    192.168.110.2    UG     000000  GIGFA
 23 OF 23 RECORDS DISPLAYED
```

Attention now turns to the principles that were applied to derive this steady state network.

3. Best Practices and Design Principles

The following principles and best practices were applied to ensure the delivery of a fast, reliable and redundant network to ABC Corp. Such features are critical to the successful operation of ABC Corp's business.

The interested reader should refer to the Bibliography for a complete list of texts relating to OSPF Design and General Design Guidelines.

Note : OSPF uses the concepts of Areas. A 32-bit number identifies these areas. Often time the Area IDs are abbreviated to Area 0, Area 1 etc. The complete Area ID for Area 1 is actually 0.0.0.1. However for readability purposes this document refers to Area 1.1.1.1 as Area 1. Please be aware of this.

3.1 Totally Stubby Areas and Default Routes

Key design goals here are :

- The IBM S/390 Servers in Area 1 should not be effected by link flaps in other OSPF Areas
- The IBM S/390 Servers in Area 1 should only require knowledge of a default route for all traffic leaving Area 1.
- The IBM S/390 Servers in Area 1 should have multiple exit points from Area 1, thus facilitating network load balancing and redundancy.

These design goals are achieved by invoking the concept of a totally stubby area. Figure 1 provides a graphical representation of the totally stubby area concept.

A totally stubby area is one that blocks external routes and summary routes (inter-area routes) from entering an area. In OSPF parlance this means that LSA Types 3, 4 and 5 are prevented from entering the Totally Stubby Area (with the exception of the default route 0.0.0.0). This leaves the default route of 0.0.0.0 and intra-area routes as the only types being advertised throughout the totally stubby area. In this network, Area 1 encompassing the S/390 Servers is thus defined as a totally stubby area.

Please note that OSPF advertises the state of links and not actual routing table entries (EIGRP is an example of a routing protocol that advertises route table entries).

Any instability or link flapping in the backbone Area 0 will not invoke recalculation / processing on the S/390 Servers in Area 1. In addition, the S/390 Servers now have a default route of 0.0.0.0 available via both ABRs (i.e. 6500 MSFCs). This meets the load balancing and redundancy requirements.

On the MSFC-1, Cisco Express Forwarding (CEF) is used in conjunction with Multi Layer Switching (MLS). When equal costs routes are available, a hashing algorithm is applied to the source / destination IP addresses. The result determines which one of the equal cost paths is used, and is the default behavior for the MSFC-1. A source / destination IP pair will always use the same path, unless a failure occurs.

The terminology of load balancing may thus be more accurately defined as Load Distribution on the MSFC1 / PFC as there is no guarantee of a mathematical balance across equal cost links / routes. The terminology of 'Load Balancing' is however ubiquitous in the industry and is thus adopted in this document. It should be noted that in order to facilitate network load balancing, equal cost per packet round robin distribution was enabled on each S/390 server via the MULTIPATH configuration statement in each S/390 Server's TCP/IP configuration profile.

The following excerpt highlights how to implement the TSA concept on the ABRs (MSFCs). Only the ABRs need to be configured as Totally Stubby. All other routers in Area 1 (in this case the S/390 Servers) are configured as standard Stub Area routers.

A key S/390 OSPF STUB Area configuration consideration should be understood to ensure that the S/390 servers are properly configured as STUB area routers. Specifically, all VIPA interfaces must be explicitly defined as OSPF_INTERFACES in the OMPROUTE configuration file. Should the VIPA interfaces not be defined at all or should they be defined as regular interfaces, OMPROUTE will view them as external routes and they won't get imported into the routing table.

It should also be noted that while *AS_Boundary_Routing* configuration statements were configured in the S/390 OSPF profiles, the *STUB=YES* parameter specified in the *AREA* configuration statement overrides the ASBR statement and prevents the S/390 server from being recognized as an ASB router.

```
msfc-top#

 router ospf 100
  area 1.1.1.1 stub no-summary
  network 192.168.2.0 0.0.0.255 area 0
  network 192.168.50.0 0.0.0.255 area 0
  network 192.168.60.0 0.0.0.255 area 0
  network 192.168.100.0 0.0.0.255 area 1.1.1.1
  network 192.168.120.0 0.0.0.255 area 1.1.1.1
  network 192.168.140.0 0.0.0.255 area 0
  network 192.168.210.0 0.0.0.255 area 0
  network 192.168.220.0 0.0.0.255 area 0
  network 192.168.230.0 0.0.0.255 area 1.1.1.1
  network 192.168.240.0 0.0.0.255 area 1.1.1.1
  distribute-list 1 in Vlan230
 distribute-list 1 in Vlan240
```

The following is an excerpt from Server S39's OSPF configuration profile (omproute.config):

```
; OMPROUTE Configuration file for OSPF for Server S39
;
; AREA
;    Sets the OSPF AREA.  If no areas are defined, the router software
;    assumes that all the router's directly attached networks belong to
;    the backbone area (area ID 0.0.0.0)
;
 AREA
   Area_Number=1.1.1.1
   Authentication_Type=None
   Stub_area=yes
;  Stub_Default_Cost=53
;  Import_Summaries=No;
;
```

The following commands can be used to verify the use of the Totally Stubby concept :

```
msfc-top#sh ip ospf 100

 Routing Process "ospf 100" with ID 192.168.140.1
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 It is an area border router
 SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
 Minimum LSA interval 5 secs. Minimum LSA arrival 1 secs
 Number of external LSA 0. Checksum Sum 0x0
 Number of opaque AS LSA 0. Checksum Sum 0x0
 Number of DCbitless external and opaque AS LSA 0
 Number of DoNotAge external and opaque AS LSA 0
 Number of areas in this router is 2. 1 normal 1 stub 0 nssa
 External flood list length 0
```

```
    Area BACKBONE(0)
        Number of interfaces in this area is 4
        Area has no authentication
        SPF algorithm executed 20 times
        Area ranges are
        Number of LSA 31. Checksum Sum 0xEDB65
        Number of opaque link LSA 0. Checksum Sum 0x0
        Number of DCbitless LSA 0
        Number of indication LSA 0
        Number of DoNotAge LSA 0
        Flood list length 0
    Area 1.1.1.1
        Number of interfaces in this area is 4
        It is a stub area, no summary LSA in this area
          generates stub default route with cost 1
        Area has no authentication
        SPF algorithm executed 27 times
        Area ranges are
        Number of LSA 12. Checksum Sum 0x6D266
        Number of opaque link LSA 0. Checksum Sum 0x0
        Number of DCbitless LSA 0
        Number of indication LSA 0
        Number of DoNotAge LSA 0
        Flood list length 0
```

Likewise the following command illustrates that Server S39 has been configured as a Stub Area :

```
d tcpip,,omproute,ospf,list,areas
EZZ7832I AREA CONFIGURATION 139
AREA ID          AUTYPE          STUB? DEFAULT-COST IMPORT-SUMMARIES?
1.1.1.1          0=NONE           YES         1             YES
0.0.0.0          0=NONE           NO          N/A           N/A
```

3.2 Use of Loopback Interfaces

Routers within an OSPF network must have a unique Router ID. Routers within an OSPF Autonomous System identify each other via this Router ID. The default method on Cisco routers uses the Highest IP address in the router as the Router ID. However, when a loopback interface is used on a Cisco Router, the Router ID for the OSPF process is automatically assigned based on the IP address of the loopback interface. Controlling the loopback address simplifies troubleshooting and is thus recommended.

Note : Some texts suggest using the loopback interface as the interface is always 'up'. In reality, Cisco routers can use the loopback address of an interface as the Router ID even if the actual physical interface is 'down'.

Selecting a Router ID on S/390 OSPF routers should also be a conscious and deliberate configuration activity. Specifically, the S/390 OSPF Router allows you the flexibility to explicitly define the router id to be any one of the interfaces defined within the OSPF configuration file. If a Router ID is not explicitly configured, the OSPF router (OMPROUTE) will arbitrarily choose one from the list of defined OSPF interfaces. Therefore, it is possible that OMPROUTE will select the IP address assigned to a dynamic VIPA interface as the Router ID, which could potentially move. To avoid this situation, it is highly recommended that the address assigned to a Static VIPA interface be chosen as the respective S/390 OSPF Router ID.

Excerpts from both MSFCs show how to utilize this loopback interface :

```
 msfc-top#

 interface Loopback0
 ip address 192.168.140.1 255.255.255.255
```

```
msfc-bottom#

 interface Loopback0
  ip address 192.168.150.1 255.255.255.0
```

The use of Router IDs is illustrated with the '**show ip ospf neighbor'** command. Issuing this on 'msfc-top' shows that a neighbor relationship has been formed with 'msfc-bottom' over VLANs 210/220/230/240. The Neighbor ID (which is the Router ID) used is 192.168.150.1. This is the loopback address of 'msfc-bottom'.

```
 msfc-top#sh ip ospf neig

 Neighbor ID     Pri   State           Dead Time   Address         Interface
 9.1.1.39          0   FULL/DROTHER    00:00:08    192.168.100.39  Vlan100
 9.1.1.10          0   FULL/DROTHER    00:00:08    192.168.120.10  Vlan120
 192.168.150.1     1   FULL/DR         00:00:07    192.168.210.2   Vlan210
 192.168.150.1     1   FULL/DR         00:00:07    192.168.220.2   Vlan220
 192.168.150.1     1   FULL/DR         00:00:07    192.168.230.2   Vlan230
 192.168.150.1     1   FULL/DR         00:00:07    192.168.240.2   Vlan240
 192.168.98.1      1   FULL/DR         00:00:06    192.168.50.3    Vlan500
```

From Server S39's viewpoint, the neighbor relationships are indeed formed using the loopback addresses of the 6500s as the Router IDs :

```
d tcpip,,omproute,ospf,nbrs
 EZZ7851I NEIGHBOR SUMMARY 121
 NEIGHBOR ADDR   NEIGHBOR ID     STATE  LSRXL DBSUM LSREQ HSUP IFC
 192.168.130.2   192.168.150.1      128     0     0     0  OFF GIG0B
 192.168.100.1   192.168.140.1      128     0     0     0  OFF GIG0A


Note: The STATE can be one of the following:
     1 (down), 2 (backup), 4 (looped back),(point-to-point),
    32 (DR other), 64 (backup DR), or 128 (designated router)
```

3.3 Controlling DR / BDR responsibilities

In OSPF, broadcast multi-access networks such as Gigabit Ethernet require the use of Designated Router (DR) / Backup Designated Router (BDR) functions. These help reduce the number of adjacencies that need to be formed and thus limit routing protocol traffic, routing overhead and the size of the Link State Database. (In reality there are only 2 router instances on any Vlan in this network, so the benefits of DR/BDR functions are somewhat limited.)

Implicit in this design is the desire to limit processing on the S/390 Servers. With this in mind all VLANs that face the S/390 Servers (VLANs 100/110/120/130) will elect the relevant MSFC as the DR and prevent the S/390 Servers from assuming either DR or BDR roles. This helps minimize processing on the S/390 Servers.

This is achieved by manually setting the OSPF Priority command on the MSFCs and the S/390 Servers. On each Vlan, the relevant MSFC's OSPF priority is set to 10, while the relevant S/390 Server's priority is set to 0. Therefore the MSFCs become the DR in each case and the S/390 Servers are prevented from assuming either DR or BDR roles.

It should be noted that each Vlan facing the S/390 Servers has only 2 logical members i.e. the relevant MSFC virtual interface and the OSA-Express interface on the S/390 Server. The reason for this is explained in detail in the Section 3.6 - Minimizing Convergence Times. The VLAN may look like a ‘point-to-point’ connection but in essence it is still a broadcast multi-access network, where multiple routers may reside.

In large scale OSPF networks, it is recommended that a router should only assume DR/DBR responsibilities on one of its interfaces. Given the limited size of this network, this was not considered a pertinent issue here.

The following excerpts from the MSFCs and the S/390 Servers show how to control the election of the DR/BDRs :

```
msfc-top#

interface Vlan100
  ip address 192.168.100.1 255.255.255.0
  ip ospf hello-interval 3
  ip ospf dead-interval 9
  ip ospf priority 10
```

S39 Server - OSPF configuration profile (**omproute.config**) :

```
OSPF_INTERFACE
   IP_address=192.168.100.39
   NAME=GIG0A
   Subnet_mask=255.255.255.0
   Demand_Circuit=no
   Attaches_To_Area=1.1.1.1
   MTU=1500
   Retransmission_Interval=5
   Transmission_Delay=1
   Router_Priority=0
   Hello_Interval=3
   Dead_Router_Interval=9
   Cost0=3
```

As an example; on VLAN 100 the MSFC will become the DR and Server S39 will be prevented from assuming either DR or BDR responsibilities.

Using the ‘sh ip ospf neighbor’ command on the MSFCs illustrates the successful invocation of this design goal.

```
msfc-top#sh ip ospf neig

Neighbor ID     Pri   State           Dead Time   Address         Interface
9.1.1.39          0   FULL/DROTHER    00:00:08    192.168.100.39  Vlan100
9.1.1.10          0   FULL/DROTHER    00:00:08    192.168.120.10  Vlan120
192.168.150.1     1   FULL/DR         00:00:07    192.168.210.2   Vlan210
192.168.150.1     1   FULL/DR         00:00:07    192.168.220.2   Vlan220
192.168.150.1     1   FULL/DR         00:00:07    192.168.230.2   Vlan230
192.168.150.1     1   FULL/DR         00:00:07    192.168.240.2   Vlan240
192.168.98.1      1   FULL/DR         00:00:06    192.168.50.3    Vlan500
```

The Neighbor IDs of 9.1.1.39 and 9.1.1.0 refer to Server S39 and Server S0A respectively. The ‘Pri’ value of 0 means that neighbors 9.1.1.39 and 9.1.1.0 have their OSPF Priority set to 0. The state of FULL means a full adjacency has been formed between the MSFC and the relevant S/390 Server over Vlans 100 / 120. The DROTHER indicates that the relevant S/390 is neither DR nor BDR.

This information is further consolidated with the command ‘**`sh ip ospf int vlan 100’`**. One can see that ‘msfc-top’ assumes DR responsibility on Vlan 100. One can also see that no BDR has been elected on the multi-access network.

```
 msfc-top#sh ip ospf int vlan 100
 Vlan100 is up, line protocol is up
   Internet Address 192.168.100.1/24, Area 1.1.1.1
   Process ID 100, Router ID 192.168.140.1, Network Type BROADCAST, Cost: 1
   Transmit Delay is 1 sec, State DR, Priority 10
  Designated Router (ID) 192.168.140.1, Interface address 192.168.100.1
  No backup designated router on this network
  Timer intervals configured, Hello 3, Dead 9, Wait 9, Retransmit 5
    Hello due in 00:00:01
  Index 1/2, flood queue length 0
  Next 0x0(0)/0x0(0)
  Last flood scan length is 3, maximum is 4
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 9.1.1.39
  Suppress hello for 0 neighbor(s)
msfc-top#
```

This information can be further reconciled with the view from Server S39's viewpoint :

```
d tcpip,,omproute,ospf,if,name=gig0a

EZZ7850I INTERFACE DETAILS 173
                INTERFACE ADDRESS:      192.168.100.39
                ATTACHED AREA:          1.1.1.1
                PHYSICAL INTERFACE:     GIG0A
                INTERFACE MASK:         255.255.255.0
                INTERFACE TYPE:         BRDCST
                STATE:                  32
                DESIGNATED ROUTER:      192.168.100.1
                BACKUP DR:              0.0.0.0

  DR PRIORITY:       0  HELLO INTERVAL:    3  RXMT INTERVAL:     5
  DEAD INTERVAL:     9  TX DELAY:          1  POLL INTERVAL:     0
  DEMAND CIRCUIT:  OFF  HELLO SUPPRESS:  OFF  SUPPRESS REQ:    OFF
  MAX PKT SIZE:   1500  TOS 0 COST:        3  DB_EX INTERVAL:    9

  # NEIGHBORS:       1  # ADJACENCIES:     1  # FULL ADJS.:      1
  # MCAST FLOODS:   30  # MCAST ACKS:     19  DL UNICAST:      OFF
  MC FORWARDING:   OFF

  NETWORK CAPABILITIES:
   BROADCAST
   DEMAND-CIRCUITS
```

In summary, the use of OSPF Priority to effect DR / BDR election on broadcast multi-access networks helps minimize processing on the S/390 Servers.

3.4 Minimizing Sub-Optimal Routing Through OSPF Route Costs

The design goal here is to minimize the potential for sub-optimal routing. Sub-optimal routing in this instance is defined as traffic destined for one S/390 Server being routed through the other S/390 Server.

Minimizing the directing of unnecessary traffic through an S/390 Server is not only good practice, but is also a more efficient use of the S/390 Servers. In essence, the S/390 Servers should only handle traffic destined for the IP addresses allocated to it. They should not be acting as transit routers. This role is best left to the 6500 / MSFCs.

It should be noted that sub-optimal routing will only occur (and is acceptable) in the final design if multiple failures occur. See Section 4.3 where multiple failures are invoked to illustrate sub-optimal routing.

There are two ways in which this design goal of minimizing sub-optimal routing is achieved :

- Controlling the costs of routes being advertised by the MSFCs and S/390 Servers.
- Using Area 0 and Area 1 links between 6500s / ABRs

The default cost of all routes being advertised by the MSFCs is 1 (Note : the MSFC is 'virtual' and has no knowledge of physical line speed, hence the cost for all interfaces defaults to 1. The default cost advertised by the OS/390s is 3. Therefore the MSFC advertises lower costs than the S/390 for the same segment. This helps to minimize the potential for sub-optimal routing, as routes with lower costs are preferred.

The second way to minimize sub-optimal routing is considered in detail next.

3.5 Minimizing Sub-Optimal Routing via Area 0 / Area 1 Links

The design goal is still to minimize sub-optimal routing. In OSPF, if multiple entries exist to a destination network or host, the longest match rule applies. If there are still multiple entries available (i.e. entries have the same longest match) then IntraArea routes are preferred over InterArea routes. In the event that there are still multiple entries available, the OSPF cost metric for the destination becomes the mitigating factor in choosing the path. At this stage, if multiple equal cost routes are still available then load balancing can occur (Section 3.1 describes how the MSFC-1 switches packets when equal cost routes occur. By default OSPF can load balance across 4 equal cost paths).

The lookup procedure described above is consistent with RFC 2328 and is illustrated further in the sections below. It is important to understand this concept before one proceeds.

In this design there are 4 x **Core** Gigabit Ethernet links between the 6500s (See Figure 1). These encompass Vlans 210, 220, 230 and 240 respectively. Figure 1 illustrates that Vlans 210 and 220 are assigned to Area 0, while Vlans 230 and 240 are assigned to Area 1

(Note: - An equally applicable design here is to use channeling features such as Gigabit EtherChannel. In this instance one could make a Gigabit EtherChannel of Vlans 210 / 220 and another Gigabit EtherChannel of Vlans 230/240. This would reduce the core links to 2 subnets as opposed to the 4 subnets used in this implementation. Channeling would help reduce the number of LSAs, adjacencies etc. in this OSPF network. Channeling was not however used at the request of the customer and a separate subnet per Core Vlan was deployed.)

One may ponder why such an arrangement was used. Let's consider what happens if all of the core Vlans are assigned exclusively to Area 0 or exclusively to Area 1.

3.5.1 Allocating All Core Links to Area 0

In this instance sub-optimal routing may occur in the event of failures. In this example all of the core links between 6500s are assigned to Area 0. In a steady state traffic from 'msfc-top' to VIPA 1.1.1.10 will traverse Vlan 120. Now consider what happens when Vlan 120 'fails'.

After convergence the routing table excerpt from 'msfc-top' looks like this :

```
msfc-top#sh ip ro
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route
```

```
Gateway of last resort is not set

     1.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
O       1.1.1.0/24 [110/4] via 192.168.100.39, 00:00:00, Vlan100
O       1.1.1.10/32 [110/8] via 192.168.100.39, 00:00:00, Vlan100
```

Traffic destined for VIPA 1.1.1.10 on Server S0A uses Vlan 100 (Server S39) as it's next hop. It may appear that the quickest route is via a next hop of 'msfc-bottom'. However 'msfc-bottom' is advertising VIPA 1.1.1.10 as an **Inter**Area route while Server S39 is advertising VIPA 1.1.1.10 as an **Intra**Area route. The IntraArea route is preferred over the InterArea route (as explained earlier), hence Vlan 100 is the correct next hop.

Traffic from 'msfc-top' destined for VIPA 1.1.1.10 would thus follow the path of Vlan 100 / Vlan 130 / Vlan 110 to reach the VIPA address of 1.1.1.10.

This is regarded as sub-optimal routing in this design. Allocating ALL of the core links to Area 0 is thus not the answer. Attention now turns to the effects of assigning all core links to Area 1 instead.

3.5.2 Allocating All Core Links to Area 1

In this instance Vlans 210/220/230/240 are now assigned exclusively to Area 1. The result of this is the potential for black holing of traffic.

Area 1 is defined as totally stubby. Therefore both MSFCs advertise 0.0.0.0 as the default over all Area 1 links (including Vlans 210/220/230/240). Once the SPF algorithm is computed, each MSFC sees the other MSFC as the default route of 0.0.0.0 over each of the core Vlans.

Routing Tables using the longest match rule (in IBM parlance this is termed 'most descriptive to least descriptive match'). If the longest match for traffic is 0.0.0.0, then the traffic will be black-holed. The packet will 'die' once the Time To Live (TTL) for the packet expires.

The above scenario assumes that no other router is originating a default route except the ABRs.

To solve the 'black-holing' problem, Distribute Lists and Access Lists can be applied to the Area 1 links, namely Vlans 210/220/230/240. These distribute lists prevent each MSFC from injecting the 0.0.0.0 entry into the routing table. Therefore the black-holing problem is resolved. As you will see in the next section however, an alternate solution was found that does **NOT** require distribute lists at all.

For informative purposes only, the following excerpts illustrate the application of Distribute Lists to prevent black-holing when ALL core links are assigned to the Totally Stubby Area, Area 1.

```
 msfc-top#wr t

 router ospf 100
  area 1.1.1.1 stub no-summary
  network 192.168.2.0 0.0.0.255 area 0
  network 192.168.50.0 0.0.0.255 area 0
  network 192.168.60.0 0.0.0.255 area 0
  network 192.168.100.0 0.0.0.255 area 1.1.1.1
  network 192.168.120.0 0.0.0.255 area 1.1.1.1
  network 192.168.140.0 0.0.0.255 area 0
  network 192.168.210.0 0.0.0.255 area 0
  network 192.168.220.0 0.0.0.255 area 0
  network 192.168.230.0 0.0.0.255 area 1.1.1.1
  network 192.168.240.0 0.0.0.255 area 1.1.1.1
  distribute-list 1 in Vlan210
  distribute-list 1 in Vlan220
  distribute-list 1 in Vlan230
  distribute-list 1 in Vlan240
 !
 access-list 1 deny   0.0.0.0
 access-list 1 permit any
```

In general, allocating ALL of the core links to Area 1 is not the ideal solution.

3.5.3 The Solution

The preferred solution is to adopt both Area 0 and Area 1 links between the 6500s. This design assigns Vlans 210 and 220 to Area 0 and Vlans 230 and 240 to Area 1.

The interesting side effect of this is to prevent the 0.0.0.0 entry being injected into the routing table of both ABRs. Therefore the distribute lists illustrated in the previous section are **NOT** required and black holing of traffic is not an issue. This simplifies the configuration.

The sub-optimal routing problem described in Section 3.5.1 is also solved as ‘msfc-bottom’ now advertises 1.1.1.10 as an IntraArea route to ‘msfc-top’. The routing table excerpt for this scenario (i.e. with Vlan 120 ‘down’) substantiates this :

```
 msfc-top#sh ip ro
 Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
        D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
        N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
        E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
        i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
        * - candidate default, U - per-user static route, o - ODR
        P - periodic downloaded static route

Gateway of last resort is not set

     1.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
O       1.1.1.0/24 [110/4] via 192.168.100.39, 00:02:02, Vlan100
O       1.1.1.10/32 [110/5] via 192.168.240.2, 00:02:02, Vlan240
                    [110/5] via 192.168.230.2, 00:02:02, Vlan230
```

The VIPA address of 1.1.1.10 is now available via IntraArea routes over Vlans 230 and 240. The traffic will no longer be forwarded to Server S39.

It is worth noting that both Server S39 and ‘msfc-bottom’ advertise the route to VIPA 1.1.1.10. Both advertise it with the same /32 mask and as an IntraArea route. Therefore the mitigating factor now comes down to the route cost as described earlier; ‘msfc-top’ sees VIPA 1.1.1.10 with a route cost of 5 via ‘msfc-bottom’ and a route cost of 8 via Server S39. Therefore ‘msfc-top’ inserts the lower cost routes into the routing table. The CEF switching algorithm will determine which of these two equal cost routes to use.

This solved the particular sub-optimal routing problem described in Section 3.5.1. Under a certain combination of failure conditions, sub-optimal routing can still occur and is acceptable. An example of this is documented in Section 4.3.

It should be noted that traffic to/from VIPA addresses will only traverse core links if failures occur. In a steady state one should only expect OSPF updates, CDP updates etc. to traverse the core links.

In summary, the use of OSPF route costs in conjunction with this distribution of core links between areas thus contributes to the design goal of minimizing sub-optimal routing. Understanding this aspect of this design is critical to the successful implementation of such an environment.

3.6 Minimizing Convergence Times

Even the best designed network is subject to failures and outages. Mission critical networks like ABC Corp's demand highly available, redundant, fault tolerant designs. With this in mind it is imperative to ensure the fastest possible convergence of the network in the event of failures.

The convergence time in an OSPF network is a function of the following :

1. Time it takes to detect topology changes / link failure
2. Time it takes routers to exchange the 'new' information and build a new routing table.

There is an additional component called the SPF Delay timer. This timer causes a delay between when OSPF receives a topology change and when it starts the SPF calculation. This is recommended where flapping links could cause repeated execution of the SPF algorithm.

The design features documented here concentrate on detecting topology changes as quickly as possible. Please bear in mind that there are trade-offs associated with this 'optimization'. In general the optimizations increase the frequency of information exchange between routers. This is an added overhead, but it is deemed acceptable in the pursuit of minimized convergence times.

Several features were deployed to minimize convergence times :

1. MSFC Autostate
2. Reduced OSPF Hello and Dead Timers
3. PortFast

Is should be stated from the outset that this network was designed with minimal need for Spanning Tree features. This is a Layer 3 design with its attendant benefits for convergence. That is not to say that Spanning Tree is not used. Spanning Tree is indeed enabled to prevent loops but Layer 3 OSPF convergence is preferred to L2 Spanning Tree convergence.

Discussion now turns to the features listed above :

3.6.1 MSFC Autostate

As was stated earlier, each Vlan facing a S/390 Server has only 2 logical members i.e. the relevant MSFC virtual interface and the OSA-Express physical interface. The key requirement is to detect link failure as soon as possible, recalculate new routes and flood LSAs. In the event of link failure on one of these Vlans, the MSFC will detect this immediately. The Vlan on the MSFC will immediately transition to the 'down' state, as there are no other members/ports in the Vlan. This will cause a recalculation of the SPF algorithm followed by LSA flooding by the MSFC.

The excerpt from the 6500-top shows that there are only 2 entities within Vlan 100, namely the MSFC Virtual Interface on logical port 15/1 and the OSA-Express Adapter connected to physical port 1/1.

```
6500-top> (enable) sh vlan 100

VLAN Name                              Status    IfIndex Mod/Ports, Vlans
---- -------------------------------- --------- ------- -----------------------
100  VLAN0100                         active    120      1/1
                                                        15/1
```

The effect of MSFC Autostate is to reduce the convergence time in the event of failure. The only downside is the need for a Vlan / subnet per OSA-Express connection. Given the availability of IP addressing in this network, this was of little concern.

3.6.2 OSPF Hello and Dead Timers

Another method for detecting topology changes / failures is through the use of OSPF Hello and Dead Timers. By default OSPF routers exchange 'Hello' packets every 10 seconds on broadcast multi-access networks. If a neighbor does not receive a Hello packet within an interval called the Dead Timer (by default this is set to 4 times the Hello timer) it presumes that the neighbor is down and will thus flood LSAs. In this design all routers within the network use the following amended timers :

Hello Timer – 3 seconds
Dead Timer – 9 seconds

All routers within an area must adhere to the same Hello and Dead timers. It should be noted that these are aggressive Hello/Dead timer values and are not to be viewed as recommended practice i.e. these values may not be applicable to all networks.

In the event of link failure one would expect the MSFC Autostate feature to 'kick in' long before the Hello / Dead Timer expires. However, in the event that the MSFC Autostate feature does not apply then the Hello / Dead Timers may be required to detect a problem with a neighboring router. The MSFC Autostate and OSPF Hello / Dead Timer mechanisms are essentially complementary to each other. (In traditional routers such as the Cisco 7500, the Ethernet keep-alive timers may supersede the OSPF Hello / Dead Timers.)

The essential feature is detecting the changes in the network topology as soon as possible and flooding LSA packets to neighbors. Only then can the neighbors' Link State Databases (LSDB) be updated correctly, thus allowing for the execution of the SPF algorithm and creation of new routing tables.

It is worth noting that the router that detects the change will recalculate the SPF algorithm first and then flood the LSAs. The receiving router will flood first, then update the LSDB, and finally run the SPF algorithm.

3.6.3 PortFast

PortFast is a feature designed for use with end-station devices. PortFast is also applied in this instance to the OSA-Express interfaces. The idea is to put a port into 'Forwarding' state as soon as link is detected. Without PortFast enabled one has to transition through the Listening and Learning stages (30 seconds in total) of Spanning Tree before the port can actually pass valid traffic.

All OSA-Express interfaces have PortFast enabled on the respective Catalyst 6500 switch port to which they connect. An even more useful command is the 'set port host' command that automatically enables portfast, plus turns off trunking and channeling negotiation. This command is thus recommended where one would use the PortFast command.

```
6500-top> (enable) set port host 1/1
2000 Jan 15 04:01:32 %SYS-6-CFG_CHG:Module 1 block changed by Console//
Port(s) 1/1 channel mode set to off.

Warning: Spantree port fast start should only be enabled on ports connected to
a single host.  Connecting hubs, concentrators, switches, bridges, etc. to a
fast start port can cause temporary spanning tree loops.  Use with caution.

Spantree port  1/1 fast start enabled.
Port(s)  1/1 trunk mode set to off.
```

It should be noted that the core links (Vlans 210/220/230/24) do not have PortFast enabled. This is because these are switch-to-switch connections and PortFast is only recommended for end station devices. It is also worth noting that OSA-Express interfaces expect Spanning Tree to be turned OFF on the switch ports they connect to. In reality, PortFast does run Spanning Tree in the background. This does not effect the OSA-Express requirement as the port is transitioned immediately into forwarding state, thus preventing the OSA-Express timers from expiring.

The following show commands highlight the use of PortFast on port 1/1, the OSA-Express port, while PortFast is disabled on port 6/1, which is one of the Core Vlan links.

```
6500-top> (enable) show spantree 1/1
Port                     Vlan Port-State    Cost  Priority Portfast
Channel_id
------------------------ ---- ------------- ----- -------- ---------- ---------
 1/1                     100  forwarding        4       32 enabled    0

6500-top> (enable) show spantree 6/1
Port                     Vlan Port-State    Cost  Priority Portfast
Channel_id
------------------------ ---- ------------- ----- -------- ---------- ---------
 6/1                     210  forwarding        4       32 disabled   0
```

From a convergence point of view, the use of PortFast allows the MSFCs and the S/390 Servers to immediately form neighbor adjacencies over these PortFast enabled links. Without PortFast these ports would have to go through the listening and learning states (by default about 30 seconds) before commencement of the OSPF adjacency process.

In summary, this section has described optimizations that will allow for the rapid convergence / convergence of the ABC Corp. network. The next section introduces a variety of failure scenarios to help consolidate one’s understanding of the network design.

3.7 Failure Scenarios

The failures below are not designed to be indicative of typical failures that may occur. Instead they were specifically invoked to help ABC Corp. understand how the network recovers given various failure domains

Please refer to Figure 1 for a review of the steady state network.

Throughout these tests ‘trace’ or ‘traceroute’ commands are utilized. When one uses a ‘trace’ command with a destination of a VIPA address on one of the S/390 Servers, the S/390 Server returns the VIPA address as the last hop to the issuing router. One would expect the IP address of the OSA-Express interface to be reported. However this is not the case so please be aware of this in the following scenarios.

3.8 Equal Cost Routing

If equal cost routes are present in a routing table, failure of the active route will cause immediate fail-over to the remaining equal cost route(s). The key point is that OSPF convergence is not required before invoking the new route. Please refer to Figure 2 for a graphical representation of this failure scenario.

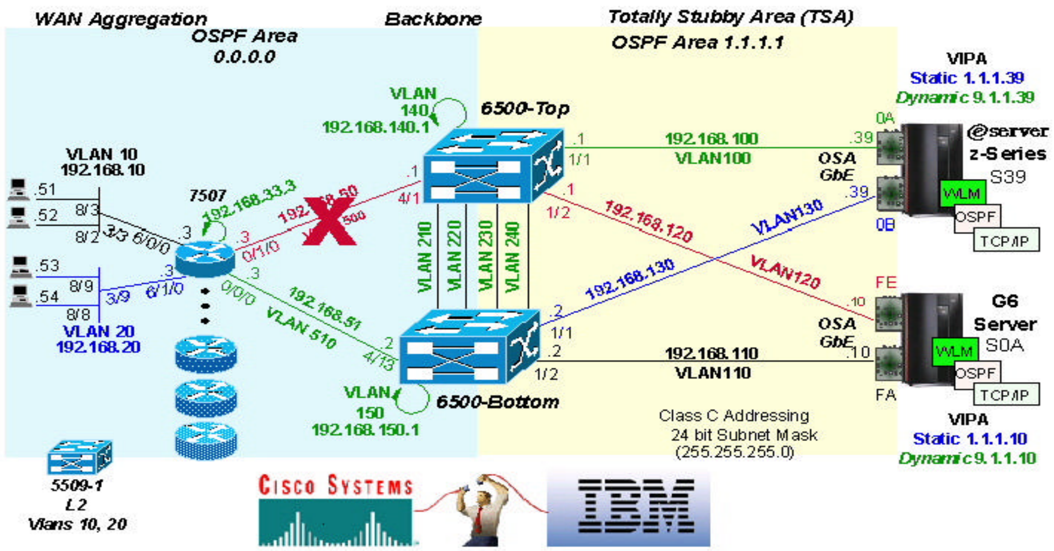

Figure 2 - Equal Cost Routing

The 7507-Access router is used to demonstrate this principle. In the steady state the routing table excerpt looks like this :

```
7507-ACCESS#sho ip ro
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

     1.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
O IA    1.1.1.0/24 [110/5] via 192.168.50.1, 00:05:33, FastEthernet0/1/0
                   [110/5] via 192.168.51.2, 00:05:33, FastEthernet0/0/0
O IA    1.1.1.10/32 [110/5] via 192.168.50.1, 00:02:46, FastEthernet0/1/0
                    [110/5] via 192.168.51.2, 00:02:46, FastEthernet0/0/0
O IA    1.1.1.39/32 [110/5] via 192.168.51.2, 00:09:39, FastEthernet0/0/0
                    [110/5] via 192.168.50.1, 00:09:39, FastEthernet0/1/0
```

The equal cost routes are highlighted in the ‘trace’ command. The ‘trace’ command shows that 192.168.50.1 is used as the next hop :

```
7507-ACCESS#trace 1.1.1.10

Tracing the route to 1.1.1.10
  1 192.168.50.1 0 msec
  2 1.1.1.10 0 msec 0 msec 0 msec /*Host reports VIPA address not OSA address*/
```

To invoke the failure condition, the active path is then shutdown on the 7500 :

```
7507-ACCESS#conf t
7507-ACCESS(config)#int fas 0/1/0 /* VLAN 500 */
7507-ACCESS(config-if)#shut
7507-ACCESS(config-if)#^Z
```

The routing table excerpt immediately reflects the presence of the alternate route :

```
7507-ACCESS#sho ip route

Gateway of last resort is not set

     1.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
O IA    1.1.1.0/24 [110/5] via 192.168.51.2, 00:06:08, FastEthernet0/0/0
O IA    1.1.1.10/32 [110/5] via 192.168.51.2, 00:03:21, FastEthernet0/0/0
O IA    1.1.1.39/32 [110/5] via 192.168.51.2, 00:10:14, FastEthernet0/0/0
```

Using the ‘trace’ command illustrates that the remaining route is now used. This route was immediately used without waiting for OSPF convergence.

```
7507-ACCESS#trace 1.1.1.10

Tracing the route to 1.1.1.10
  1 192.168.51.2 0 msec 0 msec 4 msec
  2 1.1.1.10 0 msec 0 msec 0 msec
```

Immediate fail-over to the remaining equal cost route has thus been demonstrated.

3.9 Failed routes requiring OSPF convergence

In this example, ‘msfc -top’ uses the direct route to VIPA 1.1.1.10 via VLAN 120. Failure of this link causes the ‘msfc-top’ to recalculate a new route to the VIPA address. LSA flooding also occurs to notify other routers of this link state change. Figure 3 illustrates the failure scenario.

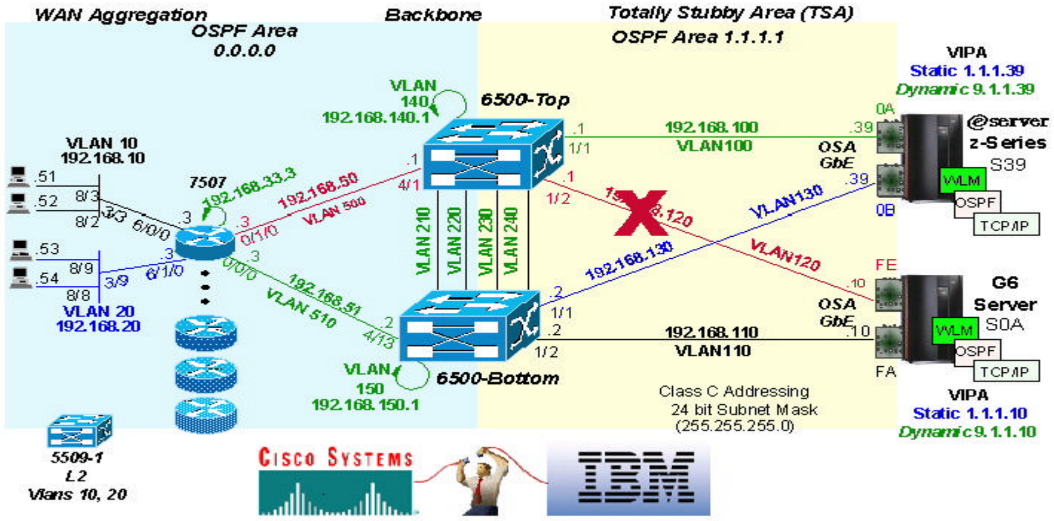

Figure 3 - OSPF Convergence

The MSFC Autostate design guideline described earlier allows ‘msfc-top’ to immediately detect the link failure and invoke the process. Once the SPF algorithm has been executed, equal cost IntraArea routes (via VLAN 230 and 240) are now available to reach the VIPA address. The CEF switching algorithm determines which of these new links is actually used.

The steps below illustrate the process. The first step is to examine the steady state routing table :

```
msfc-top#sh ip ro

Gateway of last resort is not set

     1.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
O       1.1.1.0/24 [110/4] via 192.168.100.39, 00:00:21, Vlan100
                   [110/4] via 192.168.120.10, 00:00:21, Vlan120
O       1.1.1.10/32 [110/4] via 192.168.120.10, 00:00:21, Vlan120
O       1.1.1.39/32 [110/4] via 192.168.100.39, 00:00:21, Vlan100
```

The ‘trace’ command shows that the direct route via VLAN 120 is used. No intervening hops are listed, therefore the path is direct.

```
msfc-top#trace 1.1.1.10

Type escape sequence to abort.
Tracing the route to 1.1.1.10
```

```
  1 1.1.1.10 0 msec 0 msec 4 msec
```

Failing the only port in VLAN 120 on ‘6500-top’ causes route recalculation via the SPF algorithm and LSA flooding of the state change by ‘msfc-top’. The new route(s) are injected into the routing table :

```
msfc-top#sh ip ro

Gateway of last resort is not set

     1.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
O       1.1.1.0/24 [110/4] via 192.168.100.39, 00:00:01, Vlan100
O       1.1.1.10/32 [110/5] via 192.168.240.2, 00:00:01, Vlan240
                    [110/5] via 192.168.230.2, 00:00:01, Vlan230
O       1.1.1.39/32 [110/4] via 192.168.100.39, 00:00:01, Vlan100
```

Finally the ‘trace’ command highlights how the new path invoked is over VLAN 240.

```
msfc-top#trace 1.1.1.10

Type escape sequence to abort.
Tracing the route to 1.1.1.10

  1 192.168.240.2 4 msec
  2 1.1.1.10 0 msec 0 msec 4 msec
```

This failure thus demonstrates a failure scenario where OSPF convergence is required to derive a new route to the destination.

3.10 IntraArea versus InterArea routes

This failure scenario demonstrates how sub-optimal routing can occur when certain combinations of failure conditions occur. The issue is similar to the one described earlier whereby IntraArea routes are preferred over InterArea routes.

The failure scenario is illustrated in Figure 4. This failure builds upon the previous example; 'msfc-top' is using the IntraArea routes to reach destination VIPA of 1.1.1.10 (as Vlan 120 is down). Failure of both available IntraArea routes i.e. Vlans 230 and 240 will now lead to a sub-optimal routing situation via Server S39.

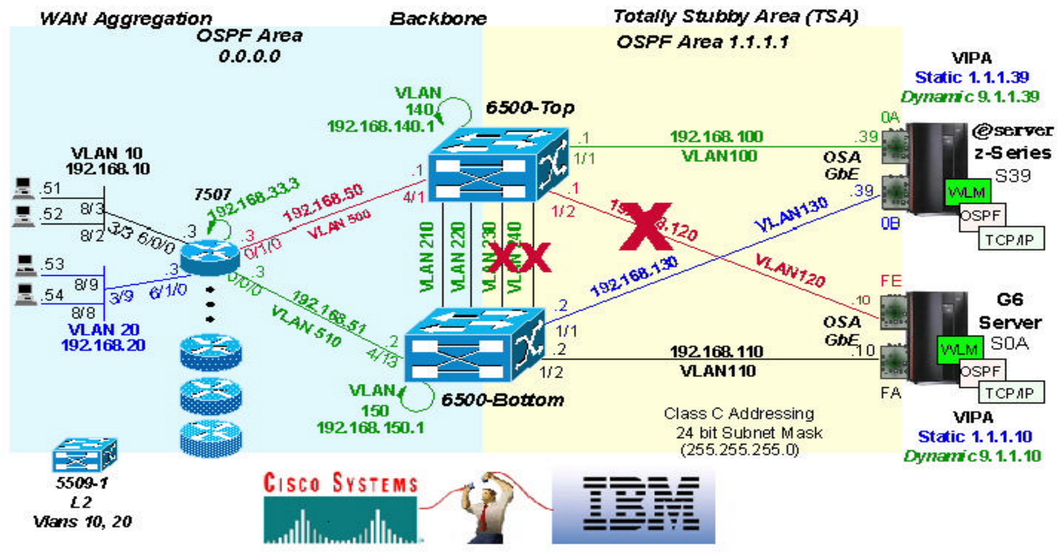

Figure 4 - IntraArea versus InterArea Routes

Here's the routing table excerpt with Vlan 120 down.

```
msfc-top#sh ip ro
```

Gateway of last resort is not set

```
     1.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
O       1.1.1.0/24 [110/4] via 192.168.100.39, 00:02:02, Vlan100
O       1.1.1.10/32 [110/5] via 192.168.240.2, 00:02:02, Vlan240
                    [110/5] via 192.168.230.2, 00:02:02, Vlan230
O       1.1.1.39/32 [110/4] via 192.168.100.39, 00:02:02, Vlan100
```

The 'trace' command shows how Vlan 240 is used.

```
msfc-top#trace 1.1.1.10

Type escape sequence to abort.
```

```
Tracing the route to 1.1.1.10

  1 192.168.240.2 4 msec
  2 1.1.1.10 0 msec 4 msec 0 msec
```

Failing both core Vlans 230 and 240 now causes OSPF convergence and the creation of a new routing table excerpt as shown below :

```
msfc-top#
00:12:24: %LINK-3-UPDOWN: Interface Vlan240, changed state to down
00:12:24: %LINK-3-UPDOWN: Interface Vlan230, changed state to down
```

```
msfc-top#
msfc-top#sh ip ro
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
   N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
   * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

     1.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
O       1.1.1.0/24 [110/4] via 192.168.100.39, 00:00:00, Vlan100
O       1.1.1.10/32 [110/8] via 192.168.100.39, 00:00:00, Vlan100
O       1.1.1.39/32 [110/4] via 192.168.100.39, 00:00:00, Vlan100
```

‘msfc-top’ now has an IntraArea route to VIPA 1.1.1.10 via the other S/390 Server, namely Server S39. The Vlan path now used to reach VIPA 1.1.1.10 is shown via the ‘trace’ command :

```
msfc-top#trace 1.1.1.10

Type escape sequence to abort.
Tracing the route to 1.1.1.10

  1 192.168.100.39 4 msec 0 msec 0 msec
  2 192.168.130.2 4 msec 0 msec 4 msec
  3 1.1.1.10 4 msec 0 msec 4 msec
```

This failure thus shows how sub-optimal routing may occur (and is acceptable) when **ALL** Area 1 Core Links between the Catalyst 6500s fail.

3.11 IntraArea Routes fail => InterArea routes used

This scenario again builds upon the previous example. It is designed to show that in the absence of Intra Area routes, InterArea routes will be used if available. Figure 5 illustrates the failure combinations used.

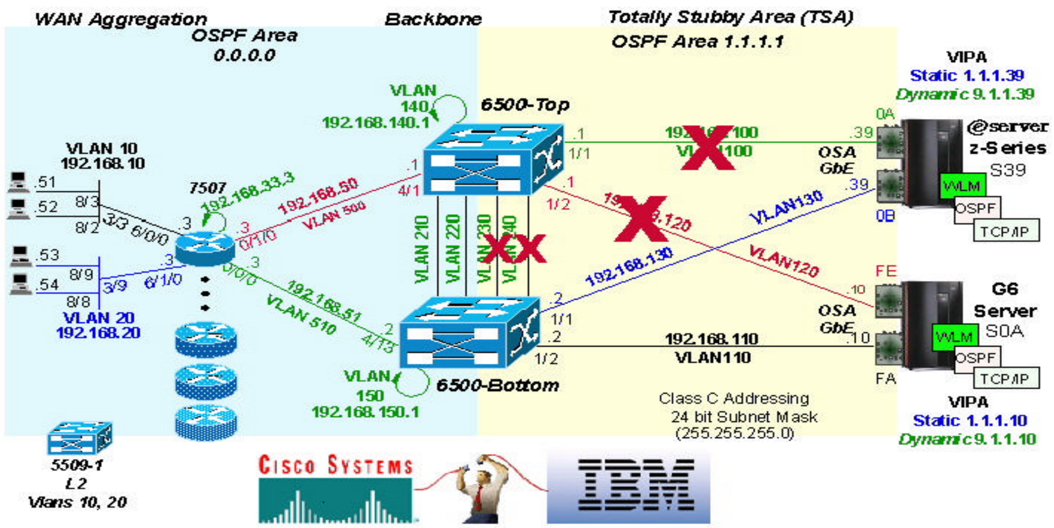

Figure 5 - InterArea Routes

With Vlans 120, 230 and 240 still in the 'down' state, the routing table excerpt for 'msfc-top' looks like this :

```
msfc-top#sh ip ro
```

Gateway of last resort is not set

```
      1.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
 O       1.1.1.0/24 [110/4] via 192.168.100.39, 00:01:39, Vlan100
O       1.1.1.10/32 [110/8] via 192.168.100.39, 00:01:39, Vlan100
 O       1.1.1.39/32 [110/4] via 192.168.100.39, 00:01:39, Vlan100
```

The active route is thus the sub-optimal route via Server S39. The 'trace' command again reflects the sub-optimal path :

```
msfc-top#trace 1.1.1.10
 Tracing the route to 1.1.1.10

   1 192.168.100.39 0 msec 4 msec 0 msec
   2 192.168.130.2 0 msec 0 msec 4 msec
   3 1.1.1.10 0 msec 4 msec 4 msec
```

This test now invokes failure of the active path from 'msfc -top' to VIPA 1.1.1.10, in this case Vlan 100

```
 msfc-top#
00:14:31: %LINK-3-UPDOWN: Interface Vlan100, changed state to down
```

In the absence of IntraArea routes to VIPA 1.1.1.10 over Area 1, the alternate InterArea routes via Area 0 are installed in the routing table after the OSPF convergence process. Note the **'IA'** definition for the route i.e. OSPF InterArea route.

```
msfc-top#sh ip ro
 Codes: C – connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
        D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
        E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
        i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
    * - candidate default, U - per-user static route, o - ODR
        P - periodic downloaded static route

Gateway of last resort is not set

      1.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
 O IA    1.1.1.0/24 [110/5] via 192.168.220.2, 00:00:04, Vlan220
                   [110/5] via 192.168.210.2, 00:00:04, Vlan210
O IA    1.1.1.10/32 [110/5] via 192.168.220.2, 00:00:04, Vlan220
                    [110/5] via 192.168.210.2, 00:00:04, Vlan210
O IA    1.1.1.39/32 [110/5] via 192.168.220.2, 00:00:04, Vlan220
                    [110/5] via 192.168.210.2, 00:00:04, Vlan210
```

CEF switching will once again determine which of the two equal cost routes to utilize to reach VIPA 1.1.1.10

```
 msfc-top#trace 1.1.1.10
 Tracing the route to 1.1.1.10

   1 192.168.220.2 0 msec
   2 1.1.1.10 0 msec 0 msec 0 msec
```

This failure scenario thus illustrates how InterArea route(s) are used in the absence of an IntraArea route(s) to the destination.

3.12 Dynamic VIPA ‘Takeover’ and ‘Takeback’

In the steady state network, dynamic VIPA address 9.1.1.10 is active on Server S0A and dynamic VIPA address 9.1.1.39 is active on Server S39. Upon failure of Server S0A’s TCPIP stack, Server S39’s TCPIP stack will assume responsibilities for the ‘failed’ dynamic VIPA address 9.1.1.10

Please refer to Figure 6 for an illustration of this failure scenario.

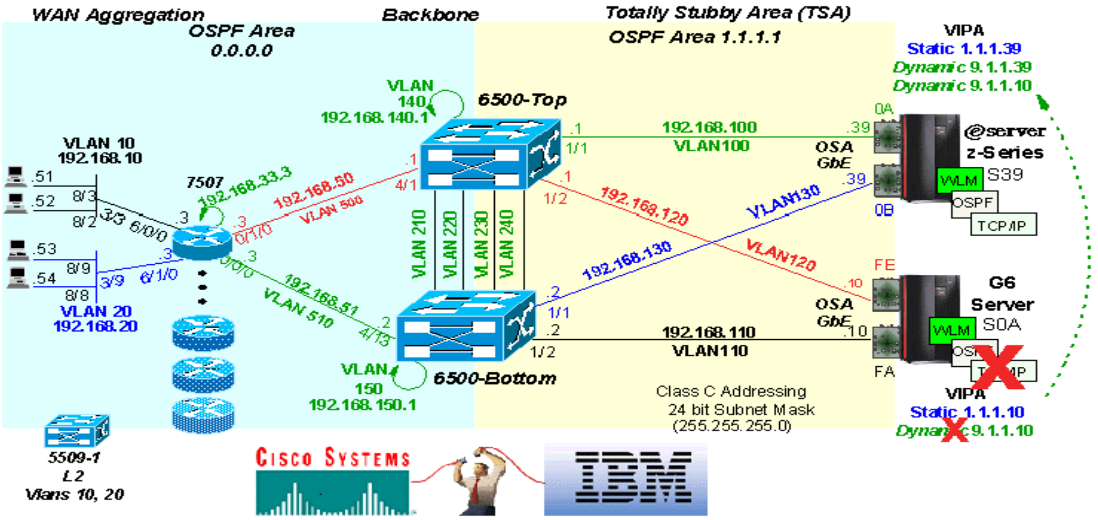

Figure 6 - Dynamic VIPA ‘Takeover’

The following steady state routing table excerpts were derived from the MSFCs :

```
msfc-top#sh ip ro
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
   N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
   E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
   i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
   * - candidate default, U - per-user static route, o - ODR
   P - periodic downloaded static route

Gateway of last resort is not set

     9.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
O       9.1.1.10/32 [110/4] via 192.168.120.10, 00:02:13, Vlan120
O       9.1.1.0/24 [110/4] via 192.168.100.39, 00:02:13, Vlan100
                   [110/4] via 192.168.120.10, 00:02:13, Vlan120
O       9.1.1.39/32 [110/4] via 192.168.100.39, 00:02:13, Vlan100
```

```
 msfc-bottom#sho ip route
 Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
```

```
        D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

9.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
O       9.1.1.10/32 [110/4] via 192.168.110.10, 17:43:46, Vlan110
O       9.1.1.0/24 [110/4] via 192.168.110.10, 17:43:46, Vlan110
                   [110/4] via 192.168.130.39, 17:43:46, Vlan130
O       9.1.1.39/32 [110/4] via 192.168.130.39, 17:43:46, Vlan130
```

The gateway address of 0.0.0.0 for VIPA **9.1.1.10** on Server S0A indicates that the address is local to the S0A stack. :

3.12.1 Server S0A :

```
D TCPIP,,N,ROUTE                                                      EZZ2500I
NETSTAT CS V2R10 TCPIP 485
 DESTINATION      GATEWAY          FLAGS  REFCNT  INTERFACE
 9.1.1.0          192.168.110.2    UG     000000  GIGFA
 9.1.1.0          192.168.120.1    UG     000000  GIGFE
 9.1.1.10         0.0.0.0          UH     000000  VIPL0901010A
 9.1.1.39         192.168.110.2    UGH    000000  GIGFA
 9.1.1.39         192.168.120.1    UGH    000000  GIGFE
```

Server S0A thus ‘owns’ the VIPA address 9.1.1.10. This is further substantiated via S39’s routing table which shows that VIPA address **9.1.1.10** is accessible via two routes, 192.168.110.1 and 192.168.120.1 respectively.

3.12.2 Server S39 :

```
D TCPIP,,N,ROUTE
EZZ2500I NETSTAT CS V2R10 TCPIP 127

9.1.1.0          192.168.130.2    UG     000000  GIG0B
9.1.1.0          192.168.100.1    UG     000000  GIG0A
9.1.1.10         192.168.130.2    UGH    000000  GIG0B
9.1.1.10         192.168.100.1    UGH    000000  GIG0A
9.1.1.39         0.0.0.0          UH     000000  VIPL09010127
```

Issuing a traceroute from Server S39 to VIPA address 9.1.1.10 shows that the S39 Server needs 2 hops to reach the destination.

```
Trace route to 9.1.1.10 (9.1.1.10)
1   (192.168.130.2) 0 ms 0 ms 0 ms
2   (9.1.1.10) 0 ms 1 ms 0 ms
```

Now the failure condition is invoked by failing the TCP/IP stack on Server S0A. The following update is displayed on the message console of Server S39

EZZ8301I VIPA 9.1.1.10 TAKEN OVER FROM TCPIP ON S0A

The routing table excerpt from Server S39 now illustrates how the Dynamic VIPA takeover has occurred. Both dynamic VIPA addresses, namely 9.1.1.39 and 9.1.1.10, are now being advertised and managed by Server S39.

```
D TCPIP,,N,ROUTE
EZZ2500I NETSTAT CS V2R10 TCPIP 127

9.1.1.10         0.0.0.0          UH     000000  VIPL0901010A
9.1.1.39         0.0.0.0          UH     000000  VIPL09010127
```

This is further reinforced by perusal of the MSFCs' routing tables. Vlans 100 and 130 are the OSA-Express links to Server S39 from the MSFCs.   The actual OSPF cost metric on both MSFCs for 9.1.1.10 and 9.1.1.39 is 4. Server S39 advertises both these routes with a default cost of 3 and the relevant MSFC adds its default cost of 1, hence a cumulative cost of 4. This proves that Server S39 does indeed 'own' both the VIPA addresses.

```
msfc-top#sh ip ro
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

9.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
O       9.1.1.10/32 [110/4] via 192.168.100.39, 00:00:32, Vlan100
O       9.1.1.0/24 [110/4] via 192.168.100.39, 00:00:32, Vlan100
O       9.1.1.39/32 [110/4] via 192.168.100.39, 00:00:32, Vlan100


 msfc-bottom#sho ip ro
 Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
        D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
    N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

9.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
O       9.1.1.10/32 [110/4] via 192.168.130.39, 00:00:51, Vlan130
O       9.1.1.0/24 [110/4] via 192.168.130.39, 00:00:51, Vlan130
O       9.1.1.39/32 [110/4] via 192.168.130.39, 00:00:51, Vlan130
```

One has thus seen how Dynamic VIPA Takeover can be achieved. To complete the process the 'Take Back' process is now invoked by recovering the TCPIP stack on Server S0A. See Figure 7 for an illustration of this process.

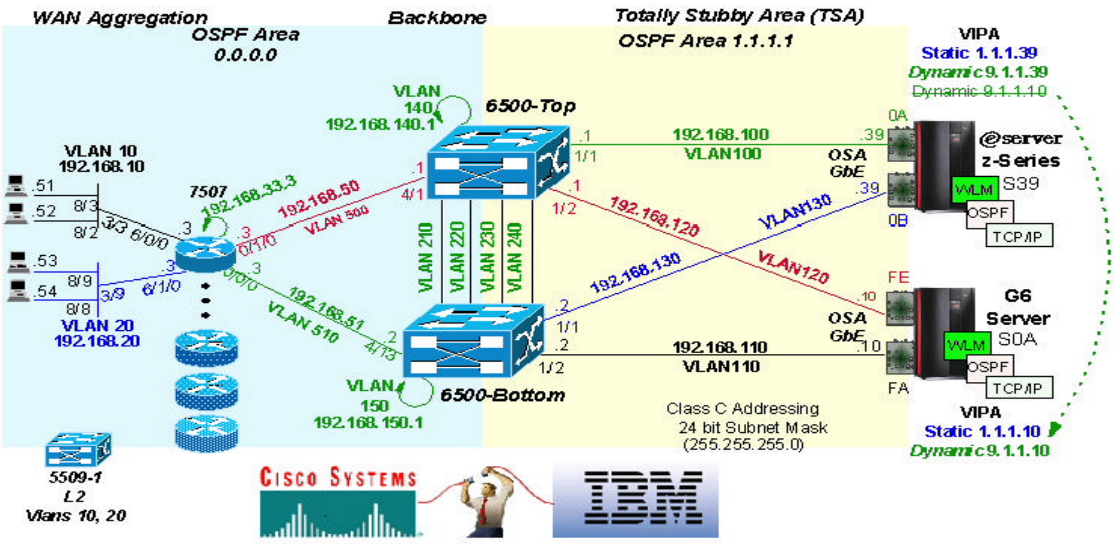

Figure 7 - Dynamic VIPA 'Takeback '

The following message received on Server S0A's message console, illustrates how Server S0A's TCP/IP stack once more takes responsibility for the Dynamic VIPA (9.1.1.10), as expected.

```
EZZ8302I VIPA 9.1.1.10 TAKEN FROM TCPIP ON S39
```

The routing table excerpts from Server S0A and S39 substantiate that Server S0A once more 'owns' the VIPA address 9.1.1.10 :

3.12.3 Server S0A

```
D TCPIP,,N,ROUTE
EZZ2500I NETSTAT CS V2R10 TCPIP 485
 DESTINATION       GATEWAY           FLAGS  REFCNT  INTERFACE
 9.1.1.0           192.168.110.2     UG     000000  GIGFA
 9.1.1.0           192.168.120.1     UG     000000  GIGFE
 9.1.1.10          0.0.0.0           UH     000000  VIPL0901010A
 9.1.1.39          192.168.110.2     UGH    000000  GIGFA
 9.1.1.39          192.168.120.1     UGH    000000  GIGFE
```

3.12.4 Server S39

```
D TCPIP,,N,ROUTE
NETSTAT CS V2R10 TCPIP 127

9.1.1.0          192.168.130.2     UG     000000  GIG0B
9.1.1.0          192.168.100.1     UG     000000  GIG0A
9.1.1.10         192.168.130.2     UGH    000000  GIG0B
9.1.1.10         192.168.100.1     UGH    000000  GIG0A
9.1.1.39         0.0.0.0           UH     000000  VIPL09010127
```

The ‘take back’ procedure has thus been successfully implemented.

This completes the analysis of the various failure domains. It is hoped that these failures have illustrated the robust network design concepts that were envisaged at the outset of this project.

5. Conclusion

This white paper has provided an insight into designing and integrating Catalyst 6500 and IBM’s OSA-Express environments using OSPF. As with all network design, various approaches are applicable. One should bear this in mind when applying these principles to production networks.

Bibliography

- OSPF Network Design Solutions – Thomas M. Thomas II
- Routing TCP/IP : Volume 1 – Jeff Doyle
- Cisco LAN Switching – Kennedy Clark & Kevin Hamilton
- Inteconnections : Second Edition – Bridges, Routers, Switches and Internetworking Protocols - Radia Perlman
- Cisco Internetwork Design – Matthew H. Birkner
- Data Communications, Computer Networks and Open Systems – Fred HalsallAppendices

Appendix A - Configuration Files

msfc-top

```
msfc-top#wr t
Building configuration...

Current configuration:
!
version 12.1
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname msfc-top
!
boot system bootflash:c6msfc-js-mz_121-2_E.bin
boot bootldr bootflash:c6msfc-boot-mz.121-2.E
enable password cisco
!
ip subnet-zero
no ip domain-lookup
!
ip cef
cns event-service server
!
!
!
interface Loopback0
 ip address 192.168.140.1 255.255.255.255
!
interface Vlan1
 ip address 192.168.1.1 255.255.255.0
!
interface Vlan100
 ip address 192.168.100.1 255.255.255.0
 ip ospf hello-interval 3
 ip ospf dead-interval 9
 ip ospf priority 10
!
interface Vlan120
 ip address 192.168.120.1 255.255.255.0
 ip ospf hello-interval 3
 ip ospf dead-interval 9
 ip ospf priority 10
!
interface Vlan210
 ip address 192.168.210.1 255.255.255.0
 ip ospf hello-interval 3
 ip ospf dead-interval 9
!
interface Vlan220
 ip address 192.168.220.1 255.255.255.0
 ip ospf hello-interval 3
 ip ospf dead-interval 9
!
interface Vlan230
 ip address 192.168.230.1 255.255.255.0
 ip ospf hello-interval 3
 ip ospf dead-interval 9
!
```

```
interface Vlan240
 ip address 192.168.240.1 255.255.255.0
 ip ospf hello-interval 3
 ip ospf dead-interval 9
!
interface Vlan500
 ip address 192.168.50.1 255.255.255.0
 ip ospf hello-interval 3
 ip ospf dead-interval 9
!
router ospf 100
 area 1.1.1.1 stub no-summary
 network 192.168.2.0 0.0.0.255 area 0
 network 192.168.50.0 0.0.0.255 area 0
 network 192.168.60.0 0.0.0.255 area 0
 network 192.168.100.0 0.0.0.255 area 1.1.1.1
 network 192.168.120.0 0.0.0.255 area 1.1.1.1
 network 192.168.140.0 0.0.0.255 area 0
 network 192.168.210.0 0.0.0.255 area 0
 network 192.168.220.0 0.0.0.255 area 0
 network 192.168.230.0 0.0.0.255 area 1.1.1.1
 network 192.168.240.0 0.0.0.255 area 1.1.1.1

!
ip classless
no ip http server
!
logging trap alerts
logging 192.168.1.100
snmp-server engineID local 000000090200000021000000
snmp-server community public RO
snmp-server community private RW
!
tftp-server bootflash:c6msfc-boot-mz.121-2.E
!
line con 0
 transport input none
line vty 0 4
 password cisco
 login
 transport input lat pad mop telnet rlogin udptn nasi
!
end
```

msfc-bottom

```
msfc-bottom#wr t
Building configuration...

Current configuration:
!
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname msfc-bottom
!
boot system flash bootflash:c6msfc-js-mz_121-2_E.bin
boot system flash bootflash:c6msfc-isv-mz.121-1.EX.bin
boot bootldr bootflash:c6msfc-boot-mz.121-2.E
enable password cisco
!
username TEST-RME-MSFC-RIGHT
ip subnet-zero
no ip domain-lookup
!
ip cef
cns event-service server
!
!
!
interface Loopback1
 ip address 192.168.150.1 255.255.255.0
!
interface Vlan1
 ip address 192.168.1.2 255.255.255.0
!
interface Vlan110
 ip address 192.168.110.2 255.255.255.0
 ip ospf hello-interval 3
 ip ospf dead-interval 9
 ip ospf priority 10
!
interface Vlan130
 ip address 192.168.130.2 255.255.255.0
 ip ospf hello-interval 3
 ip ospf dead-interval 9
 ip ospf priority 10
!
interface Vlan210
 ip address 192.168.210.2 255.255.255.0
 ip ospf hello-interval 3
 ip ospf dead-interval 9
!
interface Vlan220
 ip address 192.168.220.2 255.255.255.0
 ip ospf hello-interval 3
 ip ospf dead-interval 9
!
interface Vlan230
 ip address 192.168.230.2 255.255.255.0
 ip ospf hello-interval 3
 ip ospf dead-interval 9
!
interface Vlan240
 ip address 192.168.240.2 255.255.255.0
 ip ospf hello-interval 3
 ip ospf dead-interval 9
```

```
!
interface Vlan510
 ip address 192.168.51.2 255.255.255.0
 ip ospf hello-interval 3
 ip ospf dead-interval 9
!
router ospf 100
 area 1.1.1.1 stub no-summary
 network 192.168.51.0 0.0.0.255 area 0
 network 192.168.110.0 0.0.0.255 area 1.1.1.1
 network 192.168.130.0 0.0.0.255 area 1.1.1.1
 network 192.168.150.0 0.0.0.255 area 0
 network 192.168.210.0 0.0.0.255 area 0
 network 192.168.220.0 0.0.0.255 area 0
 network 192.168.230.0 0.0.0.255 area 1.1.1.1
 network 192.168.240.0 0.0.0.255 area 1.1.1.1
!
ip classless
no ip http server
!
logging trap alerts
logging 192.168.1.100
snmp-server engineID local 000000090200000021000000
snmp-server community public RO
snmp-server community private RW
snmp-server community read RW
snmp-server enable traps snmp
snmp-server enable traps hsrp
snmp-server enable traps config
snmp-server enable traps entity
snmp-server enable traps bgp
snmp-server enable traps rsvp
snmp-server enable traps frame-relay
snmp-server enable traps rtr
snmp-server enable traps dlsw
!
banner motd ^CABC rules^C
!
line con 0
 transport input none
line vty 0 4
 password cisco
 login
 transport input lat pad mop telnet rlogin udptn nasi
!
end
```

6500-top

```
6500-top> (enable)

begin
!
# ***** NON-DEFAULT CONFIGURATION *****
!
!
#time: Sat Jan 15 2000, 02:44:05
!
#version 5.5(3)
!
set password $2$0o8Z$GDCVUXu2Kn3mgBDKwF00h1
set enablepass $2$SqUG$jQskZG3AlnVoyIXOzgf3m/
set prompt 6500-top>
!
#system
set system name  6500-top
!
#!
#snmp
set snmp rmon enable
!
#vtp
set vtp domain ibmtest
set vlan 1 name default type ethernet mtu 1500 said 100001 state active
set vlan 100 name VLAN0100 type ethernet mtu 1500 said 100100 state active
set vlan 120 name VLAN0120 type ethernet mtu 1500 said 100120 state active
set vlan 210 name VLAN0210 type ethernet mtu 1500 said 100210 state active
set vlan 220 name VLAN0220 type ethernet mtu 1500 said 100220 state active
set vlan 230 name VLAN0230 type ethernet mtu 1500 said 100230 state active
set vlan 240 name VLAN0240 type ethernet mtu 1500 said 100240 state active
set vlan 500 name VLAN0500 type ethernet mtu 1500 said 100500 state active
set vlan 1002 name fddi-default type fddi mtu 1500 said 101002 state active
set vlan 1004 name fddinet-default type fddinet mtu 1500 said 101004 state
active stp ieee
set vlan 1005 name trnet-default type trbrf mtu 1500 said 101005 state active
stp ibm
set vlan 1003 name token-ring-default type trcrf mtu 1500 said 101003 state
active mode srb aremaxhop 7 stemaxhop 7 backupcrf off
!
#ip
set interface sc0 1 192.168.1.11/255.255.255.0 192.168.1.255

set ip route 0.0.0.0/0.0.0.0         192.168.1.1
!
#rcp
set rcp username cwuser
!
#syslog
set logging server enable
set logging server 192.168.1.100
set logging level cdp 6 default
set logging level mcast 6 default
set logging level dtp 6 default
set logging level earl 6 default
set logging level ip 6 default
set logging level pruning 6 default
set logging level snmp 6 default
set logging level spantree 6 default
set logging level sys 6 default
set logging level tac 6 default
set logging level tcp 6 default
set logging level telnet 6 default
set logging level tftp 6 default
set logging level vtp 6 default
set logging level kernel 6 default
set logging level filesys 6 default
set logging level pagp 6 default
set logging level mgmt 6 default
set logging level mls 6 default
set logging level protfilt 6 default
set logging level security 6 default
set logging level radius 6 default
set logging level udld 6 default
set logging level gvrp 6 default
set logging level cops 6 default
set logging level qos 6 default
set logging level acl 6 default
set logging level rsvp 6 default
set logging level ld 6 default
set logging level privatevlan 6 default
!
#set boot command
set boot config-register 0x2102
set boot system flash slot0:cat6000-sup.5-5-3.bin
!
#mls
set mls enable ipx
!
#port channel
set port channel 1/1-2 32
!
# default port status is enable
!
!
#module 1 : 2-port 1000BaseX Supervisor
set vlan 100  1/1
set vlan 120  1/2
set trunk 1/1  off negotiate 1-1005
set trunk 1/2  off negotiate 1-1005
set spantree portfast    1/1-2 enable
set port channel 1/1-2 mode off
!
#module 2 empty
!
#module 3 empty
!
#module 4 : 48-port 10/100BaseTX Ethernet
set vlan 500  4/1,4/3
set vlan 600  4/2
set port disable    4/14

set port speed      4/1-2,4/13  100
set port duplex     4/1-2,4/13  full
!
#module 5 : 48-port 10/100BaseTX Ethernet
set vlan 2    5/48
set vlan 160  5/1
set port disable    5/1

set spantree portfast    5/48 enable
!
#module 6 : 8-port 1000BaseX Ethernet
set vlan 2    6/5-8
set vlan 210  6/1
set vlan 220  6/2
set vlan 230  6/3
set vlan 240  6/4
set udld enable 6/6
set spantree portfast    6/7-8 enable
!
#module 15 : 1-port Multilayer Switch Feature Card
!
#module 16 empty
end
```

Cisco Equipment Configurations

6500-bottom

```
6500-bottom> (enable)

begin
!
# ***** NON-DEFAULT CONFIGURATION *****
!
!
#time: Wed Nov 1 2000, 14:45:02
!
#version 5.5(3)
!
set password $2$0o8Z$GDCVUXu2Kn3mgBDKwF00h1
set enablepass $2$CBqb$emYj5ImVlOCgbNQTg.TC31
set prompt 6500-bottom>
!
#system
set system name  6500-bottom
!
#!
#snmp
set snmp rmon enable
!
#vtp
set vtp domain ibmtest
set vtp mode transparent
set vlan 1 name default type ethernet mtu 1500 said 100001 state active
set vlan 110 name VLAN0110 type ethernet mtu 1500 said 100110 state active
set vlan 130 name VLAN0130 type ethernet mtu 1500 said 100130 state active
set vlan 210 name VLAN0210 type ethernet mtu 1500 said 100210 state active
set vlan 220 name VLAN0220 type ethernet mtu 1500 said 100220 state active
set vlan 230 name VLAN0230 type ethernet mtu 1500 said 100230 state active
set vlan 240 name VLAN0240 type ethernet mtu 1500 said 100240 state active
set vlan 510 name VLAN0510 type ethernet mtu 1500 said 100510 state active
set vlan 1002 name fddi-default type fddi mtu 1500 said 101002 state active
set vlan 1004 name fddinet-default type fddinet mtu 1500 said 101004 state
active stp ieee
set vlan 1005 name trnet-default type trbrf mtu 1500 said 101005 state active
stp ibm
set vlan 1003 name token-ring-default type trcrf mtu 1500 said 101003 state
active mode srb aremaxhop 7 stemaxhop 7 backupcrf off
!
#ip
set interface sc0 1 192.168.1.12/255.255.255.0 192.168.1.255

set interface sl0 9.67.141.250 255.255.255.240
set ip route 0.0.0.0/0.0.0.0         192.168.1.1
!
#rcp
set rcp username cwuser
!
#spantree
#vlan 1
set spantree disable    1
#vlan 110
set spantree disable    110
```

```
#vlan 130
set spantree disable    130
#vlan 210
set spantree disable    210
#vlan 220
set spantree disable    220
#vlan 230
set spantree disable    230
#vlan 240
set spantree disable    240
#vlan 510
set spantree disable    510
#vlan 1003
set spantree disable    1003
#vlan 1005
set spantree disable    1005
!
#syslog
set logging server enable
set logging server 192.168.1.100
set logging level cdp 6 default
set logging level mcast 6 default
set logging level dtp 6 default
set logging level earl 6 default
set logging level ip 6 default
set logging level pruning 6 default
set logging level snmp 6 default
set logging level spantree 6 default
set logging level sys 6 default
set logging level tac 6 default
set logging level tcp 6 default
set logging level telnet 6 default
set logging level tftp 6 default
set logging level vtp 6 default
set logging level kernel 6 default
set logging level filesys 6 default
set logging level pagp 6 default
set logging level mgmt 6 default
set logging level mls 6 default
set logging level protfilt 6 default
set logging level security 6 default
set logging level radius 6 default
set logging level udld 6 default
set logging level gvrp 6 default
set logging level cops 6 default
set logging level qos 6 default
set logging level acl 6 default
set logging level rsvp 6 default
set logging level ld 6 default
set logging level privatevlan 6 default
!
#set boot command
set boot config-register 0x102
set boot system flash bootflash:cat6000-sup.5-5-3.bin
!
#mls
set mls enable ipx
set mls nde version 8
!
#port channel
set port channel 1/1-2 1
!
# default port status is enable
!
!
#module 1 : 2-port 1000BaseX Supervisor
```

```
set vlan 110  1/2
set vlan 130  1/1
set trunk 1/1  off negotiate 1-1005
set trunk 1/2  off negotiate 1-1005
set spantree portfast    1/1-2 enable
set port channel 1/1-2 mode off
!
#module 2 empty
!
#module 3 : 4-port Multilayer Switch
clear trunk 3/1  1-1005
set trunk 3/1  off negotiate
clear trunk 3/2  1-1005
set trunk 3/2  off negotiate
clear trunk 3/3  1-1005
set trunk 3/3  off negotiate
clear trunk 3/4  1-1005
set trunk 3/4  off negotiate
!
#module 4 : 48-port 10/100BaseTX Ethernet
set vlan 510  4/13
set vlan 610  4/14
set port disable    4/2

set port speed      4/13-14  100
set port duplex     4/13-14  full
!
#module 5 : 48-port 10/100BaseTX Ethernet
set vlan 3    5/18
set vlan 170  5/1
set port disable    5/1,5/46

set spantree portfast    5/18 enable
!
#module 6 : 8-port 1000BaseX Ethernet
set vlan 3    6/6-8
set vlan 210  6/1
set vlan 220  6/2
set vlan 230  6/3
set vlan 240  6/4
set vlan 1    6/5
set port negotiation 6/5 disable
set udld enable 6/6
!
#module 15 : 1-port Multilayer Switch Feature Card
!
#module 16 empty
!
#switch port analyzer
set span 1/1 6/5 both inpkts disable multicast enable learning enable create
end
6500-bottom> (enable)
```

7507-ACCESS

```
7507-ACCESS#

Current configuration:
!
version 12.0
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
service udp-small-servers
service tcp-small-servers
!
hostname 7507-ACCESS
!
boot system flash slot1:rsp-jsv-mz_120-7_T
logging buffered 4096 debugging
no logging console
enable password cisco
!
!
!
!
!
microcode CIP flash cip27-4
microcode reload
ip subnet-zero
ip flow-cache feature-accelerate
ip telnet source-interface FastEthernet0/1/0
no ip domain-lookup
!
ip cef distributed
cns event-service server
!
!
!
interface Loopback1
 ip address 192.168.33.3 255.255.255.0
 no ip redirects
 no ip directed-broadcast
 ip route-cache flow
 no ip mroute-cache
!
interface FastEthernet0/0/0
 ip address 192.168.51.3 255.255.255.0
 no ip redirects
 no ip directed-broadcast
 ip route-cache flow
 no ip route-cache cef
 ip route-cache distributed
 ip ospf hello-interval 3
 ip ospf dead-interval 9
 no ip mroute-cache
 load-interval 30
 full-duplex
!
interface FastEthernet0/1/0
 ip address 192.168.50.3 255.255.255.0
 no ip redirects
 no ip directed-broadcast
 ip route-cache flow
 no ip route-cache cef
 ip route-cache distributed
 ip ospf hello-interval 3
 ip ospf dead-interval 9
 no ip mroute-cache
 load-interval 30
 full-duplex
!
interface Channel1/0
 no ip address
 no ip directed-broadcast
 no keepalive
 shutdown
!
interface Channel1/1
 no ip address
 no ip directed-broadcast
 no keepalive
 shutdown
!
interface Channel1/2
 no ip address
 no ip directed-broadcast
 ip route-cache flow
 no ip mroute-cache
 no keepalive
 shutdown
!
interface Channel4/0
 no ip address
 no ip directed-broadcast
 no keepalive
 shutdown
!
interface Channel4/1
 no ip address
 no ip directed-broadcast
 no keepalive
 shutdown
!
interface Channel4/2
 no ip address
 no ip directed-broadcast
 ip route-cache flow
 no keepalive
 shutdown
!
interface FastEthernet6/0/0
 ip address 192.168.10.70 255.255.255.0
 no ip redirects
 no ip directed-broadcast
 ip route-cache flow
 no ip route-cache cef
 ip route-cache distributed
 ip ospf hello-interval 3
 ip ospf dead-interval 9
 no ip mroute-cache
 load-interval 30
 full-duplex
!
interface FastEthernet6/1/0
 no ip address
 no ip redirects
 no ip directed-broadcast
 ip route-cache flow
 no ip route-cache cef
 ip route-cache distributed
 ip ospf hello-interval 3
 ip ospf dead-interval 9
 no ip mroute-cache
 load-interval 30
 shutdown
 full-duplex
!
router ospf 100
 log-adjacency-changes
 network 192.168.10.0 0.0.0.255 area 0
 network 192.168.33.0 0.0.0.255 area 0
 network 192.168.50.0 0.0.0.255 area 0
 network 192.168.51.0 0.0.0.255 area 0
!
ip classless
no ip http server
!
snmp-server engineID local 0000000902000030B6C17800
snmp-server community public RO
snmp-server packetsize 2048
snmp-server enable traps casa
!
!
line con 0
 transport input none
line aux 0
line vty 0 4
 exec-timeout 0 0
 password cisco
 login
!
end
```

IBM Server Configuration Files

S/390 Server S39

VTAM TRLEs

For OSA Chpid 0A on the S39 S/390 Server

```
TRL      VBUILD TYPE=TRL
*********************************************************************
*
* PORTNAME MUST MATCH THE PORTNAME IN THE TCP/IP GIGABIT DEFINITION
*
*********************************************************************
***
***  DEFINITION USING CHP(0A) DEVICES = 1A0-1AF
***
TRL0A00  TRLE  LNCTL=MPC,                                               X
               READ=1A00,                                               X
               WRITE=1A01,                                              X
               DATAPATH=(1A02,1A03),                                    X
               PORTNAME=GIGP0A,                                         X
               MPCLEVEL=QDIO
```

For OSA chpid 0B on the S0A S/390 Server

```
TRL      VBUILD TYPE=TRL
*********************************************************************
*
* PORTNAME MUST MATCH THE PORTNAME IN THE TCP/IP GIGABIT DEFINITION
*
*********************************************************************
***
***  DEFINITION USING CHP(0B) DEVICES = 1B00-1B0F
***
TRL0B00  TRLE  LNCTL=MPC,                                               X
               READ=1B00,                                               X
               WRITE=1B01,                                              X
               DATAPATH=(1B02,1B03,1B04,1B05),                          X
               PORTNAME=GIGP0B,                                         X
               MPCLEVEL=QDIO
```

TCP/IP Profile - S/390 Server S39

```
; TCPIP.PROFILE.TCPIP
; ===================
; Flush the ARP tables every 20 minutes.
ARPAGE 20
; ----------------------------------------------------------------------
DATASETPREFIX TCPIP
; ----------------------------------------------------------------------
TELNETPARMS
    PORT 23
    INACTIVE 600
    TIMEMARK 600
    SCANINTERVAL 120
WLMCLUSTERNAME TN3270E ENDWLMCLUSTERNAME
ENDTELNETPARMS
; -------------------------------------------------------------------
;
  IPCONFIG
```

```
     SYSPLEXROUTING
     MULTIPATH
     DATAGRAMFWD
     VARSUBNETTING
     SOURCEVIPA
; -------------------------------------------------------------------
KEEPALIVEOPTIONS
      INTERVAL 20
ENDKEEPALIVEOPTIONS
; -------------------------------------------------------------------
; Reserve low ports for servers
TCPCONFIG        RESTRICTLOWPORTS
     TCPSENDBFRSIZE 262144
     TCPRCVBUFRSIZE 262144
; -------------------------------------------------------------------
UDPCONFIG        RESTRICTLOWPORTS
; ----------------------------------------------------------------------
; AUTOLOG the following servers.
AUTOLOG 5
    FTPD JOBNAME FTPD1       ; FTP Server
;   LPSERVE                  ; LPD Server
;   NAMESRV                  ; Domain Name Server
;   NCPROUT                  ; NCPROUTE Server
    PORTMAP                  ; Portmap Server
;   OROUTED                  ; RouteD Server
    OMPROUTE                 ; OSPF ROUTER
;   WEBSERVE                 ; WEBSERVER
;   RXSERVE                  ; Remote Execution Server
;   SMTP                     ; SMTP Server
;   OSNMPD                   ; SNMP Agent Server
;   SNMPQE                   ; SNMP Client
;   TCPIPX25                 ; X25 Server
ENDAUTOLOG
; ----------------------------------------------------------------------
;
PORT
     7 UDP MISCSERV            ; Miscellaneous Server
     7 TCP MISCSERV
     9 UDP MISCSERV
     9 TCP MISCSERV
    19 UDP MISCSERV
    19 TCP MISCSERV
    20 TCP OMVS   NOAUTOLOG    ; FTP Server Control
    21 TCP OMVS                ; FTP Server Data
    23 TCP INTCLIEN            ; Telnet Server
;   25 TCP SMTP                ; SMTP Server
;   53 TCP NAMESRV             ; Domain Name Server
;   53 UDP NAMESRV             ; Domain Name Server
    80 TCP WEBSERVE            ; WEBSERVER
   111 TCP OMVS                ; Portmap Server
   111 UDP OMVS                ; Portmap Server
;  135 UDP LLBD                ; NCS Location Broker
;  161 UDP OSNMPD              ; SNMP Agent
;  162 UDP SNMPQE              ; SNMP Query Engine
;  512 TCP RXSERVE             ; Remote Execution Server
;  514 TCP RXSERVE             ; Remote Execution Server
;  515 TCP LPSERVE             ; LPD Server
   520 UDP OMVS                ; RouteD Server
;  580 UDP NCPROUT             ; NCPROUTE Server
;  750 TCP MVSKERB             ; Kerberos
;  750 UDP MVSKERB             ; Kerberos
;  751 TCP ADM@SRV             ; Kerberos Admin Server
;  751 UDP ADM@SRV             ; Kerberos Admin Server
  1500 TCP ADSM                ; ADSM Server Using TCP/IP Comm
  3000 TCP CICSTCP             ; CICS Socket
;-------------------------------------------------------------------
```

```
;Dynamic VIPA Definition
;
  VIPADYNAMIC
     VIPADEFINE     255.255.255.0    9.1.1.39
     VIPABACKUP 10                   9.1.1.10
  ENDVIPADYNAMIC
;
;
; -----------------------------------------------------------------------
; Hardware definitions:
;
;--------------------------------------
; OSA Chpid 0A
;
DEVICE GIGP0A MPCIPA PRIROUTER
LINK GIG0A IPAQGNET GIGP0A
;--------------------------------------
; OSA Chpid 0B
;
DEVICE GIGP0B MPCIPA SECROUTER
LINK GIG0B IPAQGNET GIGP0B
;----------------------------------------
; Static VIPA
;
  DEVICE DEVIPA1 VIRTUAL 0
  LINK LNKVIPA1 VIRTUAL 0 DEVIPA1
;---------------------------------------
;
; HOME Internet addresses of each link in the host.
;
HOME
;----------------------------
;   "Server Vipa"
    1.1.1.39        LNKVIPA1
;----------------------------
;  Real Address of the 0A OSA GbE on S39
  192.168.100.39  GIG0A
;----------------------------
;  Real Address of the 0B OSA GbE on S39
  192.168.130.39  GIG0B
;----------------------------
; Using OSPF gateway is not needed
;
;GATEWAY
; -----------------------------------------------------------------------
;  Network        First          Link   Packet      Subnet    Subnet
;                  hop           name    size       mask      value
TRANSLATE
; A null translate statement issues the warning message EZZ0323I
; -----------------------------------------------------------------------
ITRACE OFF
; -----------------------------------------------------------------------
;
ASSORTEDPARMS
; NOFWD
  RESTRICTLOWPORTS
ENDASSORTEDPARMS
; RESTRICTLOWPORTS issues the informational message EZZ0338I
;
BEGINVTAM
    ; Define logon mode tables to be the defaults shipped with the
    ; latest level of VTAM
    ; Define the LUs to be used for general users.
  DEFAULTLUS
      TCP00001..TCP00300
ENDDEFAULTLUS
```

```
  LUSESSIONPEND    ; On termination of a Telnet server connection,
                   ; the user will revert to the DEFAULTAPPL
  ALLOWAPPL *      ; Allow all applications that have not been
ENDVTAM
; ------------------------------------------------------------------------
;
; Start all the defined devices.
;--- S39
  START GIGP0A
  START GIGP0B
```

OMPROUTE.CONF - S/390 Server S39

```
; OMPROUTE Configuration file for OSPF
;
; AREA
;    Sets the OSPF AREA.  If no areas are defined, the router software
;    assumes that all the router's directly attached networks belong to
;    the backbone area (area ID 0.0.0.0)
;
 AREA
   Area_Number=1.1.1.1
   Authentication_Type=None
   Stub_area=yes
;  Stub_Default_Cost=5
;  Import_Summaries=No;
;
;
; COMPARISON
;    Tells the Router where external routes fit into the OSPF hierarchy.
;    Type 1 (Type1) external metrics are equivalent to the link state
;    metric
;    Type 2 (Type2) external metrics are greater than the cost of any
;    path internal to the AS.  Use of Type2 assumes routing between
;    automonous systems is the major cost of routing a packet and
;    eliminates the need for conversion of external costs to internal
;    link state metrics.
 Comparison=Type2;
;
;OSPF_INTERFACE
;    Sets the OSPF parameters for the router's network interface.
;    This statement needs to be replicated for each IP interface over
;    which OSPF will operate.
;
OSPF_INTERFACE
   IP_address=9.1.1.39
   NAME=DynVIPAAddress
   Subnet_mask=255.255.255.0
   Demand_Circuit=no
   Attaches_To_Area=1.1.1.1
   MTU=1500
   Retransmission_Interval=5
   Transmission_Delay=1
   Router_Priority=0
   Hello_Interval=3
   Dead_Router_Interval=9
   Cost0=3
;
;
OSPF_INTERFACE
   IP_address=192.168.100.39
   NAME=GIG0A
   Subnet_mask=255.255.255.0
   Demand_Circuit=no
   Attaches_To_Area=1.1.1.1
   MTU=1500
   Retransmission_Interval=5
   Transmission_Delay=1
   Router_Priority=0
   Hello_Interval=3
   Dead_Router_Interval=9
   Cost0=3
;
OSPF_INTERFACE
   IP_address=192.168.130.39
   NAME=GIG0B
   Subnet_mask=255.255.255.0
   Demand_Circuit=no
   Attaches_To_Area=1.1.1.1
   MTU=1500
   Retransmission_Interval=5
   Transmission_Delay=1
   Router_Priority=0
   Hello_Interval=3
   Dead_Router_Interval=9
   Cost0=3
;
OSPF_INTERFACE
   IP_address=1.1.1.39
   NAME=LNKVIPA1
   Subnet_mask=255.255.255.0
   Demand_Circuit=no
   Attaches_To_Area=1.1.1.1
   MTU=1500
   Retransmission_Interval=5
   Transmission_Delay=1
   Router_Priority=0
   Hello_Interval=3
   Dead_Router_Interval=9
   Cost0=3
;
;VIRTUAL_LINK
;     Configure virtual links between any two area border routers.
;     To maintain connectivity you must have all your backbone routers
;     interconnected by either permanent or virtual links.
;VIRTUAL_LINK
;   Virtual_Endpoint_RouterID=
;   Links_Transit_Area=
;   Retransmission_Interval=10
;   Transmission_Delay=5
;   Hello_Interval=30
;   Dead_Router_Interval=180;
; Authentication_Key=
;
;
;ROUTERID
;   Every router in an OSPF routing domain must be assigned a unique
;   32-bit router-id.
 RouterID=9.1.1.39
;
 AS_BOUNDARY_ROUTING
    Import_RIP_Routes=No
    Import_Static_Routes=YES
    Import_Direct_Routes=YES
    Import_Subnet_Routes=Yes
    Originate_Default_Route=No
    Originate_as_Type=2
    Default_Route_Cost=1
;   Default_Forwarding_Address =
;
;
;RANGE
```

```
;   IP_Address=
;   Subnet_Mask=255.255.0.0
;   Area_Number=0.0.0.0
;   Advertise=YES
;
;DEMAND_CIRCUIT=YES
;
;ORIGINATE_RIP_DEFAULT
;   Condition=Always
;   Cost=1
;RIP_INTERFACE
;  see config guide for all the tags
;
;ACCEPT_RIP_ROUTE
;  Accept_RIP_Route IP_Address=
;
;DEFAULT_ROUTE
;  Default_Route Name=<interface name> Next_Hop=<ip_address>
;
;INTERFACE
;  IP_Address=
;  Name=
;  Subnet_Mask=255.255.0.0
;  MTU=1500
```

S/390 Server S0A

VTAM TRLEs - S/390 Server S0A

For OSA chpid FA on the S0A S/390 Server

```
TRL      VBUILD TYPE=TRL
**********************************************************************
* START THIS ON S30 - TRLE STMT DEFINES THE OSA
*
*
**********************************************************************
***
***  DEFINITION USING CHP(FA) DEVICES = 550-55F
***
TRLFA00  TRLE  LNCTL=MPC,                                              X
               READ=550,                                               X
               WRITE=551,                                              X
               DATAPATH=(552,553,554,555),                             X
               PORTNAME=GIGPFA,                                        X
               MPCLEVEL=QDIO
```

For OSA chpid FE on the S0A S/390 Server

```
TRL      VBUILD TYPE=TRL
**********************************************************************
* START THIS ON S30 - TRLE STMT DEFINES THE OSA
*
*
**********************************************************************
***
***  DEFINITION USING CHP(FE) DEVICES = 540-54F
***
TRLFE00  TRLE  LNCTL=MPC,                                              X
               READ=540,                                               X
               WRITE=541,                                              X
               DATAPATH=(542,543,544,545),                             X
               PORTNAME=GIGPFE,                                        X
               MPCLEVEL=QDIO
```

TCP/IP Profile - S/390 Server S0A

```
; TCPIP.PROFILE.TCPIP
; ==================
; Flush the ARP tables every 20 minutes.
ARPAGE 20
; ----------------------------------------------------------------------
DATASETPREFIX TCPIP
; ----------------------------------------------------------------------
TELNETPARMS
   PORT 23
   INACTIVE 600
   TIMEMARK 600
   SCANINTERVAL 120
WLMCLUSTERNAME TN3270E ENDWLMCLUSTERNAME
ENDTELNETPARMS
; -------------------------------------------------------------------
;
 IPCONFIG
    SYSPLEXROUTING
    MULTIPATH
    DATAGRAMFWD
    VARSUBNETTING
    SOURCEVIPA
; -------------------------------------------------------------------
KEEPALIVEOPTIONS
     INTERVAL 20
ENDKEEPALIVEOPTIONS
; -------------------------------------------------------------------
; Reserve low ports for servers
TCPCONFIG       RESTRICTLOWPORTS
    TCPSENDBFRSIZE 262144
    TCPRCVBUFRSIZE 262144
; -------------------------------------------------------------------
UDPCONFIG       RESTRICTLOWPORTS
; ----------------------------------------------------------------------
; AUTOLOG the following servers.
AUTOLOG 5
   FTPD JOBNAME FTPD1       ; FTP Server
;   LPSERVE                 ; LPD Server
;   NAMESRV                 ; Domain Name Server
;   NCPROUT                 ; NCPROUTE Server
   PORTMAP                  ; Portmap Server
;   OROUTED                 ; RouteD Server
   OMPROUTE                 ; OSPF ROUTER
;   WEBSERVE                ; WEBSERVER
;   RXSERVE                 ; Remote Execution Server
;   SMTP                    ; SMTP Server
;   OSNMPD                  ; SNMP Agent Server
;   SNMPQE                  ; SNMP Client
;   TCPIPX25                ; X25 Server
ENDAUTOLOG
; ----------------------------------------------------------------------
;
PORT
    7 UDP MISCSERV          ; Miscellaneous Server
    7 TCP MISCSERV
    9 UDP MISCSERV
    9 TCP MISCSERV
   19 UDP MISCSERV
   19 TCP MISCSERV
   20 TCP OMVS   NOAUTOLOG    ; FTP Server Control
   21 TCP OMVS               ; FTP Server Data
;   21 TCP FTPD1             ; FTP Server Data
   23 TCP INTCLIEN           ; Telnet Server
;   25 TCP SMTP              ; SMTP Server
;   53 TCP NAMESRV           ; Domain Name Server
;   53 UDP NAMESRV           ; Domain Name Server
   80 TCP WEBSERVE           ; WEBSERVER
  111 TCP OMVS               ; Portmap Server
  111 UDP OMVS               ; Portmap Server
;  135 UDP LLBD              ; NCS Location Broker
;  161 UDP OSNMPD            ; SNMP Agent
;  162 UDP SNMPQE            ; SNMP Query Engine
;  512 TCP RXSERVE           ; Remote Execution Server
;  514 TCP RXSERVE           ; Remote Execution Server
;  515 TCP LPSERVE           ; LPD Server
  520 UDP OMVS               ; RouteD Server
;  580 UDP NCPROUT           ; NCPROUTE Server
;  750 TCP MVSKERB           ; Kerberos
;  750 UDP MVSKERB           ; Kerberos
;  751 TCP ADM@SRV           ; Kerberos Admin Server
;  751 UDP ADM@SRV           ; Kerberos Admin Server
 1500 TCP ADSM               ; ADSM Server Using TCP/IP Comm
 3000 TCP CICSTCP            ; CICS Socket
;-----------------------------------------------------------------------
;Dynamic VIPA Definition
;
  VIPADYNAMIC
     VIPADEFINE     255.255.255.0    9.1.1.10
     VIPABACKUP 10                   9.1.1.39
  ENDVIPADYNAMIC
;
;
; -----------------------------------------------------------------------
; Hardware definitions:
;
;-------------------------------------
; OSA Chpid FA
;
DEVICE GIGPFA MPCIPA SECROUTER
LINK GIGFA IPAQGNET GIGP0A
;-------------------------------------
; OSA Chpid FE
;
DEVICE GIGPFE MPCIPA PRIROUTER
LINK GIGFE IPAQGNET GIGPFE
;--------------------------------------
; Static VIPA
;
  DEVICE DEVIPA1 VIRTUAL 0
  LINK LNKVIPA1 VIRTUAL 0 DEVIPA1
;-------------------------------------
;
; HOME Internet addresses of each link in the host.
```

```
;
HOME
;----------------------------
;   "Server Vipa"
   1.1.1.10        LNKVIPA1
;----------------------------
;  Real Address of the FE OSA GbE on S0A
  192.168.120.10  GIGFE
;----------------------------
;  Real Address of the FA OSA GbE on S0A
  192.168.110.10  GIGFA
;----------------------------
; Using OSPF gateway is not needed
;
;GATEWAY
; ----------------------------------------------------------------------
;  Network       First        Link   Packet     Subnet    Subnet
;                hop          name    size       mask      value
TRANSLATE
; A null translate statement issues the warning message EZZ0323I
; ----------------------------------------------------------------------
ITRACE OFF
; ----------------------------------------------------------------------
;
ASSORTEDPARMS
; NOFWD
  RESTRICTLOWPORTS
ENDASSORTEDPARMS
; RESTRICTLOWPORTS issues the informational message EZZ0338I
;
BEGINVTAM
   ; Define logon mode tables to be the defaults shipped with the
   ; latest level of VTAM
   ; Define the LUs to be used for general users.
  DEFAULTLUS
     TCP00001..TCP00300
ENDDEFAULTLUS
  LUSESSIONPEND    ; On termination of a Telnet server connection,
                 ; the user will revert to the DEFAULTAPPL
  ALLOWAPPL *      ; Allow all applications that have not been
ENDVTAM
; ----------------------------------------------------------------------
;
; Start all the defined devices.
;--- S39
  START GIGPFE
  START GIGPFA
```

OMPROUTE.CONF - S/390 Server S0A

```
; OMPROUTE Configuration file for OSPF
;
; AREA
;    Sets the OSPF AREA.  If no areas are defined, the router software
;    assumes that all the router's directly attached networks belong to
;    the backbone area (area ID 0.0.0.0)
;
 AREA
   Area_Number=1.1.1.1
   Authentication_Type=None
   Stub_area=yes
;  Stub_Default_Cost=5
;  Import_Summaries=No;
;
;
; COMPARISON
;    Tells the Router where external routes fit into the OSPF hierarchy.
;    Type 1 (Type1) external metrics are equivalent to the link state
;    metric
;    Type 2 (Type2) external metrics are greater than the cost of any
;    path internal to the AS.  Use of Type2 assumes routing between
;    automonous systems is the major cost of routing a packet and
;    eliminates the need for conversion of external costs to internal
;    link state metrics.
 Comparison=Type2;
;
;OSPF_INTERFACE
;    Sets the OSPF parameters for the router's network interface.
;    This statement needs to be replicated for each IP interface over
;    which OSPF will operate.
;
;
OSPF_INTERFACE
   IP_address=192.168.110.10
   NAME=GIGFA
   Subnet_mask=255.255.255.0
   Demand_Circuit=no
   Attaches_To_Area=1.1.1.1
   MTU=1500
   Retransmission_Interval=5
   Transmission_Delay=1
   Router_Priority=0
   Hello_Interval=3
   Dead_Router_Interval=9
   Cost0=3
;
OSPF_INTERFACE
   IP_address=192.168.120.10
   NAME=GIGFE
   Subnet_mask=255.255.255.0
   Demand_Circuit=no
   Attaches_To_Area=1.1.1.1
   MTU=1500
   Retransmission_Interval=5
   Transmission_Delay=1
   Router_Priority=0
   Hello_Interval=3
   Dead_Router_Interval=9
   Cost0=3
;
OSPF_INTERFACE
   IP_address=1.1.1.10
   NAME=LNKVIPA1
   Subnet_mask=255.255.255.0
   Demand_Circuit=no
   Attaches_To_Area=1.1.1.1
   MTU=1500
   Retransmission_Interval=5
   Transmission_Delay=1
   Router_Priority=0
   Hello_Interval=3
   Dead_Router_Interval=9
   Cost0=3
;
OSPF_INTERFACE
```

```
   IP_address=9.1.1.10
   NAME=DynVIPAAddress
   Subnet_mask=255.255.255.0
   Demand_Circuit=no
   Attaches_To_Area=1.1.1.1
   MTU=1500
   Retransmission_Interval=5
   Transmission_Delay=1
   Router_Priority=0
   Hello_Interval=3
   Dead_Router_Interval=9
   Cost0=3
;
;
;
;VIRTUAL_LINK
;     Configure virtual links between any two area border routers.
;     To maintain connectivity you must have all your backbone routers
;     interconnected by either permanent or virtual links.
;VIRTUAL_LINK
;   Virtual_Endpoint_RouterID=
;   Links_Transit_Area=
;   Retransmission_Interval=10
;   Transmission_Delay=5
;   Hello_Interval=30
;   Dead_Router_Interval=180;
; Authentication_Key=
;
;
;ROUTERID
;   Every router in an OSPF routing domain must be assigned a unique
;   32-bit router-id.
 RouterID=9.1.1.10
;
 AS_BOUNDARY_ROUTING
    Import_RIP_Routes=No
    Import_Static_Routes=YES
    Import_Direct_Routes=YES
    Import_Subnet_Routes=Yes
    Originate_Default_Route=No
    Originate_as_Type=2
    Default_Route_Cost=1
;   Default_Forwarding_Address =
;
;
;RANGE
;   IP_Address=
;   Subnet_Mask=255.255.0.0
;   Area_Number=0.0.0.0
;   Advertise=YES
;
;
;
;DEMAND_CIRCUIT=YES
;
;
;ORIGINATE_RIP_DEFAULT
;   Condition=Always
;   Cost=1
;
;
;RIP_INTERFACE
;  see config guide for all the tags
;
;
;ACCEPT_RIP_ROUTE
```

```
;  Accept_RIP_Route IP_Address=
;
;
;DEFAULT_ROUTE
;  Default_Route Name=<interface name> Next_Hop=<ip_address>
;
;
;INTERFACE
;  IP_Address=
;  Name=
;  Subnet_Mask=255.255.0.0
;  MTU=1500
;
```

Appendix B - Reverting to Hybrid Mode from IOS Native Mode

The initial stages of the laboratory setup required reverting Sup/MSFCs from IOS Native mode to Hybrid Cat IOS / Router IOS mode. This process is quite lengthy. The following URL provides details on how to do this :

http://wwwin-rtpdev.cisco.com/wbu-dtl/cosmos/s.booting1.htm

Appendix C – Trademarks and Additional Disclaimers

IBM Corporation
Marketing Communications, Enterprise Systems Group
Route 100
Somers, NY 10589

IBM, IBM logo, OS/390, System/390, S/390, Parallel Sysplex,, and l ^o g logo are registered trademarks of the International Business Machines Corporation (“IBM”).

Printed in the United States of America. All Rights Reserved

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the products or services available in your area.

You can find additional information via IBM’s World Wide Web server at http://www.ibm.com.

IBM Hardware products are manufactured from new parts or new and serviceable used parts. Regardless, our warranty terms apply. Actual performance and environmental costs will vary depending on individual customer configurations and conditions.

Cisco makes no representations of any kind with respect to any IBM products, or the performance of those products in combination with any Cisco products. All sales of Cisco products, including without limitation any applicable warranty terms, will be exclusively governed by the terms of a Cisco customer agreement.

This publication was produced in the United States. Cisco may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local Cisco business contact for information on the products or services available in your area.

You can find additional information via Cisco’s World Wide Web server at http://www.cisco.com.

Actual performance and environmental costs of Cisco products will vary depending on individual customer configurations and conditions.

Cisco, Cisco logo, and Cisco-based trademarks or registered trademarks are the trademarks of Cisco Systems, Inc. or its affiliates in the U.S. and other Countries.

IBM makes no representations of any kind with respect to any Cisco Systems, Inc. (Cisco) products, or the performance of those products in combination with any IBM products. All sales of IBM products will be exclusively governed by the terms of an IBM customer agreement.

All other registered trademarks and trademarks are the properties of their respective companies.