IBM eServer BladeCenter and Cisco ESM Technology
Planning / Implementation (withdrawn product)

Main

IBM eServer BladeCenter and the Cisco Systems Intelligent Gigabit Ethernet Switch Module hints and tips
In this section, we have provided hints and tips that we discovered during our residency that will help you in the setup and configuration of your Cisco Systems Intelligent Gigabit Ethernet Switch Module (CIGESM) for the IBM eServer BladeCenter.

1. Blade server NIC numbering
The topic of which connection on a given blade server goes to which Cisco Systems CIGESM in the BladeCenter is the subject of some discussion.

In *most cases*, for Windows™ 2000, the connection named *Local Area Connection* goes to the Cisco Systems CIGESM in switch bay 1 (referred to as CIGESM1 in Chapter 7 of the Redpaper REDP3869), and the connection named *Local Area Connection 2* goes to the Cisco Systems CIGESM in switch bay 2 (referred to as CIGESM2 in chapter 7).

The phrase *most cases* is used above, as this is not always the case.

For Windows 2000, the order of the Local Area Connection *names* assigned is based on the *order* the drivers for each NIC are installed. The drivers necessary to support the NICs on a blade server are not part of a standard Windows 2000 install, and the NICs will be generically listed in Windows 2000 Device Manager as two or more *Ethernet Controllers* until the necessary drivers are loaded. For these NICs to become active, a third party driver, supplied by IBM, needs to be installed. The *normal* procedure most users follow is to install the drivers on the first *Ethernet Controller* in the list, and then install the drivers on the second *Ethernet Controller* in the list (and so on). The end result of this is the "most cases" scenario mentioned above, where the Windows 2000 connection named *Local Area Connection* goes to CIGESM1 and the one named *Local Area Connection 2* goes to CIGESM2.

If, however, the drivers are installed on the second *Ethernet Controller* in the list first, and then the first *Ethernet Controller* in the list, the connection names are reversed and the connection named *Local Area Connection* is now the one going to CIGESM2 and the connection named *Local Area Connection 2* would now be going to CIGESM1.

Important: To avoid confusion, always install drivers sequentially, from the first *Ethernet Controller* in the list to the last *Ethernet Controller* in the list.

For Linux™, eth0 goes to the Cisco Systems CIGESM in switch bay 2 (CIGESM2 in this chapter) and eth1 goes to the Cisco Systems CIGESM in switch bay 1 (CIGESM1 in this chapter). Note that this is reversed from a normal Windows 2000 install as mentioned above, and also proves to be the source of some confusion.

2. Default gateway configuration on multihomed servers
The blade servers in the eServer BladeCenter by default have two connections to the network, each usually on separate IP subnets. Using the Broadcom teaming software allows you to increase the number of subnets configured on a blade server above and beyond that available on the physical connections. The end result is, the blade server frequently has more then a single IP subnet configured.

One question that frequently comes up is, should each IP subnet configured on a multihomed system, such as the blade server, have a default gateway assigned? The answer is far from straight forward.

For all examples in this chapter, only one interface receives a default gateway, and the other interfaces are left blank for the default gateway field. This is not to suggest that this is the best solution for your environment, which will have its own unique requirements. It is only used in these examples for simplicity sake.

Microsoft has published Knowledge Base article 157025 that discusses the pros and cons of different approaches with default gateways on multihomed systems. This article can be found at:
http://support.microsoft.com/default.aspx?scid=kb;en-us;157025
We recommend that the BladeCenter system administrator review this document if there are questions on how best to address this issue in your specific environment.

3. Duplicate IP address - Part 1
Several configuration choices can result in the BladeCenter reporting duplicate IP addresses, even when there are no known conflicts. Duplicate IP address - Part 1 discusses an issue with blade servers reporting a duplicate IP address.

The most common cause of a blade server reporting a duplicate address is the result of placing one of the interfaces of a blade server in the same VLAN as the Management VLAN of the Cisco Systems CIGESMs (default is VLAN 1), with an IP address in the same subnet as the Management address used internally by the Management Modules to communicate with the Cisco Systems CIGESMs.

In this case, the Management Module will try to act as a proxy for all addresses on this subnet (as part of allowing external access to the Cisco Systems CIGESMs through its external interface) and will answer to any query for any address on the entire internal Management subnet. In this case, when the blade server checks to see if its address is available on the network (sends out an ARP request for its own address), the Management Module will respond to this ARP request for the blade server address and the blade server will assume the address is in use and report this to the user via a Windows 2000 pop-up message.

The solution is to always keep blade servers off of the Cisco Systems CIGESM Management VLAN. The default Management VLAN for the Cisco Systems CIGESMs is VLAN 1. To reduce the likelihood of blade servers being placed on this VLAN, the Cisco Systems CIGESM's sets the defaults to all ports going to the blade servers (g0/1 - g0/14) to the following: switchport access vlan 2
switchport trunk native vlan 2
switchport trunk allowed vlan 2-4094
This, however, does not prevent the user from adding VLAN 1 (or whatever the Management VLAN is) to ports going to the blade servers, and thus the consequences as described above.

Important: In light of this possible interaction between the blade servers and the Management Modules, it is highly recommended not to place blade servers on the same VLAN being used for the management VLAN on the Cisco Systems CIGESMs.

4. Duplicate IP address - Part 2
As noted in Duplicated IP address - Part 1, several configuration choices can result in the BladeCenter reporting duplicate IP addresses, even when there are no known conflicts. Duplicate IP address - Part 2 discusses an issue with the Cisco Systems CIGESM reporting a duplicate IP address.

The most common cause of a Cisco Systems CIGESM reporting a duplicate IP address is trying to change the Management IP address for the Cisco Systems CIGESM directly on the Cisco Systems CIGESM (either through CLI or through CMS).

When a user changes the Management VLAN IP address on the Cisco Systems CIGESM through means other then the Management Module Web interface, to something other then what it received from the Management module, all IP communications to the Cisco Systems CIGESM's Management IP address fail and the Cisco Systems CIGESM begins to report a duplicate IP address. Note that this duplicate IP address message only happens if you change it to an address in the same subnet that it was originally on. For example, changing from 192.168.70.127 to 192.168.70.150 would result in a duplicate IP address message, while changing from 192.168.70.127 to 10.35.15.1 would not result in a duplicate IP address message (although it would more then likely result in lost communications via IP to the Management Module).

The following sequence of events demonstrates this issue (for demonstration purposes only, do not perform on production systems):

Show current IP address of Management VLAN on the Cisco Systems CIGESM to confirm proper configuration (in this example, the default IP address is being used) CIGESM1#sh run int vlan 1
INTERFACE Vlan1
ip address 192.168.70.127 255.255.255.0
no ip route-cache
Test ping from the Cisco Systems CIGESM to the internal IP address of the Management Module to verify the connection is working:
CIGESM1#ping 192.168.70.126
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.70.126, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5)

Change IP address on Cisco Systems CIGESM to another address: CIGESM1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
CIGESM1(config)#int vlan 1
CIGESM1(config-if)#ip add 192.168.70.150 255.255.255.0
CIGESM1(config-if)#
Shortly after making this change, start receiving duplicate address messages on the console of the Cisco Systems CIGESM:
1d19h: %IP-4-DUPADDR: Duplicate address 192.168.70.150 on Vlan1, sourced by 0009.6bca.7499

Now change back to the original address:
CIGESM1(config-if)#
CIGESM1(config-if)#ip add 192.168.70.127 255.255.255.0
CIGESM1(config-if)#

And continue receiving duplicate address messages:
1d19h: %IP-4-DUPADDR: Duplicate address 192.168.70.127 on Vlan1, sourced by 0009.6bca.7499

Test ping from the Cisco Systems CIGESM to the internal IP address of the Management Module and it fails:
CIGESM1#ping 192.168.70.126
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.70.126, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

The cause of this issue is related to the duplicate IP address message reported in Part 1 (the Management Module responding to arp requests for addresses on its own internal subnet). The way to prevent this situation from occurring is to only change the IP address of the Management VLAN via the Management Module's Web interface.

The way to resolve this issue if it is already occurring:

- Resolution 1
  The safest and preferred approach is to connect into the Management Module's Web interface, click Management under I/O Module tasks on the left side of the screen, select the bay that is experiencing the issue (bay 1 in our example above), make sure the IP address shown for the Cisco Systems CIGESM in bay 1 is set as desired (if not, change it to the desired address), and click on the Save button. This will put the IP addresses back in sync between the Management Module and the Cisco Systems CIGESM, and communications should be restored (as well as the duplicate IP addresses messages will stop).
- Resolution 2
  The following will also resolve the conflict, but it will disrupt traffic while the Cisco Systems CIGESM is rebooting. Set the IP address of the Cisco Systems CIGESM back to its original value, save the config and reload the Cisco Systems CIGESM. The Cisco Systems CIGESM and the Management Module will be back in sync upon rebooting and coming back on line.

Important: To reduce the likelihood of this issue occurring, do not change the IP address of the Cisco Systems CIGESMs directly on the Cisco Systems CIGESMs. Only change the address of the Cisco Systems CIGESMs through the Management Modules Web based interface.

5. Teaming software on a blade server forces an undesired action

The Broadcom software, known as the Broadcom Advance Control Suite (BACS), is used to control NIC teaming on the blade servers. It has been noted that there are several selections within this software that do not permit you to cancel or otherwise back-out of a selection, and appear to force a user to perform an action they may not wish to perform.

Two such examples are when you click Remove VLAN at a Team Configuration window or select Tools->Delete a team from the menu bar. In these cases, a window similar to Figure-1 will be displayed for a user to specify a VLAN or a team to delete, and you will not be offered any obvious way to cancel the operation. To abort these seemingly un-abortable procedures, press the Esc key on the keyboard.

Note that if there is only one item in the list to delete (for example, only one VLAN or one team), the BACS software does not even give an option to select what is to be deleted, it simply deletes the one item and goes back to the original screen. If this deletion was undesired, the only solution is to click the Cancel button on the main BASC window, and exit the BASC software without saving the changes. Of course doing so will result in loosing any other changes you had already made since the last time the Apply or OK button was selected.

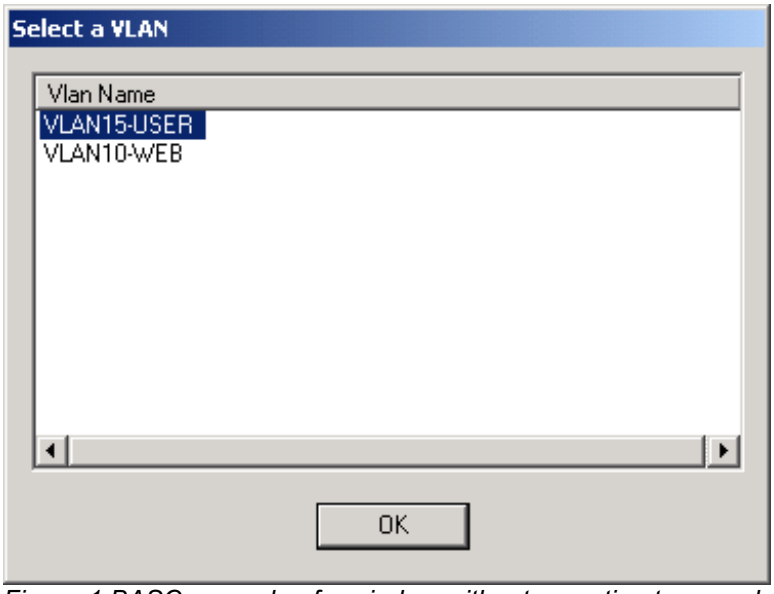

Figure-1 BASC example of a window without an option to cancel

6. Addressing issues with links not coming up

During the course of writing this document, much time was spent in the lab building and testing many configurations. During this process there were times when, no matter how the configuration was set up, the link would just not come up, or pings would not work, or any number of similar symptoms were encountered.

While the exact causes of these conditions can be debated, there was one solution that usually solved the problem and that was to take the link down and back up.

For the purposes of troubleshooting this kind of issue, the following steps may be of assistance:

- Check your cabling. Make sure it's going to the desired ports and that the cables are good
- Check your configs. Look closely as it's easy to miss a single command that will cause the link not to come up
- If all of this fails, try taking down the link and bringing it back up
    - For a Cisco Systems CIGESM or a 6500, this means doing a shut/no shut at the interface and/or any associated Port-Channel.
    - For a blade server running Windows 2000, this means disabling and then re-enabling the connection via the Windows 2000 network interface.

7. Cisco Systems CIGESM stuck at switch: prompt

Any keystrokes received via the serial console connection during the early phases of the boot up process of the Cisco Systems CIGESM will be interpreted as a break signal and will put the Cisco Systems CIGESM into an incomplete boot up state, with a prompt that simply says switch:. The following is a sample display that comes up when keys are pressed during the boot up process: The system has been interrupted prior to initializing the
flash filesystem. The following commands will initialize
the flash filesystem, and finish loading the operating
system software:
flash_init
load_helper
boot
switch:

Note that you may not see this entire message and only see the switch: prompt. Either way, the switch has not completed the boot process and will not switch traffic or otherwise function until the boot up process is completed.

If you find the switch stuck at this prompt, you can finish the boot up process by entering flash_init, waiting for the flash to initialize and return to the switch: prompt and then entering boot.

The switch should begin loading its flash image and complete the boot sequence and become fully operational shortly thereafter.

Note that the Cisco Systems CIGESM can also stop at this prompt under certain POST failure conditions. Following the procedure above may or may not help under POST failure conditions.

8. Key sequence to switch between blade servers

The BladeCenter has a Keyboard/Video/Monitor (KVM) switch built in to the Management Modules, allowing traditional access to the installed blade servers. To switch the keyboard, mouse and monitor between blade servers, perform the following key stroke combination from the keyboard attached to the active Management Module:
NumLock NumLock <blade server number> Enter

Where <blade server number> is the number of the blade server bay where the blade server is installed. For example, to select the blade server in bay 2, press the NumLock key twice, then the number 2 and then the Enter key:
NumLock NumLock 2 Enter

Note: Often times after selecting a blade server via the sequence above, the display is blank. Moving the mouse will usually bring the screen up (bring it out of screen-saver mode).

9. Native VLAN mismatch message

When changing the management VLAN on the Cisco Systems CIGESM you may receive a *native VLAN mismatch* message on the console of the Cisco Systems CIGESM. This is due to the fact that changing the management VLAN on a Cisco Systems CIGESM also changes the native VLAN on ports g0/15 and g0/16 (default native VLAN for these ports is VLAN1). Ports g0/15 and g0/16 connect, via the Management Module to all other Cisco Systems CIGESMs in the BladeCenter. When you make the initial change on one Cisco Systems CIGESM, the native VLAN will still be different on any other Cisco Systems CIGESM in the BladeCenter (on ports g0/15 and g0/16) until such time as their management VLAN is changed as well.

If you only have a single Cisco Systems CIGESM in the BladeCenter, this message will not occur when changing the management VLAN.

Note: For correct operation, the management VLAN on all Cisco Systems CIGESMs in a BladeCenter must be the same. To resolve the *native VLAN mismatch* message, change the management VLAN to the same VLAN for every Cisco Systems CIGESM in a given BladeCenter.

10. Use of RSPAN on the Cisco Systems CIGESM

Testing during the preparation of this redpaper showed an issue with configuring RSPAN on the Cisco Systems CIGESM, when utilizing IOS release 12.1(14)AY (the revision available at the time of this writing). Under certain circumstances, several interfaces, including the link to the upstream switch and the port being monitored, would begin to stream data at wire rate, resulting in, at a minimum, lost communications to the device on the port being monitored and issues on the upstream switch.

This RSPAN issue was traced to a bug with release 12.1(14)AY of IOS. A fix is under test and should be available in an upcoming release of IOS for the Cisco Systems CIGESM.

Note that the Software Configuration Guide for the Cisco Systems CIGESM (IBM document 25k8411, First Edition, June 2004) does not currently reflect the correct procedures for configuring RSPAN on the Cisco Systems CIGESM (at a minimum it neglects to mention the step to run remote-span on the VLAN being used as the destination of the RSPAN session). This is under review and should be updated soon to reflect the correct procedures.

Also note that if you have already configured RSPAN and are experiencing the issue of streaming data as described above, deleting the monitor session associated with the RSPAN will halt this condition (from config term mode run the command: no monitor session x, where x is the monitor session number configured for RSPAN use).

Important: Extreme caution is recommended when utilizing the RSPAN feature on the Cisco Systems CIGESM until this issue is resolved. Please contact your IBM technical support representative for more information or to see if a fix is available. Reference Cisco IOS caveat CSCee53625.

Other BladeCenter hints and tips
For other hints and tips for the BladeCenter, please see the following document:
http://www-306.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-45277

Related product families

Product families related to this document are the following:

- Blade Networking Modules

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION ”AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

This document, TIPS0423, was created or updated on July 6, 2004.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
  https://lenovopress.lenovo.com/TIPS0423
- Send your comments in an e-mail to:
  comments@lenovopress.com

This document is available online at https://lenovopress.lenovo.com/TIPS0423.

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at https://www.lenovo.com/us/en/legal/copytrade/.

The following terms are trademarks of Lenovo in the United States, other countries, or both:
Lenovo®
BladeCenter®

The following terms are trademarks of other companies:

Linux® is the trademark of Linus Torvalds in the U.S. and other countries.

Microsoft® and Windows® are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.