

Objective of This Document	1
Interoperability within Cisco Environment.....	1
Tolly Group Interoperability Certification.....	1
About BLADE Network Technologies.....	1
Introduction	2
1. Physical Connection	3
2. Spanning-Tree	3
3. VLANs.....	8
4. Link Aggregation.....	9
5. Routing Protocols	10
6. Management and Administration:.....	11
7. Additional Information.....	13
Independent Interoperability Certification by the Tolly Group	13
For More Information.....	14

Objective of This Document

Interoperability within Cisco Environment

This document highlights the integration process of BLADE Network Technologies' Gigabit Ethernet Switch Modules for IBM BladeCenter in a Cisco environment. Guidelines and recommendations are proposed to assist the network administrator in his or her implementation. The goal of this document is to give the reader a basic and practical understanding on how to integrate BLADE switching technology within the BladeCenter into a Cisco infrastructure, and to point out where to go to gain more detailed information.

Tolly Group Interoperability Certification

This document also covers results of recent certification by Tolly Group for the BLADE Network Technologies Layer 2-3 and L2-7 GbESM switches for interoperability with Cisco and other manufacturer's products.

About BLADE Network Technologies

On February 13, 2006, Nortel and Garnett & Helfrich Capital announced that the two companies have entered into an agreement to establish Nortel's Blade Server Switch Business Unit ("BSSBU") as a private enterprise—BLADE Network Technologies, Inc.—the industry's first vendor to focus exclusively on serving the network infrastructure requirements of the rapidly growing blade server market.

BLADE Network Technologies retains the intellectual property and Nortel branding for the Gigabit Ethernet Switch Modules for IBM BladeCenter. While the term "Nortel" may be used as a brand throughout this document, the switches mentioned are designed and manufactured by BLADE Network Technologies. When comparing commands between the BLADE Network Technologies GbESM switch operating system and the Cisco IO operating system, the terms "AOS" and "Cisco" will be used in this document.

Introduction

The GbESM switch operating system, Alteon OS (AOS), has been tested in many environments with many different network vendors, including Cisco. **To date, with over 150,000 units in production, no inter-operability incidents have been reported.** The Tolly Group and BLADE Network Technologies have tested and certified the GbESM to work with most of Cisco's equipment. The majority of the protocols used by BLADE are based on industry standards and therefore, are compatible with other vendors adhering to these standards. In addition, Nortel also includes support for several Cisco proprietary protocols.

Confusion can sometimes arise due to the giving of different names to the same features. The table below shows that many of the BLADE Network Technologies and Cisco terms often refer to the same feature, which adheres to the same standard.

Different terminology, same functionality

Nortel GbESM from BLADE Network Technologies	Cisco
Port Fast Forwarding	Portfast
Spanning Tree Protocol Group (STG)	Per VLAN Spanning Tree (PVST), PVST+
Fast Uplink Convergence	UplinkFast
Link Aggregation, Multi-Link Trunking (MLT), Trunking	EtherChannel, Channeling
VLAN Tagging, 802.1Q Tagging	Trunking, .1q (dot one q) Tagging, Intra Switch Link (ISL)
Port VLAN Identification (PVID)	Switchport Access VLAN (untrunked port) Native VLAN (trunked port)
Virtual Router Redundancy Protocol (VRRP)	Hot Standby Router Protocol (HSRP)
Forwarding Database (FDB)	Content-Addressable Memory (CAM)
Port Mirroring	Switched Port Analyzer (SPAN)

The GbESM switch has two command line interface (CLI) entry modes—AOS and ISCLI. Alteon OS is a Unix-type menu based command system, while ISCLI has a command structure more familiar to users of Cisco switches.

1. To access the ISCLI, enter the following command from the AOS CLI:

```
/boot/mode ISCLI
```

2. Then reboot the switch:

```
/boot/reset
```

The default command-line interface for the GbESM switch is the AOS CLI.

3. To revert to the AOS CLI from the ISCLI, enter the following command:

```
>> ISCLI# boot cli-mode aos
```

4. Then reboot the switch:

```
>> ISCLI#reload
```

A browser-based interface (BBI) using a standard web browser is also support for configuration of the GbESM switches.

The remainder of this document is divided into the following sections:

5. Physical Connection
6. Spanning Tree
7. VLANs
8. Link Aggregation
9. Routing Protocols
10. Management and Administration
11. Additional Information

1. Physical Connection

The IBM BladeCenter infrastructure allows the initial configuration of the GbESM via the Management Module. It is recommended to proceed with the initial configuration before connecting the external ports to your current infrastructure. It is also recommended to connect only one port at the time to avoid potential configuration errors and disruption of your network.

All internal interfaces on the 1Gb Ethernet switch modules are default configured to 1000 Mbps Full-Duplex. Auto-negotiation may be enabled on these internal ports, if required. (Auto-negotiation may be required for blade servers running Solaris or for QS20 blades.)

All external interfaces are auto-sensing on both the L2-3 Copper and L2-7 GbESM. If auto-negotiation is enabled, a straight CAT5 cable can be used. If auto-negotiation is turned off, a cross over cable is required to connect to another switch (a straight-through cable is required in this case to connect a free-standing server to an external port). The L2-3 Layer 2-3 GbESM with fiber interfaces requires 1000BaseSX Multi-Mode cables and LC connectors; these fiber connectors run only at 1 Gb, and will not auto-negotiate down to 10/100 Mb.

For the 10GbESM switch, all 10Gb internal and external ports are hard configured to 10 Gb full-duplex.

Typically, auto-negotiation between Nortel and Cisco will be successful. To validate the port speed use the following commands:

Cisco: **show interface**

AOS **/info/link**

If unsuccessful, you can also configure the ports manually by disabling auto-negotiation:

AOS **/cfg/port Ext1/gig**

Cisco: **switch#(config-if): speed 1000**
switch#(config-if): duplex full

In this last case, a cross connect cable will be required.

2. Spanning-Tree

Both Cisco and Nortel's Spanning-Tree implementation are both are fully compatible with respect to IEEE 802.1d, 802.1q (VLAN tagging), 802.1w (Rapid STP) and 802.1s (Multiple STP) standards. Cisco's RPVST protocol is not industry standard, and represents a proprietary

implementation of STP which is not compatible with non-Cisco products, including the Nortel BladeCenter switches. However, interoperability can still be attained in all scenarios.

Cisco uses two approaches to multiple instances of Spanning Tree: RPVST or MSTP. RPVST is Cisco-proprietary and is not supported by other networking vendors, including the Nortel BladeCenter switches. However RPVST can interoperate with PVST.

In general, the following options are available for STP interoperability with Cisco uplink switches:

If the Cisco is running (see below), then you can run this on the L2-3 GbESM

=====	=====
RPVST	PVST
PVST	PVST
STP	STP
MSTP	MSTP
RSTP	RSTP

Nortel GbESMs support up to 128 instances of Spanning-Tree per switch. You can configure one instance to only have one VLAN (equivalent to Cisco PVST+) or to have multiple VLAN sharing the same Spanning-Tree instance (equivalent to Cisco MSTP). The Spanning Tree functions are configured on the Nortel GbESM with commands from the `/cfg/l2/stg` and `/cfg/l2/mrst` menus.

RPVST

If the Cisco core is running the proprietary RPVST version of spanning-tree protocol, this will intercommunicate with PVST BPDUs:

- 1) Run PVST on the Nortel side, and assign each VLAN to its own spanning-tree-group.

Nortel:

For multiple VLANs, assign each one to a dedicated spanning tree group.

```
/cfg/l2/vlan 1/stg 1
/cfg/l2/vlan 302/stg 2
/cfg/l2/vlan 304/stg 3
/cfg/l2/vlan 4095/stg 128
```

Increase the bridge priority to the maximum value as a precaution; the location of the IBM BladeCenter switch on the edge of the network makes it poor choice for the STP root bridge:

```
/cfg/l2/stg 1/brg/prio 65535
/cfg/l2/stg 2/brg/prio 65535
/cfg/l2/stg 3/brg/prio 65535
/cfg/l2/stg 128/brg/prio 65535
```

..
Cisco also supports a bridge identifier mode called “MAC Address Reduction” where the same MAC address is used for all PVST+ or MISTP instances on the same switch. The bridge priority of the Nortel switch, like any other switch it that same broadcast domain, should be using values that are a factor of 4096. (i.e.: 4096, 8192, 12228, etc.)

If Spanning-Tree is enabled, it is most likely that the Nortel switch in the BladeCenter should NOT be the ROOT bridge for the network. Under normal conditions, the switch election will use the bridge priority and MAC address to determine which switch will become ROOT Bridge. To be sure the Nortel switch doesn't become ROOT Bridge, you can increase the bridge priority to the maximum value (do this across all enabled ST groups).

PVST+

In the GbESM implementation, an administrator creates a spanning tree group (STG) and then assigns a VLAN to it. This differs from the Cisco implementation where an administrator creates a VLAN, and then a spanning tree instance (i.e. STG) is automatically created and assigned to the VLAN.

The PVST+ interoperability implemented on the GbESM Ethernet Blade Switch includes the following features:

- Tagged ports which are members of multiple VLANs may belong to more than one STG (and thus transmit tagged Bridge Protocol Data Units (BPDUs) – used to elect a root bridge), but untagged ports can belong to only one STG.
- When a tagged port belongs to more than one STG, egress BPDUs are tagged to identify their STG membership.
- An untagged port cannot span multiple STGs.
- Each GbESM Ethernet Blade Switch supports 128 STGs operating simultaneously (STG 128 is reserved for switch management).
- The default STG 1 can hold multiple VLANs; all other STGs (groups 2–128) can hold one VLAN.

MSTP

You can also configure multiple instances of Spanning-Tree based on the standard IEEE 802.1s.

The GbESM Ethernet Blade Switch meets the IEEE 802.1W (Rapid STP) standard and is compatible with Cisco switches that are 802.1W compliant. The bridge priorities, port costs, port priorities, link type and edge port may be manually assigned on the GbESM Ethernet Blade Switch, allowing the core or other Catalyst switches to be the root bridge.

Note: When Multiple Spanning Tree is turned on, VLAN 1 is moved from Spanning Tree Group 1 to the Common Internal Spanning Tree (CIST). □When Multiple Spanning Tree is turned off, VLAN 1 is moved back to Spanning Tree Group 1.

Another Option: Switching OFF STP in the BladeCenter

The architecture of the IBM BladeCenter and its position on the edge of the network mean that there is practically no danger of a loop internal to the BladeCenter being formed. Spanning-tree can in many cases be turned off on the IBM BladeCenter switches. If required, they can continue to pass BPDU packets, but only as passive elements in the STP hierarchy. Using this design allows the Cisco uplink switches to determine the STP topology and treat the BladeCenter switches as either stub or transparent elements which can pass, but do not interpret its BPDUs.

To turn off spanning tree globally on the Nortel/IBM BladeCenter switches:

Nortel:

```
/cfg/l2/stg x off
```

If the external ports need to pass the uplink switch BPDUs (in the case that a cross-connect is realized using a cable between the two BladeCenter switches), the external ports only must be enabled for spanning tree:

```
/cfg/l2/stg x /port ext1 on  
/cfg/l2/stg x /port ext2 on
```

To verify the Spanning-Tree status:

Cisco: **show spanning-tree sum**

Nortel: **/info/l2/stg**

Note: Spanning-Tree values are often displayed in Hexa-Decimal.

When an interface becomes active, the Spanning-Tree phases (Blocking, Listening, and Learning) will begin. To accelerate that process, Cisco modifies the STP behavior to start forwarding frames immediately (PortFast) and verify any BPDU packets (BPDU Guard) that would put the port in “error disable”. Nortel’s equivalent is called “FastForwarding” and is perfectly inter-operable with Cisco PortFast.

Cisco: **spanning-tree portfast**

Nortel: **/cfg/port EXT1/fastfwd ena**

Rapid Spanning-Tree is based on industry standard IEEE 802.1w and is fully compatible between Nortel L2-3 GbESM and Cisco.

Cisco: **spanning-tree mode mst**

Nortel: **/cfg/l2/mrst/mode rstp**

Cisco: `spanning-tree mode mst`
Nortel: `/cfg/l2/mrst/mode mstp`

Nortel switches for the BladeCenter disable Spanning-Tree on all internal ports since they are directly connected to servers.

3. VLANs

VLANs can be defined on a Nortel switch and exchange packets with a Cisco infrastructure and vice-versa. The standard protocol to exchange VLAN information is IEEE 802.1q. Both Cisco and Nortel support this standard. Cisco also has an older proprietary protocol called Intra Switch Link (ISL) that is only supported in older Cisco switches and is **not** supported by Nortel GbESM. Also note that VTP is **not** supported on BLADE GbESM switches.

Although Cisco typically refers to 802.1Q VLAN tagging as “VLAN trunking” or “dot1q trunking,” the technologies are the same and, therefore, completely interoperable. There are two key items to ensure interoperability: ports on both ends of the tagged link (or dot1Q trunk) are assigned to same VLANs, and that the Cisco “native VLAN” matches the PVID assigned by the GbESM switch port.

You should first set the port for tagging

```
Cisco: interface <n>
        # switchport
        # switchport trunk encapsulation dot1q
        # switchport mode trunk
```

GbESM: `/cfg/port <n>/tag enable`

Then set the PVID. The same VLAN must be selected on both sides of a link for normal operation.

```
Cisco: switchport trunk native vlan <n>
AOS:   /cfg/port <n>/pvid <n>
```

Additional VLANs can be assigned to the port using the `/cfg/l2/vlan <n>/add` command. A port can have only one PVID at a time. One good practice is to select a VLAN that is not used for substantive traffic to be the native VLAN and use it to provide management access to the switches.

The 802.1q standard allows for a “native VLAN” to be defined on a tagged link. Frames for this VLAN will not be tagged. Cisco’s Native VLAN (by default, VLAN 1) is called Default VLAN (by default VLAN 1) on BLADE GbESM switches. A Native or Default VLAN is a VLAN that is part of an 802.1q link but doesn’t have an 802.1q header. The Cisco native VLAN is the only VLAN assigned to the port(s) that sends frames untagged. If you want to change the Native or Default VLAN:

```
Cisco: switchport mode trunk native x
AOS:   /cfg/port ext1/pvid x
```

To enable an 802.1q link:

```
Cisco: switchport mode trunk encap dot1q
AOS:   /cfg/port <n>
        tag ena
```

By default every MAC address is learned and stored in the Forwarding Database with the corresponding VLAN id and Port.

4. Link Aggregation

There are three basic ways to configure a link aggregation between a Nortel and Cisco switch. All links in a multi-link trunk must have identical characteristics, including:

- the same tagging parameter (enabled or disabled)
- the same set of VLANs and the same native VLAN (PVID)
- the same speed and duplex
- the same spanning tree parameters
- the same flow control parameters

Refer to the BLADE Network Technologies Nortel GbESM “Application Guide” and Cisco’s documentation for more configuration details and examples.

Method 1) This is the recommended configuration. Create a static EtherChannel (Port-Channel mode ON) on a Cisco switch and assign the desired interfaces to the Port-Channel. Create a static Trunk on the Nortel switch and assign the desired interfaces.

```
Cisco: router (config) interface port-channel 1
      router (config-if) exit
      router (config) interface fastethernet 0/1
      router (config-if) channel-group 1 mode ON
      router (config) interface fastethernet 0/2
      router (config-if) channel-group 1 mode ON
```

```
AOS: /cfg/l2/trunk 1
      ena
      add Ext1
      add Ext2
```

(Note: AOS **trunk 1** and IOS **port-channel 1** are not required to have the same values; these instance numbers are of local significance only.)

This simple configuration will create a static link aggregation between most Cisco and Nortel equipment.

To verify the Trunk Status:

```
Cisco: show int port-channel 1
      or
      Show int Gig <n/n> etherchannel
```

```
AOS: /info/l2/trunk
```


Method 2) Both Cisco and BLADE GbESM switches support the industry standard IEEE 802.3ad (LACP). You can create an LACP aggregation on both sides. Configure the administration key to be the same for the ports that you wish to group together and have at least one port to be configured as active (usually all ports are configured to be active). It is a best practice to configure the administration key to be the same on both ends of the link.

To view the state of the ports as well as the local and remote administration keys:

```
Cisco:      int Gig <n/n>
            channel-group <n> active (or passive)
```

(Note: do not use the other options (**always, preferred**, etc) which use the Cisco proprietary PAgP in place of LACP.)

```
AOS:      /info/12/lacp
```

Method 3) BLADE GbESM also allows a manual configuration of all possible parameters for creating a link aggregation. This method may also be used in combination with method 1 and 2, as described above. This configuration could potentially be useful if supporting older proprietary equipment or simply for optimizing traffic distribution over multiple links. When enabled, several L2 and L3 values are taken into account by the hashing algorithm used by the link distribution.

When ports are aggregated, the GbESM uses a hashing algorithm to choose which physical port a packet will be sent through. The parameters used for this hashing algorithm are configurable so that customers can ensure equal use of all of the physical ports. The defaults are:

- For a Layer 2 trunk, Source MAC and Source IP
- For a Layer 3 trunk, Source MAC and Destination IP

One or two of the following parameters can be included in the hash:

- Source IP or Destination IP
- Source MAC or Destination MAC (useful for Layer 2 only)
- L2 and L3 hashing can each have different parameters in use.

```
Cisco:      port-channel load-balance
```

```
AOS:      /cfg/12/thash
```

5. Routing Protocols

The Tolly Group and Blade Network Technologies have tested, and certified RIP, OSPF and BGP for compatibility with Cisco equipment. There has been full regression testing between BLADE Network Technologies and Cisco for all three of these routing protocols.

6. Management and Administration:

Management

BLADE GbESM switches can be monitored within any SNMP Management Framework, including Cisco's. Nortel's MIBs are available on the IBM website along with the latest firmware. You will find standard RFC MIBs and switch-specific MIBs that can easily be compiled within your SNMP console. After configuring the SNMP parameters, the switch will begin reporting SNMP traps to your framework.

```
AOS: /cfg/sys/ssnmp
```

Syslog servers can also be configured for reporting events information:

```
AOS: /cfg/sys/syslog
```

Remote Monitoring (RMON) is also supported for every interface for even further statistics:

```
AOS: /cfg/port Ext 1/rmon ena
```

Administration

BLADE GbESM switches support seven different levels of login access, each with related privileges. Refer to the BLADE "Command Reference" for specific details. These login accounts can be configured with local passwords or configured on a centralized security server. BLADE GbESM switches support both industry RADIUS and Cisco's TACACS+ protocols to communicate with a centralized security server.

Refer to the BLADE GbESM "Application Guide" for RADIUS and TACACS+ configuration. If you're using Cisco Access Secure (ACS) along with RADIUS for user authentication, make sure the Cisco ACS "User" or "Group" RADIUS IETF 006 Service-Type is "Administrative" instead of the default "Login".

BladeHarmony Manager

BladeHarmony™ Manager is a Java-based application that runs on Microsoft Windows®. It provides a graphical user interface (GUI) for Simple Network Management Protocol (SNMP)-based remote management of BLADE Network Technologies' Nortel® brand blade switch modules from BLADE Network Technologies. It is similar in functionality to CiscoWorks.

Requirements

- Nortel Layer 2/3 10Gb Uplink Ethernet Switch Module, firmware version 1.0
- Nortel Layer 2/3 Gigabit Ethernet Switch Module (GbESM), firmware version 1.2.3.11

Features and Benefits

BladeHarmony Manager offers a centralized point of administration of blade switch modules across multiple chassis and racks, including configuration management tools that facilitate time-consuming tasks, including bulk software/image upgrade to load and activate software updates and simplified management of complex configurations such as moving, or assigning ports to a large block of filters.

In addition, monitoring and graphing of switch functions for one or many blade switch modules is available.

This tool can be downloaded at:

http://www.bladenetwork.net/pages/products_blade_harmony.html

7. Additional Information

A useful source of information on Nortel Blade Server Switching can be found at:
<http://www.bladeserverswitching.org>

Documents

The IBM Redbook “**Nortel Networks L2/3 Ethernet Switch Module for IBM BladeCenter**”, in particular Chapter 7 covers in further details the integration of Nortel BladeCenter switches in a Cisco environment, and includes a number of sample configurations:
<http://www.redbooks.ibm.com/abstracts/redp3586.html?Open>



Independent Interoperability Certification by the Tolly Group

One question often asked by our customers is whether the Nortel BladeCenter switch is compatible and interoperable with their existing networking infrastructure (i.e.: Cisco, Enterasys, 3Com, etc.). The answer in almost every case is “yes!”.

Despite the fact that we have more than 100,000 units in production since 2003 and NO issues reported; we have recently gained an official, independent certification of our interoperability with other vendor’s equipment. This verification by an independent laboratory can be used to assure IBM Blade Center customers using Nortel GbESM switches that the switches will interoperate without issues when connected to the access network.

The Tolly Group was commissioned to certify the Nortel BladeCenter switch with many other vendors using several different protocols – the ones mostly likely to be found within the access network.

The following networking features were tested for interoperability:

- 802.1p/Q VLAN Tag Propagation Switch Interoperability
- Link Aggregation 802.3ad Switch Interoperability
- Accelerated Convergence Spanning Tree 802.1w Switch Interoperability
- IPv4 - RIP v2 - Routing Protocol Support Switch Interoperability
- IPv4 - OSPF - Routing Protocols Support Switch Interoperability
- IPv4 - BGP 4 - Routing Protocols Switch Interoperability
- Virtual Router Redundancy Protocol (VRRP) Switch Interoperability
- Jumbo Frame Support Switch Interoperability
- 10 Gigabit LAN Phy Support Switch Interoperability

- Multiple Spanning Tree Protocol (MSTP) Support (IEEE 802.1s)

You can review the full certification matrix on the web sites listed below or see page 112 of this document for more details.

- **Nortel Layer 2/3 Copper GbE Switch Module for IBM BladeCenter (32R1860):**
<http://www.tolly.com/TVDetail.aspx?ProductID=296>
- **Nortel 10Gb Ethernet Switch Module for IBM BladeCenter (39Y9267)**
Fixed port Layer 2-3 blade switch module for IBM BladeCenter H and BladeCenter HT chassis, with 14 internal and 6 external 10GbE ports:
<http://www.tolly.com/TVDetail.aspx?ProductID=294>
- **Nortel 10Gb Uplink Ethernet Switch Module for IBM BladeCenter (32R1783):**
Fixed port Layer 2-3 blade switch module for IBM BladeCenter H and BladeCenter HT chassis, with 15 internal GbE ports and 3 external 10GbE ports
<http://www.tolly.com/TVDetail.aspx?ProductID=295>

For additional Tolly Group reports on Blade Network Technologies IBM products, please go to:
<http://www.tolly.com/DocDetail.aspx?DocNumber=206168>

For More Information

For more information about BLADE Network Technologies, please visit
<http://www.bladenetwork.net>.

For technical information or for a list of BLADE Sales Network Architects, visit
<http://www.bladeserverswitching.org>.

Blade Network Technology's Sales Network Architect Team

Cole Kiblinger - SNA US West

Work +1 510 984 2814

E-mail cole.kiblinger@bladenetwork.net

Fred Rabert - SNA US Central, Canada and Wall Street

Cell +1 678 492 1528

E-mail fred.rabert@bladenetwork.net

Scott Lorditch - SNA US East and Federal

Cell +1 720 839 7192

E-mail scott.lorditch@bladenetwork.net

Derek Beattie - SNA EMEA North

Cell +44 (0) 7956 415 302

E-mail derek.beattie@bladenetwork.net

William Scull - SNA EMEA South

Cell +49 (0) 1718 549 172

E-mail william.scull@bladenetwork.net

Kavish Shah - SNA Asia Pacific
Cell +61 (4) 126 251 12
E-mail kavish.shah@bladenetwork.net

Charles Ferland - Resellers, EMEA
Cell +49 151 1265 0830
E-mail charles.ferland@bladenetwork.net

Peter Hall - Resellers, Asia Pacific
Cell +61 (4) 177 333 75
E-mail peter.hall@bladenetwork.net

Tracy Hickox - Resellers, North America
Cell +1 443 995 8486
E-mail tracy.hickox@bladenetwork.net

BLADE NETWORK TECHNOLOGIES: The industry's #1 provider of Ethernet, IP, and application switches for blade server systems. Our products represent a strategic control point in a blade server system. They are the gateway for all traffic directed to and from a blade server system. As a strategic control point, our products have a direct influence on a blade server system's performance, high availability, scalability, manageability, security, and total cost-of-ownership.

©2007 Blade Network Technologies, Inc. All rights reserved. All trademarks and copyrights are property of their respective owners. Information in this document is subject to change without notice. Blade Network Technologies assumes no responsibility for any errors that may appear in this document.

<http://www.bladenetwork.net>

MKT070630