

WHITE PAPER

3Com[®]
Solutions:
XRN TECHNOLOGY
AND CLUSTERED
STACKING



3Com Solutions: XRN TECHNOLOGY AND CLUSTERED STACKING

CONTENTS

| | |
|-------------------------------------|----|
| About this Guide | 2 |
| Overview | 2 |
| Configuring XRN Stacking..... | 7 |
| Managing XRN Stacks | 8 |
| Configuring Clustered Stacking..... | 8 |
| Managing Clustered Stacks | 10 |
| Summary..... | 11 |
| List of Terms | 12 |

Since 3Com invented the stackable form factor in the early 1990's, it has become the de facto standard for networking vendors because of its flexibility, scalability, and the ability it provides for companies to cost-effectively upgrade their networks as they grow. Over the years, 3Com has continued to build on its innovation with successive generations of the universally known SuperStack® brand. The most recent evolution, which includes 3Com's unique eXpandable Resilient Networking (XRN®) technology, enables up to eight individual switches to be combined as a single switching and routing entity. It allows network managers to build affordable networks that are highly resilient, deliver exceptional performance and meet today's converged networking demands.

ABOUT THIS GUIDE

This XRN Technology and Clustered Stacking Configuration Guide provides best practices for any technical pre- or post-sales engineer implementing stacking or clustering within 3Com's managed switch portfolio running 3Com OS. It is organized in five main sections:

- › Overview, Including Definitions, Deployment Examples and Product Compatibility
- › Step-by-Step XRN Stacking Configuration
- › XRN Stack Management
- › Step-by-Step Clustered Stacking Configuration
- › Clustered Stack Management

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

OVERVIEW

As organizations become increasingly dependent on converged networks to provide essential business services to office-bound and mobile users, they must ensure that their communications infrastructures effectively support advanced data, voice and multimedia applications. Selecting and deploying the right switching technology is imperative because it can help companies optimize communications while lowering total cost of ownership (TCO).

3Com XRN stacking and clustered stacking, used independently or in tandem, enable companies to centralize the management of their network switches, and provide them with the control and flexibility they need at a price they can afford.

XRN STACKING DEFINED

XRN technology enables multiple interconnected Gigabit switches to behave as a single management switching entity, across Layer 2 and Layer 3, acting together to form a Distributed Fabric. XRN technology logically binds all the functions, performance and resilience together, allowing stackable switches to act like a chassis. And, while traditional solutions offer little more than Distributed Device Management (DDM) for stacking, 3Com XRN technology consists of the following three main technical components, which help companies to minimize their overall resource requirements and reduce associated costs:

- › Distributed Device Management (DDM) – enables the multiple switches in an XRN stack to operate as if they were a single device on the network. DDM lets customers configure and manage all ports and devices in the XRN fabric as a single entity via the command line interface (CLI), web interface, or simple network management protocol (SNMP). In the event of a failure in one of the switches, management access to the remaining switch is retained on the same Internet Protocol (IP) address.
- › Distributed Link Aggregations (DLA) – allows wiring closet switches or hosts to create an aggregated link that is dual-homed across the switches in the XRN fabric. If a port in the aggregated link fails, traffic is forwarded via the remaining ports. DLA provides intelligent local forwarding, enabling each switch in the fabric to make its own forwarding decisions without sending traffic across the aggregated link.

Features such as DLA enable companies to make sure that their networks continue to operate in the event of equipment or cabling failures. DLA also allows managers to schedule maintenance during planned times that are least disruptive to the company's operations, rather than undertaking costly emergency repairs during peak times.

- › Distributed Resilient Routing (DRR) – lets all switches in the XRN fabric act as a single logical router. The switches use the same router interfaces and mirror each other's routing tables. This protects the network against unit failure and enables each switch to locally route traffic for greater Layer 3 forwarding performance.

Using DRR functionality allows managers to configure routing for a stack of switches so that multiple switches can function as a single routing engine. DRR significantly reduces the overall complexity related to configuring separate routers. It also helps to ensure that there are fewer network issues resulting from misconfigurations.

In order to realize the above-mentioned attributes, all members of the XRN stack must be joined together via special stacking-enabled connections.

CLUSTERED STACKING DEFINED

3Com’s clustered stacking delivers benefits similar to the single IP address management available with XRN stacking. But, rather than the advanced level of resiliency and network availability contained in XRN technology, clustered stacking offers increased scalability and flexibility. Clustered stacking simplifies the method by which a group of otherwise individual switches is managed, via a mechanism that is built into the switches themselves. It enables single IP management for up to 32 devices in a single cluster, and allows devices from different product families to be part of the same cluster. Clustered stacking, which does not require the use of additional management applications, is an ideal way to manage a network in logical groups to ensure a consistent management interface, even when the groups contain devices from different product families.

A clustered stack includes the following two types of devices:

- › Cluster Commander - initiates management control for the whole cluster. A cluster commander must be configured with an IP address that is visible from the network management station.
- › Cluster Member – all devices that are not cluster commanders are considered to be cluster members. Members join a cluster at the request of a commander once clustering has been correctly enabled on that device.

COMBINING XRN STACKING AND CLUSTERED STACKING

3Com OS makes it possible for companies to experience the benefits of both XRN stacking and clustered stacking. Network administrators can use XRN technology to create highly-resilient topologies and then group XRN stacks into clusters, managing them as single entities appropriate to the physical configurations of their networks.

The ability to group products from different families is particularly beneficial when technologies, such as 10/100 and gigabit are mixed on parts of the network. Grouping products enables various technologies to be managed as if they were part of the same product family.

DEPLOYMENT EXAMPLES

XRN stacking simplifies management, and is primarily deployed to maximize a network’s availability and performance. In contrast, when deployed by itself, clustered stacking does not have a bearing on a network’s operation, but significantly impacts the ease with which a network can be managed. Different rules should be applied to a network design depending on whether XRN stacking or clustered stacking is in use.

Deploying XRN Stacking

XRN stacking is specifically designed to provide maximum availability for mission critical networks. The ability to split aggregated uplinks across multiple devices in a stack, combined with a distributed routing engine and distributed management, allows for the creation of network topologies where failures in devices or links do not result in noticeable interruptions to the general operation of a network.

By adding the ability to split aggregated XRN stacking uplinks across multiple devices in a resilient core topology, XRN technology provides a network with unparalleled levels of availability – in excess of the 5 “9s” recommended for networks that carry Voice over IP (VoIP) traffic. See Figure 1 for an example of this type of network.

Deploying Clustered Stacking

Clustered stacking allows for multiple devices from various product families to be grouped and managed in a way that best suits the logical topology of the network.

For example, in Figure 2, the engineering researchers on a University campus can take advantage of the bandwidth they need to rapidly access applications from email and PowerPoint presentations for engineering research, simulations and modeling, enabling them to quickly share data and retrieve information on the web. Multiple individual switches spread across several buildings provide students with reliable, high-speed network access from their dormitories.

While this topology has multiple switches in numerous locations and includes various switch families, it is likely that configurations will be consistent for all students. In this scenario, grouping all dormitory switches into a single cluster enables the network administrator to efficiently change the network configuration for all student access rather than requiring each switch to be managed individually.

PRODUCT FAMILY COMPATIBILITY

In any network deployment, there are different options available for implementing XRN stacking versus clustered stacking. The following table demonstrates how the 3Com switch families offer support for XRN stacking and clustered stacking:

| | 4200 | 4200G | 4500 | 4500G | 5500 | 5500G |
|--------------------|------|-------|------|-------|------|-------|
| Clustered Stacking | | ✓ | | ✓ | ✓ | ✓ |
| XRN Stacking | ✓* | | ✓* | | ✓ | ✓ |

*Switch Family does not support DLA or DRR

FIGURE 1: XRN Stacking Deployed in a Mission-Critical Environment

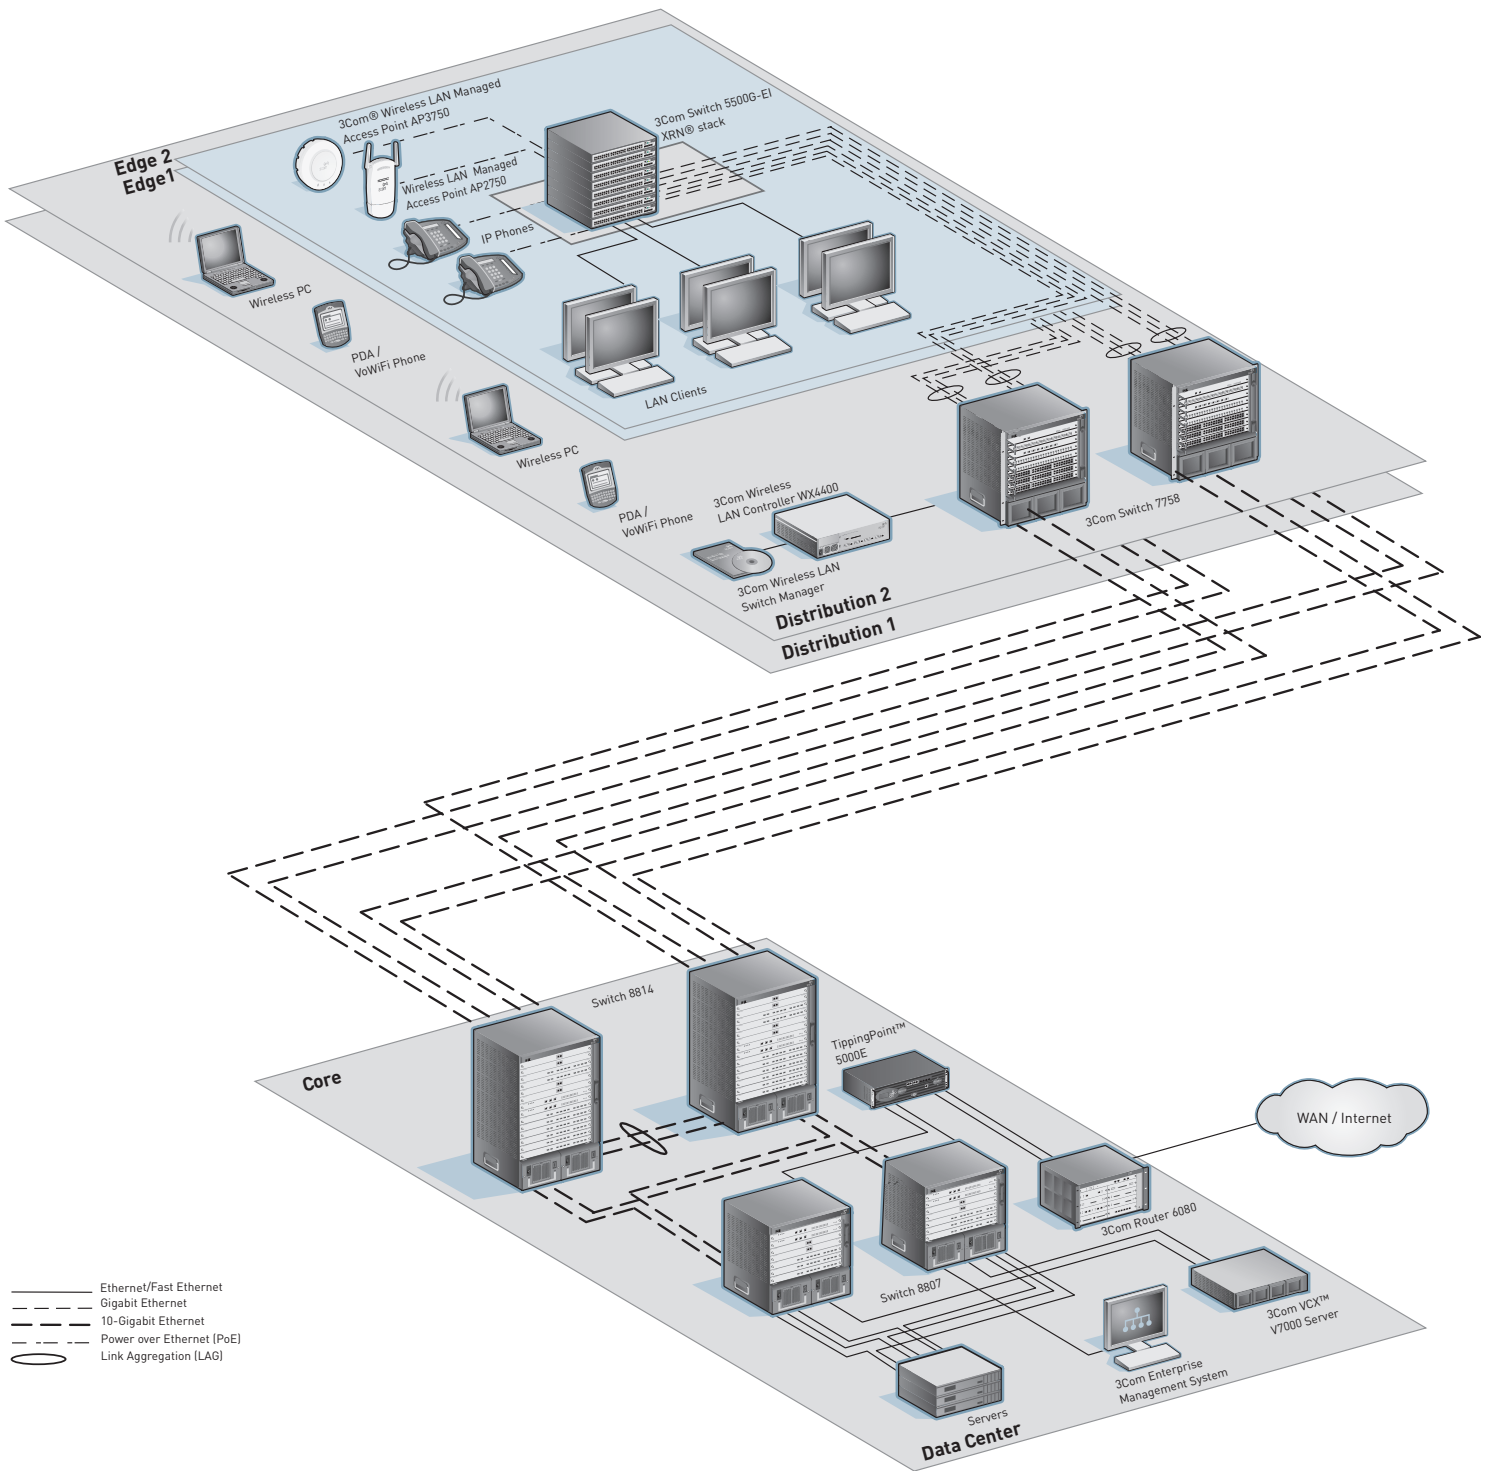
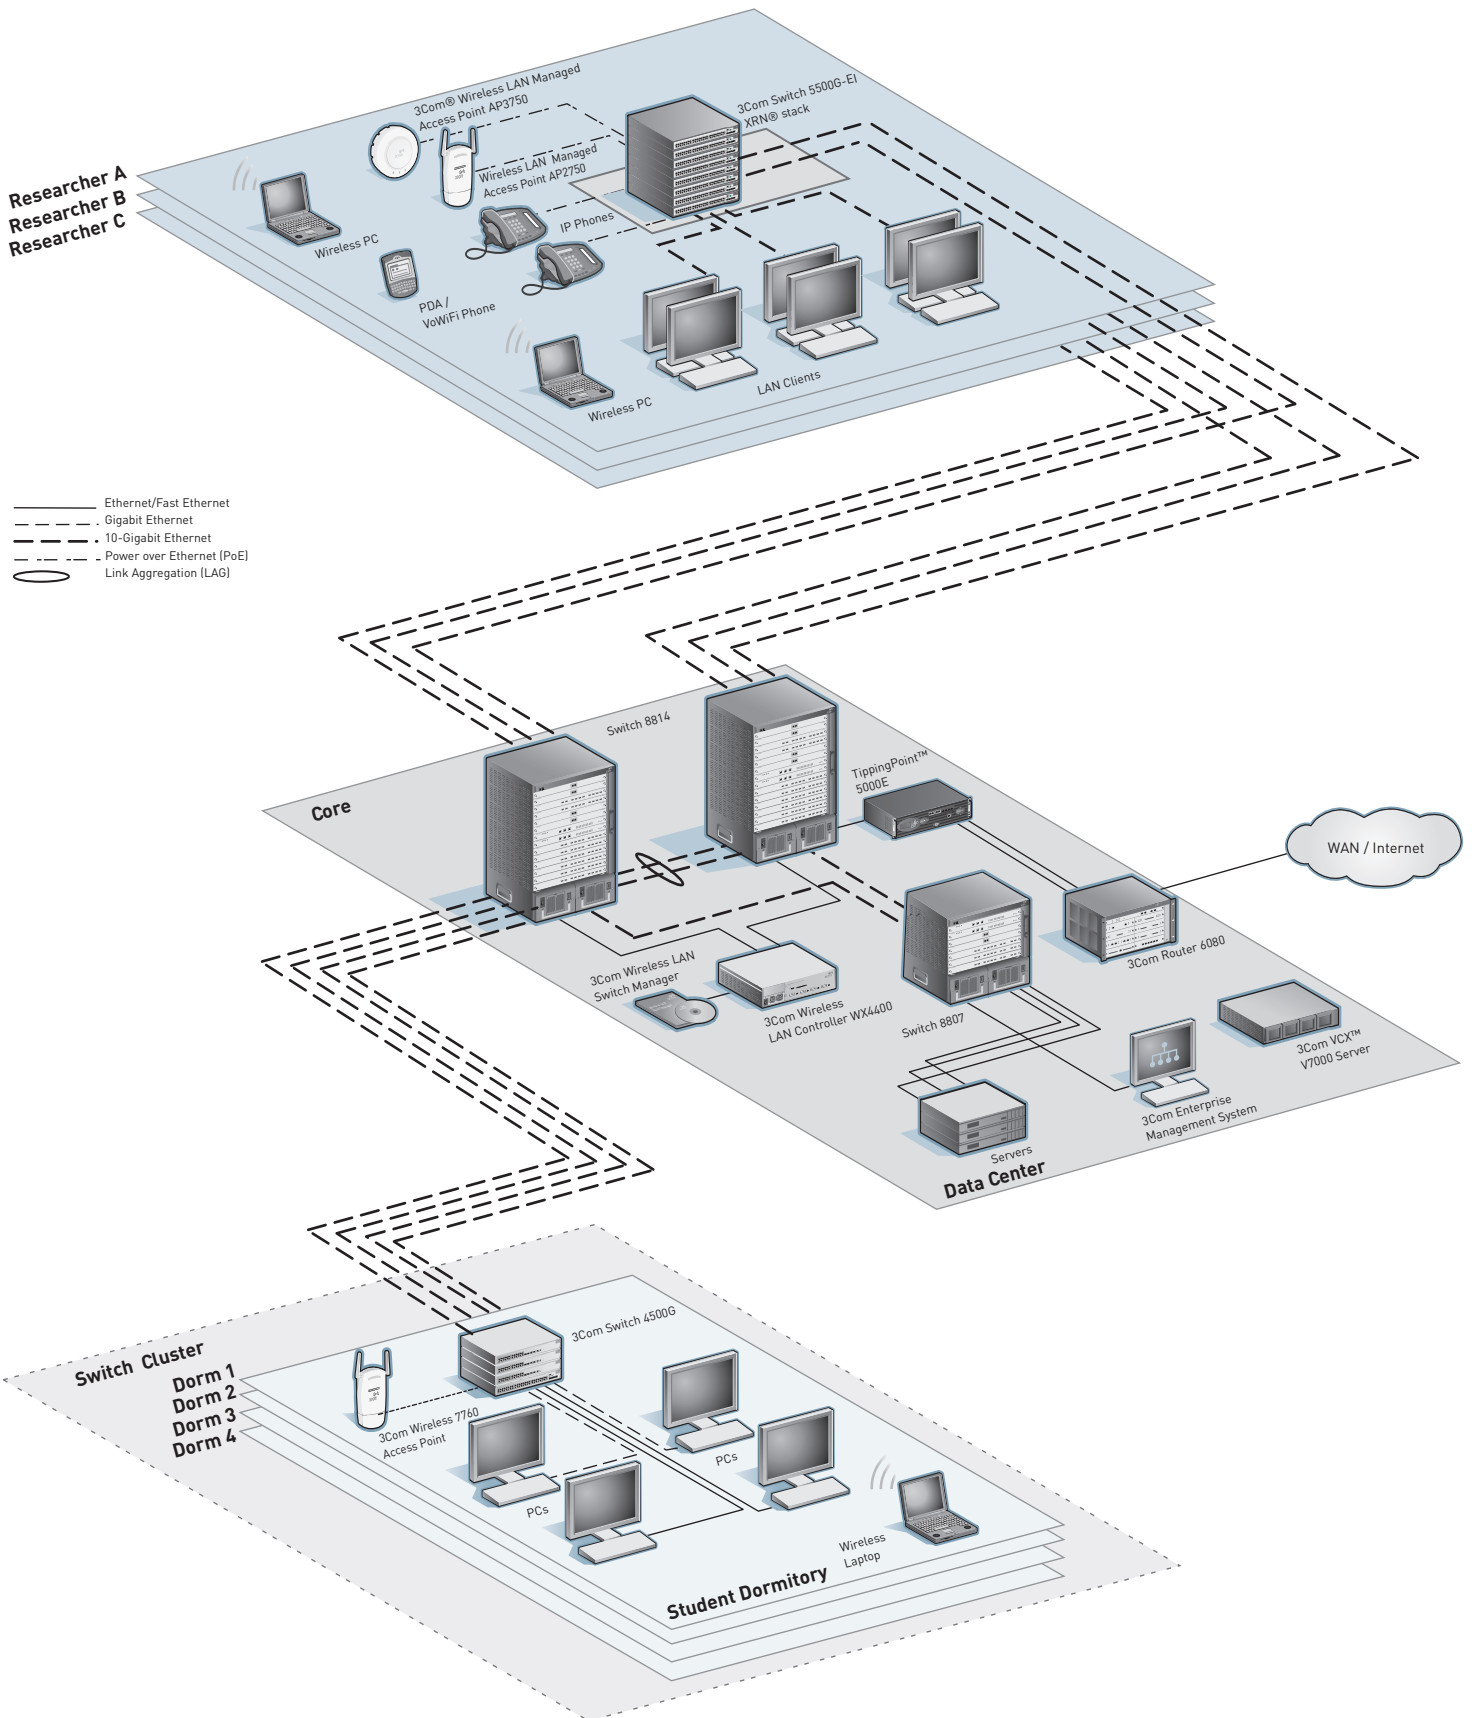


FIGURE 2: University with XRN and Clustered Stacking



CONFIGURING XRN STACKING

XRN technology can be used throughout an enterprise network to provide high-performance, highly fault-tolerant backbones, Gigabit-over-copper aggregation, and Layer 3 switching. After configuration, all switches actively share routing intelligence and network loads—eliminating the wasted bandwidth and added expense of a passive standby unit. Ultra-fast failover recovery automatically redistributes traffic among the other active units if a switch becomes disconnected.

There are two basic ways to implement XRN stacking:

- **Local XRN stacking connection.** Switches can be configured locally in a resilient XRN stacking configuration, with both switches typically residing in the same rack.
- **Distributed XRN stacking connection.** Switches can be remotely connected up to 70Km apart using a fiber connection.

Regardless of whether local or distributed XRN stacking connections are used, 3Com switches that support XRN stacking must adhere to the following four main rules:

1. All devices must be from the same product family.

Only models from the same product family can be stacked together. The 3Com Switch 5500-EI series and the 3Com Switch 5500G-EI series are the two 3Com families that currently support full XRN stacking.

2. All devices must have the same installed versions of software.

Before deploying a stack, all devices should be checked to ensure they are running the same version of software. If they are not the same, a common software version should be chosen, copied to all units and set as the next boot code before the units are deployed in the stack.

New software files should be uploaded to switch flash memory and then configured as the next boot versions. All specific configuration files that are loaded at this time should ensure that the same system name is used across all units, and that there are no conflicting unit numbering assignments.

3. All devices must have the same system name.

All units destined for each particular stack and loaded with the same version of software, must have the same system name. This will be the factory default, or it can be set manually via the CLI.

```
[5500-EI] sysname <system_name>
```

4. All devices must have interconnected ports enabled as fabric ports.

Depending on the switch, different methods are used to physically stack the units together; this affects how the fabric ports are enabled.

Switch 5500 10/100 variants are connected using standard front panel Gigabit ports. To ensure that all units can successfully join a stack, the user must set interconnecting ports to their factory default setting, therefore enabling them as fabric ports.

```
[5500-EI] fabric-port GigabitEthernet 1/0/27 enable
[5500-EI] fabric-port GigabitEthernet 1/0/28 enable
```

[Note: on a 52-Port Switch 5500, ports 27 and 28 will be 51 and 52.]

Switch 5500G Gigabit variants are connected via dual dedicated 48 Gbps high-bandwidth stacking ports on the rear panel of the units.

For both 10/100 and Gigabit models, fully-resilient stacking is enabled when the units are daisy chained together and connected in a full ring by adding a stacking connection from the top unit in the stack down to the bottom unit in the stack.

Switch 4500 10/100 models can be stacked using the same method as the Switch 5500 10/100 models, except without supporting the resilient stack connection. This allows the two free Gigabit connections at the top and bottom of the stack to be used as resilient uplinks.

NUMBERING DEVICES FOR STACKING

When stacking switches together, users should decide whether they are going to set the unit numbering to ensure that the same units always have the same number within the stack, or if, for ease of deployment, they are going to allow units to automatically number themselves.

Unit numbering can be controlled on individual units or on units that have already joined a stack.

Return a unit to its factory default of auto-numbering with this:

```
[5500-EI] change self-unit to auto-numbering
```

If manual unit numbering is required, use the following to ensure that numbering within a stack remains consistent when units are added to or removed from a stack:

```
[5500-EI] change self-unit to <unit_number>
```

Use the following if other units in a stack are to be re-numbered:

```
[5500-EI] change unit-id <current_unit_number> to <new_unit_number>
```

Once numbering changes for a unit or units in a stack is complete, the changes should be saved. Use the following to preserve the changes across a reboot:

```
<5500-EI> fabric save-unit-id
```

STACKING STATUS

Once a stack has been successfully created and all units have joined the stack, use the following to check on the status of the stacking fabric from the CLI:

```
[5500-EI] display xrn-fabric
Fabric name is 5500-EI, system mode is L3.
```

| Unit Name | Unit ID |
|-----------|---------|
| 5500-EI | 1 |
| 5500-EI | 2 |
| 5500-EI | 3{(*)} |

After the stack has been built, the complete stack can be managed from the IP address of any single unit within that stack.

MANAGING XRN STACKS

XRN stacking enables multiple physical devices from the same product family to operate as a single logical entity that can be managed from a single IP address via the CLI, web interface or SNMP. In the event of a failure in one of the switches, management access to the remaining switch is retained on the same IP address.

3COM OS CLI

Features such as routing, multicast support and spanning tree, which are global to a stack of switches, only need to be enabled once in the 3Com OS CLI for the entire stack. All commands in the 3Com OS CLI relating to port specific configuration are structured so that they can easily scale when an individual unit is connected to an XRN stack of devices. All such commands use the xx/yy/zz port numbering format, where xx is the unit number within the stack, yy is the slot number for modular products (always “00” for stackable switches) and zz is the port number within unit xx. The xx/yy/zz format enables all ports in a stack to be configured without having to connect to multiple units. When connected to one device in a stack, it is possible to control features and apply attributes to ports on other units, purely by changing the value of xx in the command line. For example, to display the status of port 10 for unit 3 in a stack the following command is used:

```
[5500G-EI] display interface ethernet 3/0/10
```

SNMP

The single logical entity that is created when multiple switches are linked to create a stack enables that stack to be managed from any standard SNMP management software as a single entity.

For example, when viewing a stack of Switch 5500G-EIs on a 3Com Network Director topology, each stack is shown as a single device, rather than a group of connected devices as is the case with a cluster.

CONFIGURING CLUSTERED STACKING

Clustered stacking technology from 3Com enables effective centralized management of 4500G switches, and also clusters that include switches from the 4200 and 5500 families. Deploying clustering is a relatively simple process that involves the following steps:

- › Divide the switches into logical groups.
- › Select a commander for each cluster.
- › Configure the cluster management VLAN.
- › Define the scope of the cluster.
- › Enable clustering on the commander.
- › Build the cluster.
- › Enable clustering on the other switches in the cluster.
- › Re-build the cluster.

DIVIDE THE SWITCHES INTO LOGICAL GROUPS.

All of the elements related to the division of switches into logical groups are beyond the scope of this paper. However, the below items need to be considered:

1. All devices must be from 3Com Switch Families that support clustering.

Models from the various product families that support clustering can be stacked together. The 3Com Switch 4500G series, 5500 series and 5500G series all support clustered stacking for up to 32 devices; the Switch 4200G series supports up to 16 devices.

2. All switch devices must have direct access to the same Layer 2 broadcast domain.

All units intended to be part of a particular cluster must have direct access to the same Layer 2 broadcast domain. Providing this access ensures that devices can communicate directly with one another at Level 2, and helps companies avoid security issues that can arise when trusted and untrusted devices share the same broadcast domain.

3. Ports connecting clustered devices must be members of the same “cluster management VLAN.”

Ports connecting devices that intend to be part of the same cluster must be members of the same “cluster management VLAN.”

4. Each cluster must be configured with its own unique “cluster management VLAN.”

Multiple clusters can exist on the same broadcast domain, but each cluster must be configured with its own unique “cluster management VLAN.”

SELECT A COMMANDER FOR EACH CLUSTER.

There are no designated rules for selecting a commander, but certain factors can be considered, such as choosing the aggregation switch for a wiring closet.

CONFIGURE THE CLUSTER MANAGEMENT VLAN.

By default, 3Com OS uses VLAN 1 as the cluster management VLAN. If an alternative VLAN must be used, it should be configured before clustering is enabled. An alternate VLAN could be needed due to security or other reasons, such as the deployment of multiple clusters in the same broadcast domain.

If an alternate VLAN is required, it must first be created. For example:

```
[5500G-EI] vlan 2
[5500G-EI-vlan2] description Cluster_Management_VLAN
```

Next, an IP address must be added to the interface of the chosen management VLAN so that the cluster can be accessed across the network:

```
[5500G-EI-vlan2] interface vlan 2
[5500G-EI-Vlan-interface2] ip address 192.168.1.1 24
[5500G-EI-Vlan-interface2] quit
```

Note: If there is a requirement to manage all devices in a cluster from a centralized SNMP management tool, then the ports that link the devices in that cluster must be set to the hybrid type.

The new VLAN must be specified as the management VLAN of the switch:

```
[5500G-EI] management-vlan 2
```

The new VLAN must be specified as the VLAN interface through which the cluster management will be performed:

```
[5500G-EI] cluster
[5500G-EI-cluster] nm-interface vlan-interface 2
[5500G-EI] quit
```

DEFINE THE SCOPE OF THE CLUSTER.

The scope of the cluster should be considered before the cluster is built. The scope can be constrained by defining how many devices the commander will look across to join to the cluster. This is specified in terms of hop count:

```
[5500G-EI] ntdp hop 2
```


ENABLE CLUSTERING ON THE COMMANDER.

Once the management VLAN has been chosen and configured and the scope of the cluster has been defined, clustering can be enabled on the designated commander:

```
[5500G-EI] cluster enable
```

Then, assign the private IP address range that the commander will use to communicate with all devices in the cluster:

```
[5500G-EI] cluster
[5500G-EI-cluster] ip-pool 192.168.1.1 24
```

Note: although the pool of addresses can be re-used, it is recommended that each cluster in a network is given a unique pool of IP addresses to ensure no address conflicts occur when the clusters are managed via SNMP.

BUILD THE CLUSTER.

The commander can now be instructed to build the cluster:

```
[5500G-EI-cluster] auto-build
Restore topology from local flash file, for there is no base topology.
(Please confirm in 30 seconds, default No). (Y/N) N
```

The commander will prompt the user to enter an identifying name for the cluster:

```
Please input cluster name: GR_1
```

If there are no connected devices that support clustering, or clustering has not been enabled in devices that are connected, the following result will be returned:

```
Collecting candidate list, please wait...
Candidate list:
Name Hops  MAC Address  Device
Processing...please wait
Cluster auto-build Finish!
      0 member(s) added successfully.
```

At this point, if the commander is asked to display a list of all members of the cluster, it will show itself as the only member:

```
[GR_1_0.5500G-EI-cluster] display cluster member
SN Device          MAC Address      Status  Name
0 3Com Switch 5500G-EI 48-Port 0016-e0d6-7d00 Admin GR_1_0.5500G-EI
```

Notice that an additional “0” is added to the cluster name. It denotes the device number within the cluster, with “0” always being reserved for the commander.

ENABLE CLUSTERING ON THE OTHER SWITCHES IN THE CLUSTER.

Because clustering is disabled on the switches by default, the other members of the logical group now have to be enabled for clustering:

Member 1:

```
[5500-EI] cluster enable
```

Member 2:

```
[4200G] cluster enable
```

When the commander is again asked to list all members of the cluster, it will show that these units have joined the cluster:

```
[GR_1_0.5500G-EI] display cluster member
SN Device          MAC Address      Status  Name
0 3Com Switch 5500G-EI 48-Port 0016-e0d6-7d00 Admin GR_1_0.5500G-EI
1 3Com Switch 5500-EI          0012-a992-9d01 Up      GR_1_1.5500-EI
2 3Com Switch 4200G 48-Port 0016-e0f7-c480 Up      GR_1_2.4200G
```

In certain deployments, it may be easier to globally enable all devices for clustering and then exclude some devices from being added to the cluster. In these cases, it is possible to use the “black-list” feature to exclude specific devices by MAC address:

```
[GR_1_0.5500G-EI] cluster
[GR_1_0.5500G-EI-cluster] black-list add-mac 0012-a992-9d01
      Inserted the MAC to the black-list successfully!
[GR_1_0.5500G-EI-cluster]
      Member 0012-a992-9d01 in cluster is deleted.
```

Displaying the member list confirms that the unwanted device has been removed from the cluster:

```
[GR_1_0.5500G-EI-cluster] display cluster member
SN Device          MAC Address      Status  Name
0 3Com Switch 5500G-EI 48-Port 0016-e0d6-7d00 Admin GR_1_0.5500G-EI
2 3Com Switch 4200G 48-Port 0016-e0f7-c480 Up      GR_1_2.4200G
```

Devices explicitly excluded from a cluster can be confirmed from any device in that cluster by viewing the contents of the black-list:

```
[GR_1_0.5500G-EI-cluster] display cluster black-list
Device ID          Access Device ID  Access port
0012-a992-9d01     0016-e0d6-7d00   GigabitEthernet1/0/14
```

If the network topology changes and a device previously excluded from a cluster is removed from the network or subsequently required to join that cluster, the device can be removed from the black-list:

```
[GR_1_0.5500G-EI-cluster] black-list delete-mac 0012-a992-9d01
      Deleted MAC from the black-list successfully!
[GR_1_0.5500G-EI-cluster] display cluster black-list
      The black-list is empty now!
```

RE-BUILD THE CLUSTER.

Once clustering has been successfully enabled on the other switches, the commander can be instructed to rebuild the cluster with all available members:

```
[GR_1_0.5500G-EI-cluster] auto-build
Collecting candidate list, please wait...
Candidate list:
Name  Hops  MAC Address      Device
5500G-EI  1     0012-a992-9d01   3Com Switch 5500G-EI
Processing...please wait
Member 0012-a992-9d01 is joined in cluster GR_1.
      Cluster auto-build Finish!
      1 member(s) added successfully.
```

Displaying the member list confirms that the previously excluded device has joined the cluster:

```
[GR_1_0.5500G-EI-cluster] display cluster member
SN Device          MAC Address      Status  Name
0 3Com Switch 5500G-EI 48-Port 0016-e0d6-7d00 Admin GR_1_0.5500G-EI
1 3Com Switch 5500-EI          0012-a992-9d01 Up      GR_1_1.5500-EI
2 3Com Switch 4200G 48-Port 0016-e0f7-c480 Up      GR_1_2.4200G
[GR_1_0.5500G-EI-cluster]
```

After the cluster is configured and the correct members are included and excluded, the configuration is complete and the cluster can be managed from the single IP address that was assigned to the cluster commander at the beginning of the cluster configuration procedure.

MANAGING CLUSTERED STACKS

Clustered stacking enables multiple devices from different product families to be managed as a single logical entity either via their CLI, web interface or SNMP.

3COM OS CLI

Once a cluster has been created, all devices in the cluster can be managed by connecting to a single unit. Clusters are implemented so that all members of a cluster become an extension of the CLI of the cluster commander. That is, once a CLI connection has been established with the commander, it is possible to manage all other members of the cluster without having to connect directly to them.

A combination of display cluster member and switch-to cluster member commands can be used to identify and then manage the members of the cluster:

```
[GR_1_0.5500G-EI] display cluster member
```

| SN | Device | MAC Address | Status | Name |
|----|------------------------------|----------------|--------|-----------------|
| 0 | 3Com Switch 5500G-EI 48-Port | 0016-e0d6-7d00 | Admin | GR_1_0.5500G-EI |
| 1 | 3Com Switch 5500-EI | 0012-a992-9d01 | Up | GR_1_1.5500-EI |
| 2 | 3Com Switch 4200G 48-Port | 0016-e0f7-c480 | Up | GR_1_2.4200G |

Once the switch has been identified, it is simple to view and manage that device using its member number in the cluster. For example, to view the Switch 5500-EI in this cluster, or member "1.":

```
[GR_1_0.5500G-EI] cluster switch-to 1
```

Once the command is entered, the command line of the target switch is accessed and that device can be configured as normal:

```
[GR_1_1.5500-EI]
```

After the member device configuration changes have been completed, control can be returned to the commander by quitting from the member device's command line:

```
[GR_1_1.5500-EI] quit
```

```
[GR_1_0.5500G-EI]
```

Using this method, all members of the cluster can be managed and configured from a single management session.

WEB INTERFACE

It is possible to carry out cluster management via a standard web browser for all 3Com product families that support cluster management from the web interface.

The status of the cluster can be monitored and reconfigured from a few simple to use browser screens. Figure 3 shows the summary screen.

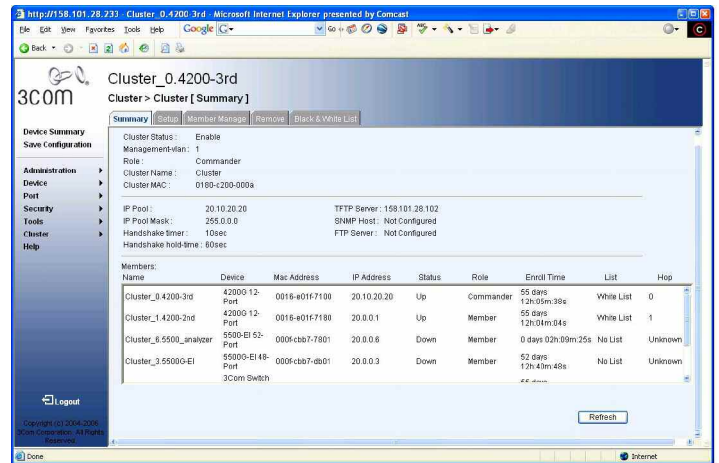


FIGURE 3: Cluster management summary screen

The web interface can also be used to display, as shown in Figure 4, the topology of the network that contains the cluster.

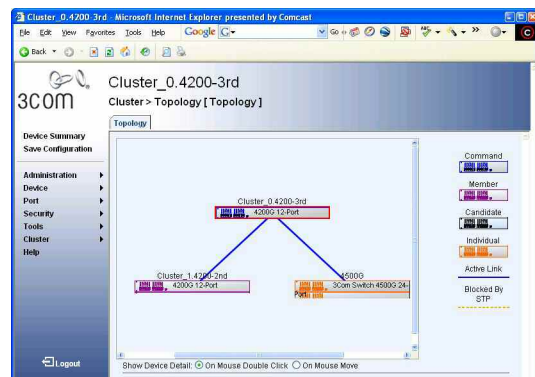


FIGURE 4: Cluster topology

When units are grouped and managed as a cluster via a web interface, the file system can be easily managed on those devices.

Software or unit configurations can be updated across the cluster, as shown in Figures 5 and 6.

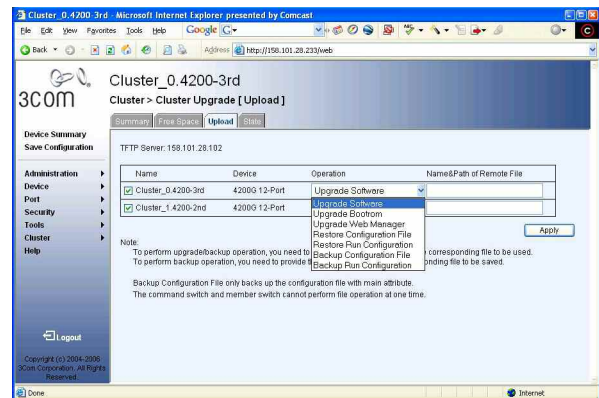


FIGURE 5: Cluster software upgrade

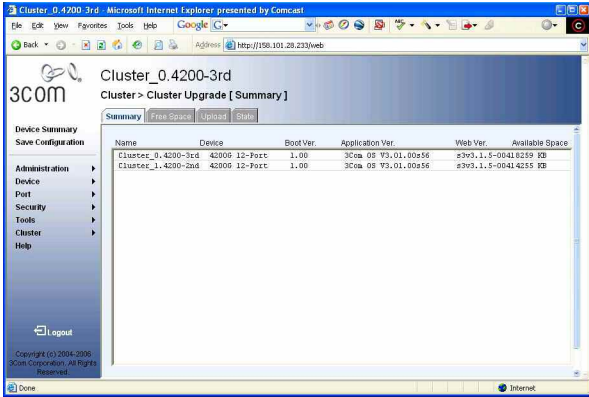


FIGURE 6: Cluster device upgrade

SNMP

Once a cluster has been created, it can be identified and managed via any common SNMP management tool if the management IP interface has been set-up on all cluster members.

Figure 7 shows a 3Com Network Director discovered topology map for a network consisting of two separate clusters and the aggregating switch that links them. To simplify the example, other interfaces on the network have been left as undiscovered.

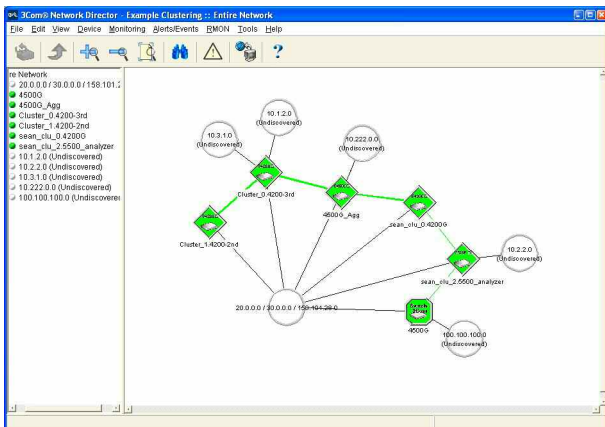


FIGURE 7: 3Com Network Director discovered topology map

The cluster name and member number are appended to the system name when a unit joins a cluster. Therefore, it is easy to use the topology map to see which units are clustered.

For remote SNMP management with the added benefit of configuring units via a graphical interface similar to the web interface, it is possible to use the 3Com Device Manager, either as part of Network Director or standalone, to view and manage a cluster. Device Manager combines a graphical view of the cluster with graphical representations of the front panels of the units in that cluster, simplifying the overall configuration procedure for port-based features. See Figure 8.

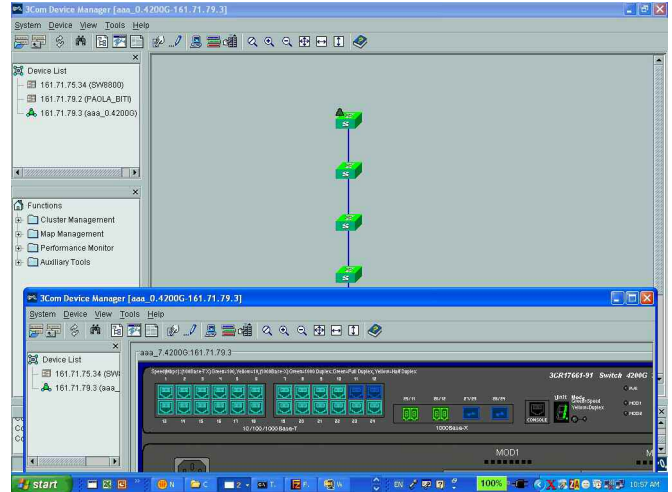


FIGURE 8: 3Com Device Manager

Using the Device Manager allows groups of ports to be selected and enables changes to be applied to all ports in that group with a single management action.

SUMMARY

Companies can use 3Com XRN technology to build affordable networks that are highly resilient and deliver exceptional availability. XRN technology, which includes DDM, DLA, and DRR, enables up to eight switches to be combined as a single switching and routing entity, and allows companies to design and implement core Gigabit backbones that feature no single point of failure – a safeguard against hardware, cabling and software failure. The distributed nature of XRN technology takes full advantage of the high-performance, non-blocking switching capacity of each switch in the Distributed Fabric.

On the other hand, clustered stacking provides companies with increased scalability and flexibility. It allows single IP management for up to 32 devices in a single cluster, and enables various product devices from different product families to be part of the same cluster. Because clustered stacking does not require the use of additional management applications, it is an ideal way to manage a network in logical groups to ensure a consistent management interface.

3Com XRN technology and clustered stacking can be deployed separately, or implemented collectively to provide companies with a powerful, cost-effective solution that centralizes the management of network switches and delivers an unprecedented amount of control and flexibility.

LIST OF TERMS

3Com OS – 3Com Operating System

CLI – Command Line Interface

DDM – Distributed Device Management

DLA – Distributed Link Aggregations

DRR – Distributed Resilient Routing

IP – Internet Protocol

MAC – Media Access Control

SNMP – Simple Network Management Protocol – The IETF protocol is used between network management stations and manageable network devices. It is used to read (GET) or write (SET) information in an NIC, repeater, switch, router, server, workstation or other network device. Using SNMP, a management station can configure devices, display statistics, detect error rates, identify network changes and gather status information.

TFTP/FTP – Trivial File Transfer Protocol/File Transfer Protocol

VLAN – Virtual Local Area Network – 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 specific 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.

XRN Technology - eXpandable Resilient Networking



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.

Copyright © 2007 3Com Corporation. All rights reserved. 3Com, the 3Com logo and XRN are registered trademarks of 3Com Corporation. All other company and product names may be trademarks of their respective companies. While every effort is made to ensure the information given is accurate, 3Com does not accept liability for any errors or mistakes which may arise. All specifications are subject to change without notice. 503183-001-001 02/07