WHITE PAPER

3COM SOLUTIONS: INTEROPERATING WITH CISCO SYSTEMS

CONTENTS

WHY IS INTEROPERABILITY IMPORTANT?

In today's fast paced and highly competitive world, IT is a critical enabler of business. Properly designed business IT systems enable increased adaptability and responsiveness, reduce capital and operational costs and ultimately improve company efficiency. However, to optimize all the benefits of technological innovations, businesses must ensure their networks are designed with the future in mind. It is no longer cost effective to be trapped in wholly proprietary systems where each new advance creates adoption challenges rather than opportunities.

To maintain a competitive edge, companies must invest in best-of-breed technology that integrates or overlays into their existing network, with the best network design invariably resulting in a multi-vendor IT system. While the advantages clearly justify this approach, the added complexity cannot be denied. To minimize this complexity, it is incumbent on vendors and their partners to facilitate the development of multi-vendor networks through the adoption of standards-based protocols promulgated by the IEEE and IETF.

3Com accepts this challenge: as the inventor of the Ethernet standard, the company repeatedly delivers against its mission to build standards-based connectivity, delivering solutions with the latest technologies, such as SIP (Session Initiation Protocol) and Power Over Ethernet (PoE) IEEE 802.3af.

3Com, whose name is partly derived from "compatibility", is committed to interoperability so its customers can fully utilize their existing infrastructure while leveraging the advances in technology and future-proofing their investments.

With a few straightforward and lab-proven examples, this guide will demonstrate 3Com's committment to its customers and interoperability. The key is utilizing standards-based protocols in both 3Com and Cisco products, with all necessary conversion and translation being processed on the Cisco switches and routers. This strategy will provide for seamless integration with 3rd-party equipment, including that from 3Com. Finally, should a 3Com customer so choose, this approach also will assist in migration towards a single, efficient autonomous system with standards-based protocols.

This interoperability guide specifically focuses on two important areas:

- Deploying standards-based OSPF IP routing and integrating it with a Cisco network running EIGRP routing;
- Using 3Com Power over Ethernet LAN switches with Cisco Call Manager IP Telephony system and Cisco IP phones.

INTENDED USERS OF THIS GUIDE

This best practices guide is intended for any technical pre- or post-sales engineer intending to deploy a 3Com network alongside a Cisco LAN or Call Manager IP Telephony system.

A glossary of terms used in this guide appears at the end of this document.

DEPLOYING OSPF ROUTING WITHIN A CISCO EIGRP NETWORK

INTRODUCTION

The Internet developed rapidly, relying upon the common TCP/IP suite as its mainstream protocols. Simultaneously, the expanding choice of routing hardware led to an array of routing protocols so that the many different routers could work together. Emerging during this time were such common unicast routing protocols as Routing Information Protocol (RIP), Open Shortest Path First (OSPF), Intermediate System-to-Intermediate System (IS-IS) and Border Gateway Protocol (BGP), as well as Cisco-proprietary routing protocols, including Enhanced Interior Gateway Routing Protocol (EIGRP).

Choosing a suitable routing protocol is critical in large network design. There is the fundamental requirement of connectivity and a host of secondary factors, including network topology and network management. Every routing protocol also has its own strengths and weaknesses. Some best practices for choosing protocols are outlined as follows:

Compatibility › Compatibility of protocols ensures connectivity and expandability of the network, as different manufacturers can support them while customers gain more choice.

Network topology › Network topology has a direct influence on protocol choice. For example, RIP is unsuitable for a complex network because its coverage is limited; complex networks need more powerful protocols, such as OSPF, EIGRP.

Strength, stability › As the signaling protocol to ensure network connectivity, the routing protocols must be strong and stable. For long periods they must bear various abnormalities that appear in the network, such as hardware error or extremely heavy loads. And because routers are located at the decision-making points of the network, they themselves must not cause unpredictable network behavior.

Best path selection › Routing protocols aim at finding the best path in the network to ensure connectivity. Each routing protocol has its own standard to judge route quality, and uses such parameters as next hop number, bandwidth and delay. Generally, these parameters are quantified with a metric for route data. To ensure the best network path, we should select the proper measurement for different network environments.

Management and security › Dividing the autonomous system into different areas eases network management and decreases the possibility of routes being unreachable. Thus, open standard and robust protocols are best able to handle management and security requirements.

With these critical factors in mind, customers are considering whether the IETF-recommended OSPF (Open Shortest Path First) or Cisco's proprietary EIGRP is better suited to their future network. Both of these dynamic routing protocols support fast convergence from link failure, both are loop-free, secure and take up little bandwidth, and both are widely used in today's network.

However, going forward the OSPF routing protocol has considerable advantages over its proprietary EIGRP counterpart. The IETF developed OSPF as an interior routing protocol for IP networks and through such updates as OSPFv3, which supports IPv6, the protocol offers unmatched future-proofing for major IT deployments.

The table below compares the advantages and disadvantages of OSPF and EIGRP.

	OSPF	EIGRP
Industry Standards	IETF recommendation; supported by most vendors; IPv6 support for OSPF V3	Cisco-proprietary; not supported by other vendors
Popularity	Most popular IGP in the world, mature	Declining popularity; few new networks being designed with EIGRP
Algorithm	Shortest Path First algorithm, fast convergence, loop-free	DUAL algorithm could be in stuck in active status (DUAL-3-SIA error) for its querying to spread out over big network
Topology	Build hierarchical, scalable network through partitioning into ASs	Limited scalability; cannot build a hierarchical network
Future	Supports OSPF traffic engineering	Does not support traffic engineering

Table 1: Comparison of OSPF and EIGRP advantages and disadvantages

HOW OSPF WORKS

Overall, OSPF routers collect and forward the link state data of autonomous systems (AS), then use this data in running the shortest path first algorithm to calculate routes. To collect and forward link data, each OSPF router employs a link state advertisement (LSA), which describes the local network connection state, including valid interfaces and reachable neighbors. This description is advertised across the AS, and the collection of LSA data thus forms the link state database (LSDB). Because each LSA describes the surrounding network topology of a router, the LSDB taken as a whole thereby reflects the AS network topology.

To calculate routes without loops, the router refers to the LSDB and runs the shortest path first algorithm to build the shortest path tree, taking itself as the starting point. This shortest path tree gives the route to nodes in the AS and by definition, the route will not have any loops.

Integral to OSPF functionality are the concepts of designated router, backup designated routers, area and routing priority.

- In multi-access networks where there are two or more routers, the network designates a router (DR) to respond to the LSDB synchronization of all routers in the network segment. This frees non-DR routers from that role, thus reducing bandwidth overhead in the same network segment. However, to ensure complete redundancy, OSPF uses the concept of a backup DR, which mirrors the DR, exchanges keep-alives and, in the event of a DR failure, automatically becomes the DR for the multi-access network.
- The OSPF protocol is capable of dividing the AS into different areas according to the topology. When the area border router (ABR) transmits routing information to other areas, it generates the brief LSA with this unit of the network segment, thus decreasing the LSA numbers in the AS as well as the complexity of route calculation.
- The OSPF protocol adopts four classes of routes in order of priority as follows:

  1. Inter-area routing
  2. Intra-area routing
  3. Type 1 exterior routing
  4. Type 2 exterior routing

Inter- and intra-area routes describe the interior network structure of the autonomous system, while the exterior routes describe how to choose routes to destinations outside the AS. Type 1 exterior routes correspond to the information introduced by OSPF from other interior routing protocols; costs of these routes and of the OSFP route itself are comparable. Type 2 exterior routes correspond to the information introduced by OSPF from exterior routing protocols; costs of these routes are much larger than costs of the OSFP route itself, therefore only exterior costs are considered in the calculation.

The OSPF protocol is well developed and especially suitable to enterprise networks. Its advantages include the following:

- Truly loop-free routing, due to the link state and shortest path first algorithm itself.
- Fast convergence quickly transmits routing change information throughout the AS and re-calculates routes.
- Supports equal-cost load balancing.
- Has the capabilities to divide the AS into different areas according to topology. When the ABR transmits routing information to other areas, it generates the brief LSA with the unit of segment, thus decreasing the number of LSAs and the complexity of route calculation (route information will not increase very rapidly in an expanding network).
- Reduces transmission and bandwidth overhead associated with hello packets and other routing information traffic.
- Restricts routes to four classes, providing more reliable routing choice.
- Supports two types of packet authentication modes: clear text authentication mode; and encrypted authentication mode with MD5 algorithm.
- Suits any size network, supporting up to thousands of routers.
- Unlike EIGRP, OSPF is an open standard routing protocol developed by IETF. It is supported by many mainstream vendors, guaranteeing wide compatibility and promising ongoing improvement as a protocol.
- Expands to support traffic engineering because of link-state awareness.

The OSPF protocol is large and complex; its network attributes division of areas requires well-trained personnel. In addition, OSPF does not support unequal load balance; it creates by default the metric of a link, based on the bandwidth of that link, choosing the path with the smallest metric towards the same destination.

DEPLOYMENT EXAMPLE

3Com understands the customer's reluctance in changing, all at once, their entire network's routing protocol from a proprietary one like EIGRP to an open standard protocol like OSPF. The purpose of this example is to provide a migration path which initially allows OSPF to run on the new routers and switches integrating with EIGRP and, at a later time, replacing the old hardware with new, running OSPF. This will protect the customer's investment and allows for more competitive offerings in the future instead of being locked into a proprietary offering from one company.

1. Design the OSPF network

Customers should continue to run EIGRP on all of their Cisco switches. OSPF should be run on one of the Cisco switches that connects to the 3Com switch.

Standard-practice OSPF design guidelines should be followed:

- Determine OSPF network topology (numbers of routers in an area, neighbors for any one router, areas support per router, designated router selection)
- Determine OSPF addressing and route summarization *(optional)*
- Determine OSPF route selection (metrics and load balancing) *(optional)*
- Determine OSPF convergence timers *(optional)*

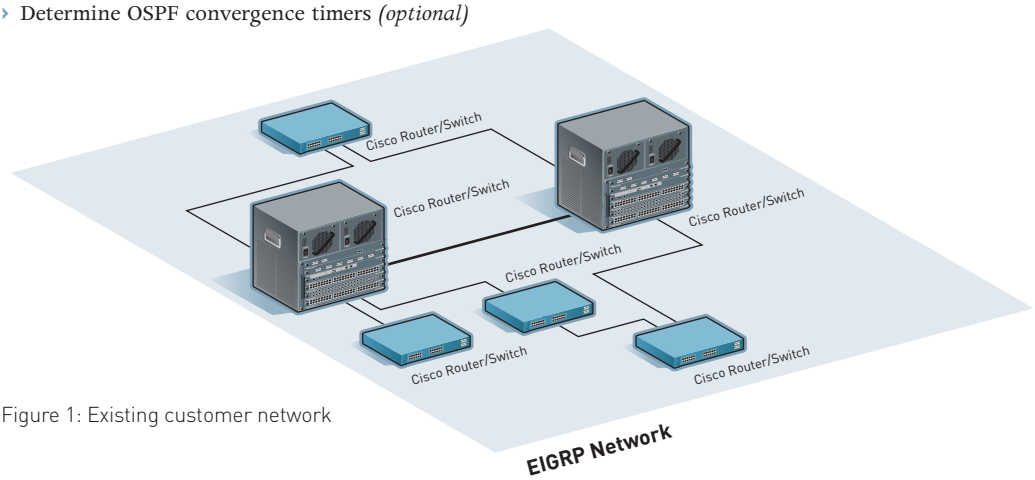

Figure 1: Existing customer network

2. Configure OSPF and EIGRP

Configure the 3Com switches to run OSPF. Configure the Cisco switch to import EIGRP routes into OSPF and import OSPF routes into EIGRP. The 3Com switches will now interoperate with the Cisco network.

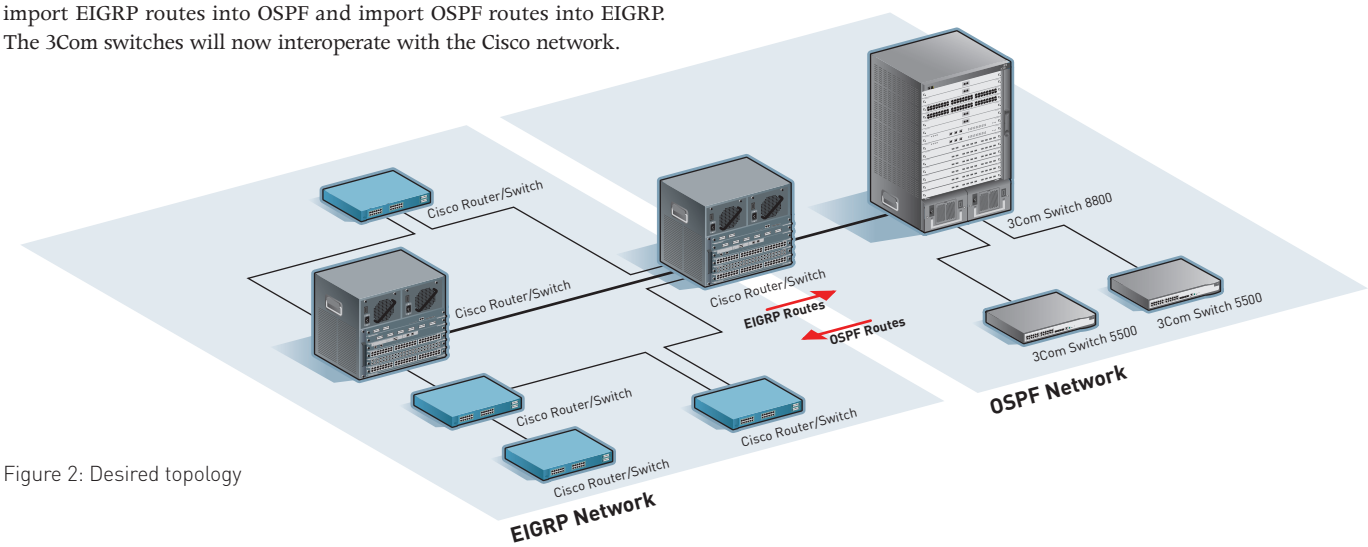

Figure 2: Desired topology

3. Sample OSPF and EIGRP Interoperation

The following example describes a typical scenario of enabling OSPF on 3Com switches, and setting up the translation between EIGRP and OSPF on a single or limited number of Cisco routers or switches wherein 3Com interconnects.

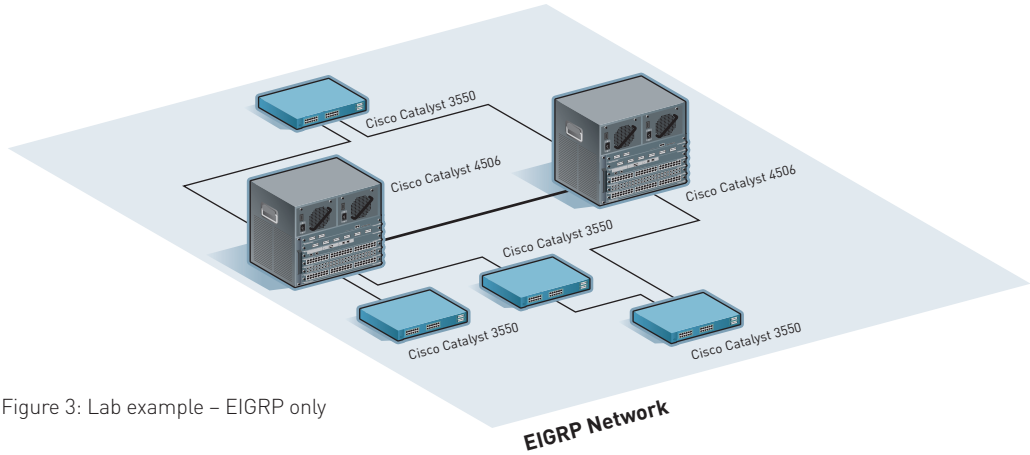

Figure 3: Lab example – EIGRP only

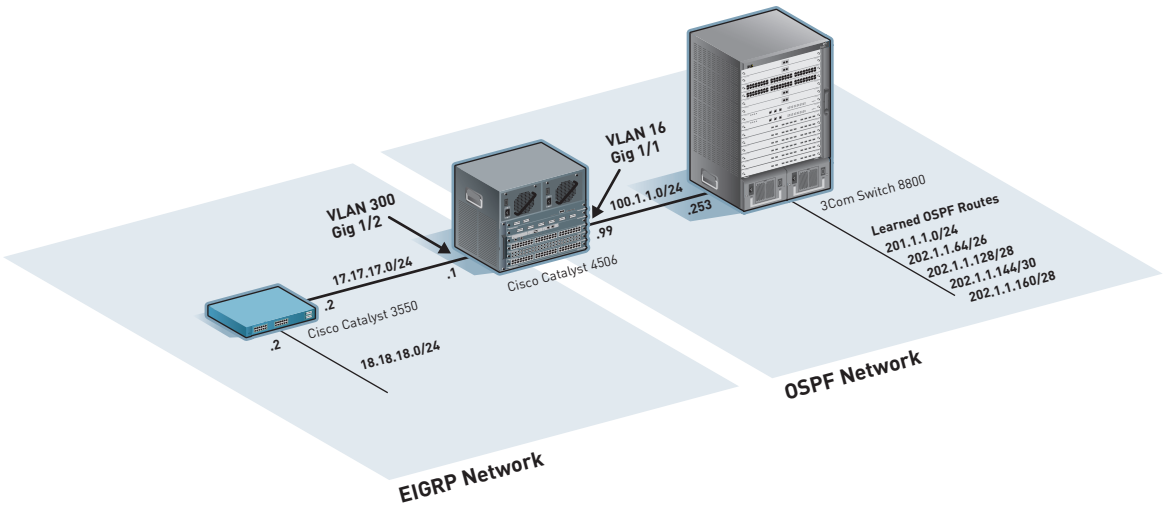

Figure 4: Lab example – EIGRP and OSPF

Catalyst 4506

```
Switch#show version
Cisco Internetwork Operating System Software
IOS (tm) Catalyst 4000 L3 Switch Software (cat4000-I5S-M), Version 12.2(20)EW, EARLY
DEPLOYMENT RELEASE SOFTWARE (fc1)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2004 by cisco Systems, Inc.
Compiled Wed 02-Jun-04 18:21 by hqluong
Image text-base: 0x00000000, data-base: 0x011F88F4

ROM: 12.1(20r)EW1
Dagobah Revision 93, Swamp Revision 6

Switch uptime is 10 minutes
System returned to ROM by reload
System image file is "bootflash:cat4000-i5s-mz.122-20.EW.bin"

cisco WS-C4506 (MPC8245) processor (revision 8) with 524288K bytes of memory.
Processor board ID FOX08230032
MPC8245 CPU at 400Mhz, Supervisor V
Last reset from Reload
12 Virtual Ethernet/IEEE 802.3  interface(s)
24 FastEthernet/IEEE 802.3 interface(s)
2 Gigabit Ethernet/IEEE 802.3 interface(s)
511K bytes of non-volatile configuration memory.

Configuration register is 0x2

Switch#show running-config
Building configuration...

Current configuration : 3932 bytes
!
version 12.2
.
.
.
!
vlan 16,300
!
.
.
.
!
interface GigabitEthernet1/1
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 16
 switchport mode trunk
!
interface GigabitEthernet1/2
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 300
 switchport mode trunk
!
.
.
.
.
!
interface Vlan16
 ip address 100.1.1.99 255.255.255.0
 ip pim dense-mode
 ip ospf cost 10
!
interface Vlan300
 ip address 17.17.17.1 255.255.255.0
 ip pim dense-mode
!
.
.
.
.
!
router eigrp 1
 redistribute connected
 redistribute ospf 1
 redistribute odr
 network 17.17.17.0 0.0.0.255
 default-metric 10000 100 255 100 1500
 no auto-summary
!
router ospf 1
 log-adjacency-changes
 area 0.0.0.0
 redistribute connected subnets
 redistribute eigrp 1 subnets
 network 100.1.1.0 0.0.0.255 area 0.0.0.0
!
.
.
.
!
end

Switch#show ip route
Codes: C - connected, S - static, 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, su - IS-IS summary, 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

     17.0.0.0/24 is subnetted, 1 subnets
C       17.17.17.0 is directly connected, Vlan300
     100.0.0.0/24 is subnetted, 1 subnets
C       100.1.1.0 is directly connected, Vlan16
     18.0.0.0/24 is subnetted, 1 subnets
D       18.18.18.0 [90/3072] via 17.17.17.2, 00:14:17, Vlan300
O IA 201.1.1.0/24 [110/20] via 100.1.1.253, 00:11:56, Vlan16
     202.1.1.0/24 is variably subnetted, 4 subnets, 3 masks
O IA    202.1.1.128/28 [110/30] via 100.1.1.253, 00:11:57, Vlan16
O E2    202.1.1.144/30 [110/1] via 100.1.1.253, 00:11:57, Vlan16
O E2    202.1.1.160/28 [110/1] via 100.1.1.253, 00:11:57, Vlan16
O IA    202.1.1.64/26 [110/40] via 100.1.1.253, 00:11:57, Vlan16
```

Catalyst 3550

```
Switch#show version
Cisco Internetwork Operating System Software
IOS (tm) C3550 Software (C3550-I5Q3L2-M), Version 12.1(19)EA1c, RELEASE SOFTWARE (fc2)
Copyright (c) 1986-2004 by cisco Systems, Inc.
Compiled Tue 03-Feb-04 05:31 by yenanh
Image text-base: 0x00003000, data-base: 0x0080DFF0

ROM: Bootstrap program is C3550 boot loader

Switch uptime is 1 minute
System returned to ROM by power-on
System image file is "flash:c3550-i5q3l2-mz.121-19.EA1c/c3550-i5q3l2-mz.121-19.EA1c.bin"

cisco WS-C3550-24 (PowerPC) processor (revision M0) with 65526K/8192K bytes of memory.
Processor board ID CHK0647W0SF
Last reset from warm-reset
Bridging software.
Running Layer2/3 Switching Image

Ethernet-controller 1 has 12 Fast Ethernet/IEEE 802.3 interfaces

Ethernet-controller 2 has 12 Fast Ethernet/IEEE 802.3 interfaces

Ethernet-controller 3 has 1 Gigabit Ethernet/IEEE 802.3 interface

Ethernet-controller 4 has 1 Gigabit Ethernet/IEEE 802.3 interface

24 FastEthernet/IEEE 802.3 interface(s)
2 Gigabit Ethernet/IEEE 802.3 interface(s)

The password-recovery mechanism is enabled.
384K bytes of flash-simulated non-volatile configuration memory.
Base ethernet MAC Address: 00:11:21:16:31:80
Motherboard assembly number: 73-5700-11
Power supply part number: 34-0966-04
Motherboard serial number: CAT08200G54
Power supply serial number: LIT080707L1
Model revision number: M0
Motherboard revision number: A0
Model number: WS-C3550-24-EMI
System serial number: CHK0647W0SF
Configuration register is 0x10F


Switch#show running-config
Building configuration...

Current configuration : 2716 bytes
!
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
```

```
!
hostname Switch
!
!
ip subnet-zero
ip routing
!
.
.
.
!
interface FastEthernet0/1
 switchport access vlan 300
 switchport mode dynamic desirable
 no ip address
!
interface FastEthernet0/2
 switchport access vlan 301
 switchport mode dynamic desirable
 no ip address
!
.
.
.
!
interface Vlan1
 no ip address
 shutdown
!
interface Vlan300
 ip address 17.17.17.2 255.255.255.0
!
interface Vlan301
 ip address 18.18.18.2 255.255.255.0
!
router eigrp 1
 network 17.17.17.0 0.0.0.255
 network 18.18.18.0 0.0.0.255
 no auto-summary
 no eigrp log-neighbor-changes
!
ip classless
ip http server
!
.
.
.
!
end

Switch#sho ip rou
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

     17.0.0.0/24 is subnetted, 1 subnets
C       17.17.17.0 is directly connected, Vlan300
     100.0.0.0/24 is subnetted, 1 subnets
D EX    100.1.1.0 [170/3072] via 17.17.17.1, 00:12:36, Vlan300
     18.0.0.0/24 is subnetted, 1 subnets
C       18.18.18.0 is directly connected, Vlan301
D EX 201.1.1.0/24 [170/281856] via 17.17.17.1, 00:12:37, Vlan300
     202.1.1.0/24 is variably subnetted, 4 subnets, 3 masks
D EX    202.1.1.128/28 [170/281856] via 17.17.17.1, 00:12:39, Vlan300
D EX    202.1.1.144/30 [170/281856] via 17.17.17.1, 00:12:39, Vlan300
D EX    202.1.1.160/28 [170/281856] via 17.17.17.1, 00:12:39, Vlan300
D EX    202.1.1.64/26 [170/281856] via 17.17.17.1, 00:12:39, Vlan300
```

Switch 8800

```
[8800]display version
3Com Corporation
3COM OS SW8800 V3.01.21s168rec
Copyright (c) 2004-2006 3Com Corporation and its licensors. All Rights Reserved.
Switch 8800 uptime is 2 weeks, 6 days, 6 hours, 34 minutes

[8800]display current-configuration
#
 config-version 3.01.21s168rec
#
 sysname 8800_7slot
#
.
.
.
#
vlan 1
#
vlan 16
#
interface Vlan-interface16
 ip address 100.1.1.252 255.255.255.0
 igmp enable
 pim dm
#
.
.
.
#
interface GigabitEthernet3/1/9
 port link-type hybrid
 port hybrid vlan 16 tagged
 port hybrid pvid vlan 16
#
.
.
.
#
ospf 1
 import-route direct
 import-route static
 area 0.0.0.0
  network 100.1.1.0 0.0.0.255
#
.
.
.
#
return

[8800]display ip routing-table
 Routing Table: public net
Destination/Mask   Protocol Pre  Cost        Nexthop         Interface
17.17.17.0/24      O_ASE    150  20          100.1.1.99      Vlan-interface16
18.18.18.0/24      O_ASE    150  20          100.1.1.99      Vlan-interface16
100.1.1.0/24       DIRECT   0    0           100.1.1.252     Vlan-interface16
100.1.1.252/32     DIRECT   0    0           127.0.0.1       InLoopBack0
201.1.1.0/24       OSPF     10   20          100.1.1.253     Vlan-interface16
202.1.1.64/26      OSPF     10   40          100.1.1.253     Vlan-interface16
202.1.1.128/28     OSPF     10   30          100.1.1.253     Vlan-interface16
202.1.1.144/30     O_ASE    150  1           100.1.1.253     Vlan-interface16
202.1.1.160/28     O_ASE    150  1           100.1.1.253     Vlan-interface16
```

INTEROPERATING WITH CISCO PHONE SYSTEMS

INTRODUCTION

With previous generations of LAN products, it was easy for vendors to claim their Voice over IP (VoIP) solutions only worked with their own LAN infrastructure. This is no longer the case: 3Com is first and foremost an open standards networking vendor and has delivered LAN switches truly capable of operating in a mixed-vendor environment, as this guide shows.

With the latest generation of 3Com LAN switches, customers can deploy the VoIP solution they want and still have a choice of infrastructure equipment to run it on. 3Com switches can detect VoIP phones from any vendor, automatically power them, and segregate and prioritize their traffic to ensure traditional call quality and reliability.

VOICE AND DATA CONVERGED NETWORKING BEST PRACTICES

VoIP technology offers a wide range of benefits, including reduction of telecom costs, management of one network instead of two, simplified provisioning of services to remote locations and the ability to deploy a new generation of converged applications. In pursuit of these benefits, however, every organization must control costs and—even more importantly—risk. No organization can afford to have its voice services compromised.

Thus, those running VoIP must take steps to ensure their converged network delivers acceptable call quality and non-stop availability. In particular, these steps should include:

- Validating the performance and stability of the initial VoIP deployment; and
- Safeguarding that performance and stability as changes are made going forward.

This latter requirement is often overlooked amid concerns about initial deployment; yet ongoing vigilance is essential for successfully maintaining a converged voice and data environment on an ongoing basis, especially in light of the highly dynamic character of today's enterprise networks. Many companies also overlook the potential impact that a real-time application like VoIP can have on their other networked applications and services. This impact should not be underestimated: VoIP can and often does affect the way other critical applications behave.

For voice and data traffic to behave optimally on a shared network, a few best practices should be followed:

- Separate voice and data traffic. The use of a dedicated voice VLAN enhances security and allows both voice and data traffic to be controlled more easily across the network.
- Enforce Quality of Service (QoS) policies. All switches and routers within the environment should participate in the QoS infrastructure, allowing the flow of voice and data traffic to be controlled throughout the network topology.
- Power the IP phone handset. The use of Power over Ethernet provides both power and data to the phone, removing the need for an additional power brick at the desk. Power also can be managed centrally with a UPS or RPS to provide greater availability.

DEPLOYMENT EXAMPLES

1. Configure Legacy PoE Mode for Cisco Phones

Power over Ethernet support was not standardized across the industry until IEEE 802.3af was ratified in July 2003. Cisco PoE devices released to market prior to this date supported a proprietary mechanism for detecting when a device requiring power is connected to a PoE enabled port. As well as being fully 802.3af standards-compliant, all current 3Com PoE switches also are capable of detecting and powering Cisco legacy PoE devices.

The first step is to ensure PoE is enabled on all ports that need to power PoE devices:

```
[5500G-EI] interface GigabitEthernet 1/0/1

[5500G-EI-GigabitEthernet1/0/1] poe enable
```

If detection of Cisco legacy devices is required, this must be enabled globally for each switch:

```
[5500G-EI] poe legacy enable
```

2. Configure Legacy Voice VLAN Mode for Cisco Phones

A major concern for organizations deploying VoIP arises when considering the impact of adds, moves and changes on the network. As organizations add locations, shift users, and make other modifications to the structure of the network and its traffic patterns, end-to-end voice service levels may be affected in unanticipated ways.

3Com has reduced these concerns by introducing its auto voice-VLAN feature, which automatically adds or removes the dedicated voice VLAN from an edge port as a VoIP phone is connected or removed from a switch access port. This feature preserves voice security by only adding the dedicated voice VLAN to ports where VoIP equipment is connected and automatically removing it when that equipment is removed.

Configure Auto Voice-VLAN

1. Define the dedicated voice VLAN:

```
[5500G-EI] vlan 5
```

2. Enable the voice VLAN on the switch:

```
[5500G-EI] voice vlan 5 enable
```

3. Define OUIs for any VoIP equipment that will be connected to the network (if they are not already set in the switch's default configuration):

```
[5500G-EI] voice vlan mac-address 0003-6b00-0000 mask ffff-ff00-0000
description Cisco Phone 1

[5500G-EI] voice vlan mac-address 0015-2b00-0000 mask ffff-ff00-0000
description Cisco Phone 2

[5500G-EI] voice vlan mac-address 0013-1900-0000 mask ffff-ff00-0000
description Cisco Phone 3
```

4. Enable the voice VLAN feature on any edge ports where VoIP devices may be connected:

```
[5500G-EI] interface GigabitEthernet 1/0/1

[5500G-EI-GigabitEthernet1/0/1] port link-type trunk

[5500G-EI-GigabitEthernet1/0/1] voice vlan enable
```

Option 184

When a VoIP device connects to a port with voice VLAN enabled, the administration can be further simplified by using DHCP IP address allocation with Option 184. This DHCP option allows a VoIP device to do the following:

- Receive the IP address it should use on the network to which it has been attached; and
- Receive information about the tagged VLAN it should use for transmitting and receiving.

If this option is configured in the DHCP server to be the same voice VLAN as is configured in the switch, then the whole process of securing the voice traffic on the network can be automated. Most VoIP vendors support DHCP option 184, so this is an excellent open standard way of allowing organizations to consider mixed-vendor networks.

One exception to this practice is Cisco Systems, which instead utilizes its own proprietary network discovery protocol to inform the VoIP devices of the VLAN they should use for voice traffic. 3Com's auto voice-VLAN function includes a Cisco-compatible feature[1] to allow the switch to communicate the VLAN information to a Cisco VoIP device when it detects the discovery protocol request from that device. This enhancement ensures that organizations have the choice of deploying a mixed-vendor solution if that best suits their requirements.

This mode is available on the Switch 5500 series and is configured as follows:

```
[5500G-EI-GigabitEthernet1/0/1] voice vlan legacy
```

3. Configure Automatic QoS for Cisco Phones

Controlling QoS for voice traffic on the network is an essential part of ensuring that calls are not dropped and voice quality remains at acceptable levels. Acceptable VoIP quality requires a bi-directional latency—or delay—of not more than 80 milliseconds (ms) for true toll-quality voice communication. Voice quality degrades as latency increases, but even with a delay of 150-180ms each way, voice quality is still in the acceptable range. If the appropriate QoS policy is applied, voice traffic can pass across the network ahead of less critical data and ensure the VoIP system operates at an acceptable level of reliability and quality.

3Com switches support a number of methods for detecting and automatically prioritizing voice traffic. The traffic can be identified (classified) by several methods, including MAC address OUIs specific to VoIP devices (e.g. 3Com OUI = "00-E0-BB", Cisco OUIs = "00-03-6b", "), Protocol type (e.g., NBX Ethertype = 0x8868), or by traffic marked with the DCSP "EF" code point, as follows:

```
[5500G-EI] acl number 4999

[5500G-EI-acl-ethernetframe-4999] rule 0 permit type 8868 ffff

[5500G-EI-acl-ethernetframe-4999] rule 1 permit source 00e0-bb00-000
ffff-ff00-0000

[5500G-EI-acl-ethernetframe-4999] rule 2 permit source 0003-6b00-000
ffff-ff00-0000

[5500G-EI-acl-ethernetframe-4999] rule 3 permit source 0015-2b00-000
ffff-ff00-0000

[5500G-EI-acl-ethernetframe-4999] rule 4 permit source 0013-1900-000
ffff-ff00-0000

[5500G-EI-acl-ethernetframe-4999] quit

[5500G-EI] acl number 3997

[5500G-EI-acl-adv-3000] rule 0 permit ip dscp ef
```

Once classified, the VoIP traffic needs to be marked to ensure it can be correctly prioritized within the switch and across the rest of the network. This prioritization should be applied to all edge ports that could become connected to VoIP equipment, as shown in the following example:

```
[5500G-EI] interface GigabitEthernet 1/0/1

[5500G-EI-GigabitEthernet1/0/1] traffic-priority inbound ip-group
rule 0 cos voice

[5500G-EI-GigabitEthernet1/0/1] traffic-priority inbound link-group
4999 rule 0 dscp ef cos voice

[5500G-EI-GigabitEthernet1/0/1] traffic-priority inbound link-group
4999 rule 1 dscp ef cos voice

[5500G-EI-GigabitEthernet1/0/1] traffic-priority inbound link-group
4999 rule 2 dscp ef cos voice

[5500G-EI-GigabitEthernet1/0/1] traffic-priority inbound link-group
4999 rule 3 dscp ef cos voice

[5500G-EI-GigabitEthernet1/0/1] traffic-priority inbound link-group
4999 rule 4 dscp ef cos voice
```

1 Requires software version v3.02.02 or higher.

LIST OF FIGURES

GLOSSARY

ABR – Area Border Router – a router used to interconnect the backbone area of an OSPF network to one or more OSPF areas, whose main function is to summarize the routes from the non-backbone OSPF areas to the backbone and to pass information from the backbone area into the areas attached to the ABR.

AS – Autonomous System – a routing domain, many of which make up the Internet. As a collection of networks under a common administration the AS shares a common routing strategy. Each AS is subdivided by areas and must be assigned a unique 16-bit number by the Internet Assigned Numbers Authority (IANA), which delegates authority for IP address-space allocation and domain-name assignment to the InterNIC and other organizations. IANA also maintains the database of AS numbers and assigned protocol identifiers used in the TCP/IP stack.

BPDU – Bridge Protocol Data Unit – a type of packet for communicating spanning tree information between IEEE 802.1D-compatible bridges. There are two types of BPDU packets: configuration BPDUs, used to configure and maintain the spanning tree domain; and Topology Change Notification BPDUs (TCN BPDU), used to indicate a change has occurred in the spanning tree domain.

DHCP – Dynamic Host Configuration Protocol – a protocol from the TCP/IP set of protocols used for the automatic assignment of IP addresses to network devices.

DUAL – Diffusing Update Algorithm – a convergence algorithm used in Enhanced IGRP (EIGRP) that provides loop-free operation at every instant throughout a route computation. It allows routers involved in a topology change to synchronize at the same time, without involving other routers not affected by the change.

EIGRP – Enhanced Interior Gateway Routing Protocol – a Cisco-proprietary distance vector protocol which is totally loop-free and has the fastest convergence speed of all routing protocols; relies on Diffused Update Algorithm (DUAL) to calculate the shortest path to a destination within the network.

IAS – Internet Authentication Service – Microsoft IAS controls who connects to the network using either IEEE 802.1X (Network Login) or MAC address-based authentication (RADA).

IEEE 802 – the Institute of Electrical and Electronic Engineers series of LAN standards

IGP – Interior Gateway Protocol – a routing protocol used within an autonomous system, such as RIP, IS-IS and OSPF. In networks using such exterior gateway protocols as BGP, the routes received have a next hop that is not necessarily connected; the IGP is used to resolve these next hops.

IP – Internet Protocol – the most common protocol used on the Internet, IP is an IETF (Internet Engineering Task Force) standard, first specified in RFC 791. IP protocol functions are equivalent to those of the Network Layer of the OSI Model. It is primarily responsible for directing a data packet to a specific network (subnet) on the Internet. IP also is responsible for formatting the packet for the network on which it is transmitted. The Ethernet protocol's MTU (Maximum Transmission Unit) size of 1518 bytes, for instance, requires a specific Layer 2 format. The IP protocol will break the data stream down into 1518-byte-size packets, each of which is affixed to Layer 2 information. IP is a connectionless protocol – it is not IP's responsibility to verify if the packets transmitted were received properly by the destination device; a higher level protocol or an application will verify a transmission was properly received.

ISL – Inter-Switch Link – Cisco-proprietary protocol that maintains VLAN information as traffic flows between switches and routers.

L3/VPN – Layer 3 Virtual Private Network – Internet Protocol Virtual Private Network (IP VPN) is a group of technologies widely used by corporations and service providers to provide secured, private and scalable communications with proper QoS, over a public IP-based infrastructure, such as the Internet and Service Provider-shared IP networks. IP VPN is replacing the traditional VPN technologies, such as ATM VPN, Frame Relay VPN and TDM-based VPN to become the mainstream of VPN services, though interfaces to the existing technologies exist in some cases. The core technology of VPN is the encapsulation or tunneling algorithms. There are three types of VPNs: MPLS-based Layer 3 VPNs, CPE-based Layer 3 VPNs using the IPSec protocol, and Secure Sockets Layer (SSL) remote-access VPNs (as compared to site-to-site VPNs).

Latency – the delay between the time a device requests access to a network and the time it is granted permission to transmit; and delay between the time a device receives a frame and the time that frame is forwarded out the destination port.

LSA – Link-State Advertisement – a packet of information exchanged among link-state routers about the status and location of networks. The LSAs are stored in each link-state router's LSDB.

LSDB – Link-State Database – a database of LSAs (link-state advertisements) residing on a router configured for link-state routing communications.

MAC – Media Access Control – a sub-layer in the OSI data link layer responsible for accommodating the access methods required to transmit data onto the communication media at the OSI physical layer. IEEE 802.3 CSMA/CD and IEEE 802.5 Token Ring are two of the more common types of media access methods. The MAC layer is also the layer at which MAC addressing occurs.

MISTP – Multiple Instance Spanning Tree Protocol – a Cisco protocol allowing several VLANs to be mapped to a reduced number of spanning-tree instances, made possible because most networks do not need more than a few logical topologies; each instance handles multiple VLANs that have the same Layer 2 topology.

MSTP – Multiple Spanning Tree Protocol – IEEE-recommended protocol (IEEE 802.1s) arising from Cisco's MISTP, relies heavily on RSTP (IEEE 802.1w).

OSPF – Open Shortest Path First – IETF-developed and recommended interior routing protocol uses shortest-path-first algorithm to build a loop-free, shortest-path route calculation, in combination with AS, DR (including backup DR), LSA, LSDB and routing priority to communicate and synchronize.

OUI – Organizational Unique Identifier – three octets assigned by the IEEE in a block of 48-bit LAN addresses.

PoE – Power over Ethernet – the ability for the LAN switching infrastructure to provide power over a copper Ethernet cable to an endpoint powered device (a feature once known as "inline power"). PoE enables scalable and manageable power delivery and simplifies deployments of IP telephones and wireless access points. IEEE 802.3af standardization of PoE is encouraging more new device applications, such as network-attached automation controls, video cameras, point-of-sale devices, and card scanners.

PVST+ – Per VLAN Spanning Tree Plus – an enhancement to IEEE 802.1Q specification that provides support for Dot1q trunking, rather than ISL, to map multiple spanning trees to a single spanning tree; not supported on non-Cisco devices.

QoS – Quality Of Service – a performance measure for a transmission system reflecting its transmission quality and service availability, providing bandwidth ensuring data flow is received within a specified time interval and with minimum errors. QoS is guaranteed on a circuit-based network because a path is dedicated to the data being transmitted between the endpoints. Packet-switched networks however, have a problem with QoS because each of the packets may take a different path to the end device and must contend with packets being transmitted from other devices. As a result, some packets may be delayed and received out of sequence.

RADA – RADIUS Authentication Device Access – a media access control (MAC)-based procedure, used in conjunction with a TippingPoint™ Intrusion Preventions Systems (IPS), which allows a 3Com switch to implement remedial action when suspicious or unauthorized activity is detected

RADIUS – Remote Authentication Dial-In User Service – an AAA (authentication, authorization and accounting) protocol for applications such as network access or IP mobility

RSTP – Rapid Spanning Tree Protocol – an evolution of the Spanning Tree Protocol (802.1D standard) providing for faster spanning tree convergence after a topology change. The standard also includes features equivalent to Cisco's PortFast, UplinkFast and BackboneFast offerings for faster network re-convergence.

STP – Spanning Tree Protocol – implements the 802.1D IEEE algorithm by exchanging BPDU messages with other switches to detect loops, and then removes the loop by shutting down selected bridge interfaces, thereby preventing loops from being formed when switches or bridges are interconnected via multiple paths and guaranteeing only one active path between two network devices.

TCP/IP – Transmission Control Protocol/Internet Protocol – a protocol corresponding to the transport layer and session layers of the OSI model, used to provide connection-oriented services between the process on the transmitting side of a communication link and the destination process on the receiving side of the communication link. TCP provides a reliable connection between communication devices, by providing a mechanism for detecting lost, duplicated, or corrupted packets. On the transmitting side, it is used to break up transmissions into sequenced blocks of data; on the receive side, used to properly sequence and reassemble the transmissions.

UDP – User Datagram Protocol – a Layer 4 TCP/IP protocol providing reliable, connectionless delivery using IP, with the ability to distinguish among multiple destinations within a given host.

UPS/RPS – Uninterrupted Power Supply/Redundant Power Supply

VLAN – Virtual Local Area Network – a Layer 2 protocol enabling the grouping of bridge ports, MAC addresses, or IP subnets (depending on VLAN type) into VLAN domains, which then can be assigned specific network privileges. Benefits include confining a broadcast to a specfic VLAN and providing a base onto which services can be offered or denied to specific VLANs. The IEEE 802.1Q recommendation contains guidelines on VLAN operation.

VoIP – Voice over Internet Protocol – enables a phone conversation to be digitized and transported over an IP network, to reduce long-distance costs and increase universality. Because the signaling needed to map calls to the public network is proprietary, most VoIP installations provide calling only between users on corporate intranets.

Visit www.3com.com for more information about 3Com secure converged network solutions.

3Com Corporation, Corporate Headquarters, 350 Campus Drive, Marlborough, MA 01752-3064
3Com is publicly traded on NASDAQ under the symbol COMS.

503179-001 10/06