

Sun Fire™ Blade Application Journal

Issue 2/2004



Configuration Tips & Techniques for the Sun Fire B1600 Blade Platform



In This Issue...

Load-Balancing Architectures	2
Installation Overview	3
Install Solaris OS and Web Server	5
Configure the Load Balancing Blade	7
Set Up VLANs	7
Set Up Load Balancing Blade Failover	9
Troubleshooting Tips & Techniques	9
Conclusion	10
Acknowledgements	10

Content Load Balancing In Your Sun Fire B1600 Blade Platform

The Sun Fire B1600 Blade Platform offers more choice than any other blade server on the market today. You can choose between UltraSPARC® and x86 architecture server blades. You can choose between single and dual processors on x86 architecture server blades. And you can choose between running Sun's Solaris™ Operating System on either architecture — and you can run Linux on Sun's x86 architecture server blades.

This range of choice is enhanced with specialty blades designed to make it easy to configure high-availability, high-performance, and highly-secure networks and services all on a single blade platform, across multiple blade platforms, and even across external, stand-alone servers. The first two of



Figure 1: Sun Fire B10n Content Load Balancing Blades

Sun's specialty blade offerings are the Sun Fire B10n Content Load Balancing Blade and the Sun Fire B10p SSL Proxy Blade, which allow you to integrate load-balanced services and hardware-assisted Secure Socket Layer (SSL) encryption into your blade environment.

The Sun Fire B10n Content Load Balancing Blade incorporates a custom hardware Layer 4-7 classification engine that supports content-based load balancing based on URLs, CGI scripts, and cookies. You can use the hardware-assisted classifier, for example, to split loads for different types of content onto different load-balanced service groups, allowing you to tune your services for best response times regardless of Web site structure. After classification, the load balancing blade can balance load across multiple server blades in a service group on the basis of weighted round robin, response time, weighted least connection, and static algorithms. Finally, you can improve service response times further with the blade's server-to-client direct response that

Setting Up Web Services With Sun Fire B10n Content Load Balancing Blades

Learn how to balance workloads across horizontally-scaled service groups of Sun Fire B100s Server Blades. In the next issue, learn how to add hardware-accelerated SSL encryption with Sun Fire B10p SSL Proxy Blades.

allows the you to use the full gigabit Ethernet bandwidth of each server blade to respond to client requests.

This issue of the *Sun Fire Blade Application Journal* helps you through the process of setting up load balancing using Sun Fire B10n Content Load Balancing blades. We'll discuss how to set up a single server group using the Solaris OS running on UltraSPARC processor-based Sun Fire B100s server blades. (Setting up the Linux OS is similar). We'll focus on how to set up component-level redundancy with a pair of load balancing blades in a failover configuration and at least two server blades to provide server-level redundancy.

In a future issue, we'll cover path-level redundancy, using two SSCs in the blade platform and Solaris IP Multi-Pathing (IPMP) to provide continuous service despite the failure of a single switch or network connection.

Load-Balancing Architectures

Before delving into the details of how to set up load balancing on your blade platform, we'll take a look at the differences between traditional load-balanced services and load balancing using the Sun Fire B10n Content Load Balancer.

Traditional Load Balancing

A traditional load-balancing arrangement using horizontally-scaled Web servers is illustrated in Figure 2. A load-balancing switch distributes incoming service requests across a set of servers known as a *service group*. Two service groups are shown: one group of three servers supports a Web site and the other group of two servers support an FTP site. The servers respond to requests and pass their responses through the load-balancing switch and back to clients on the Internet.

This example illustrates the multiple networks used to support load balancing configurations: a *data network* carries requests to servers and responses back to the clients.

Service networks carry request and response

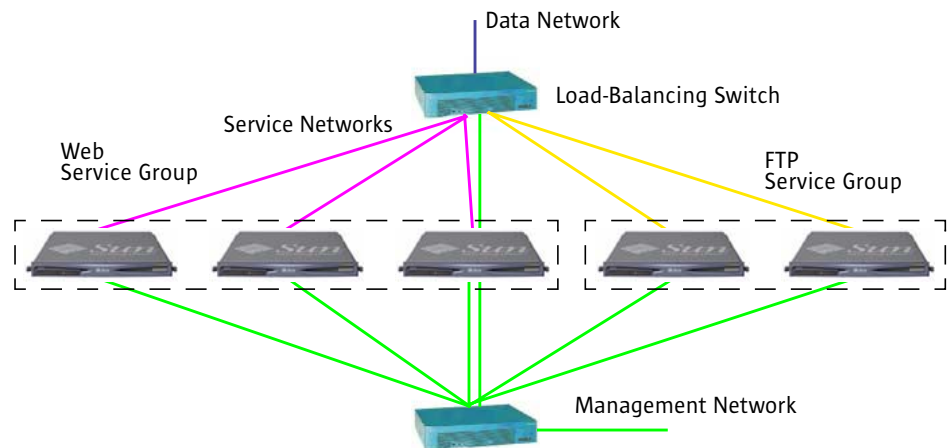


Figure 2: Traditional load balancing across two service groups using discrete components.

traffic from the load-balancing switch to the servers supporting each service. And a *management network* supports network management tools and allows you to install software and content on the servers.

A typical load-balancing configuration would include redundant pairs of load-balancing switches. Even though this example only illustrates a single load balancer, you can see that you must rack multiple discrete components, cable multiple physical networks, adding to the infrastructure you must maintain. If you change a configuration (like adding or removing a server) you must again change cabling, where mistakes can cause outages and significant security issues.

Benefits of the Sun Fire B1600 Blade Platform

Using Sun Fire B10n Content Load Balancing Blades in the Sun Fire B1600 Blade Platform, you can support the network illustrated in Figure 2 in only half of the space used by discrete components, with room to add or scale services as your business needs dictate.

Using the blade platform, you also increase quality of service because each server and content load balancer is supported by dual power supplies and a redundant network fabric — both of which contribute to increased uptime. The more servers you add to the platform, the more economical each one

becomes, because the cost of the dual power supplies — usually prohibitive in 1U servers — is amortized across a larger number of servers (up to 16 in a Sun Fire B1600 Blade Platform).

All of the interconnections illustrated in Figure 2 (except the data network) are contained within the platform using Virtual Local Area Networks (VLANs). Using VLANs in a Sun Fire B1600 Blade Platform means that you can change your network configuration without touching a network cable, reducing the possibility of cabling errors compromising security or taking a service offline. For system integrators and service providers it means being able to duplicate configurations for different customers simply by downloading a pre-defined switch configuration file.

Using the Content Load Balancing Blade

The Sun Fire B10n Content Load Balancing Blade uses the blade platform's built-in Switch and System Controllers (SSCs) to route incoming requests to the proper server over a set of VLANs. In a sense, the load balancing blade acts as the brains for the switch that's already part of the platform, economically leveraging its resources.

Physically, the load balancing blade has two gigabit Ethernet interfaces that connect to each of up to two SSCs installed into a blade platform. Logically, the load balancing blade




Figure 3: The Sun Fire B10n Content Load Balancing Blade connects to VLANs rather than to physical networks.

interfaces with VLANs that correspond to the physical networks shown in the discrete networking example (Figure 3).

- Requests from clients are accepted by the load balancing blade on the data network.
- The load balancing blade passes requests to servers over the appropriate service network VLAN (one network per service).
- The blade uses the management network to communicate with the content load balancing module installed on each server blade.

One of the key differences between load balancing with the Sun Fire B10n Content Load Balancing Blade and traditional load-balancing architectures is the use of server-to-client direct response. As Figure 4 illustrates, requests from clients arrive on the data network and are passed to server blades on the services network. Responses return directly to clients via the data network, allowing you to use the full bandwidth of the blade platform’s switches to deliver content.

Blade Platform VLAN Configuration

When you set up your Sun Fire B1600 Blade Platform for load balancing, you can choose whether or not to use VLANs to separate traffic between networks. Whether or not you use VLANs, you it’s best to configure three different subnets for the data, management, and services networks.

Figure 5 shows a single SSC with its eight external ports (NETP0-7), its management port (NETMGT), and its internal, blade-facing ports (SNP0-15). Five slots are occupied by blades: three single-processor server blades are assigned to the Web service group; two server blades are assigned to the FTP service group, and one slot is occupied by the Sun Fire B10n Content Load Balancing Blade.

Each blade is connected to the switch by its first interface, and several VLANS are illustrated using colored lines:

- The (blue) data network extends to the data-center network through external port NETP0, and the VLAN to which it is assigned is connected to each of the five blades.
- The (green) management network VLAN extends to each of the blades and is connected through the SSC’s built-in packet filter to its NETMGT port.
- The (magenta) Web service network connects the load balancing blade with the three Web server blades. Note that the VLAN illustrated in magenta is not connected to any external port — the service

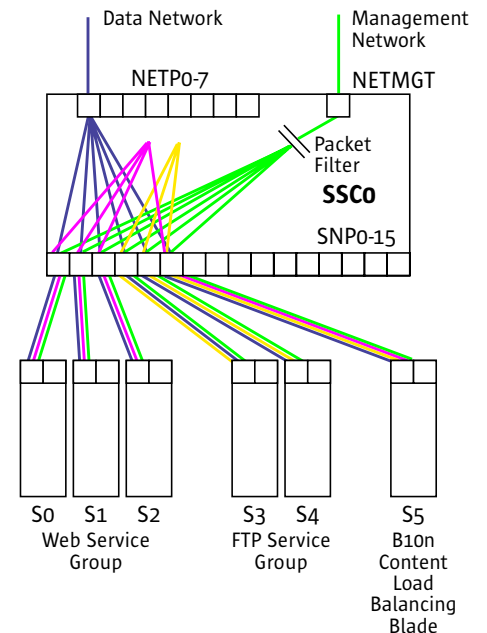


Figure 5: Example Load Balancing configuration using VLANs in the Sun Fire B1600 Blade Platform.

networks can be completely contained within the blade platform itself.

- The (yellow) FTP service network connects the load balancing blade with the two server blades in the FTP service group.

Although neither of the service groups are shown connected to an external switch port, they might be connected to an external network for two reasons:

- The Sun Fire B10n Content Load Balancing Blade can support multiple blade platforms depending on the service and the workload. If more than one blade platform is to participate in a load-balancing configuration, data, service, and management networks can be extended to additional platforms.
- Likewise, external stand-alone servers can also participate in load-balanced service groups. Today, Sun supports stand-alone Sun Fire V60, V65, V210, and V240 servers.

Installation Overview

In this issue of the *Sun Fire Blade Application Journal* we’ll show how to implement a simple load-balanced service group with an emphasis on component-level redundancy.




Figure 4: Server-to-client direct response sends large amounts of data directly to clients, using the full bandwidth of the switch and reducing latency.

We'll set up a service group containing two Sun Fire B100s server blades running the Solaris Operating System, and two Sun Fire B10n Content Load Balancing Blades in a fail-over configuration. This simple configuration will illustrate the basics of configuring the platform, server blades, and load balancing blades, making it easy for you to extend the basic configuration to one that meets your specific needs.

The blade platform VLAN configuration that we'll set up is illustrated in Figure 6:

- The (blue) data network connects the external NETPo port with the load balancing blades and the two server blades running a Web server. This network supports the server-to-client direct response. It is assigned VLAN 21.
- The (magenta) services network is contained entirely within the blade platform, and it connects the load balancing blades with the two servers in the service group. It uses VLAN 23.
- The (green) management network is also contained within the blade platform, connecting the load balancing blades with the two servers in the service group. It is assigned the default VLAN 2.

Using this configuration, you can access each blade's serial console through the SSC accessible through the platform's network management port. In real life, you might wish to establish a separate, private management network that's extended to one of the external switch ports. You also might wish to open specific ports on the SSC's built-in packet filter so that management traffic can pass to the external NETMGT port. These options are illustrated by the dashed green lines in Figure 6. We won't discuss how to configure these networks in this issue, however if you wish to use the platform's management network, you can find instructions for configuring the packet filter in the *Sun Fire B1600 Blade System Chassis Software Setup Guide*.

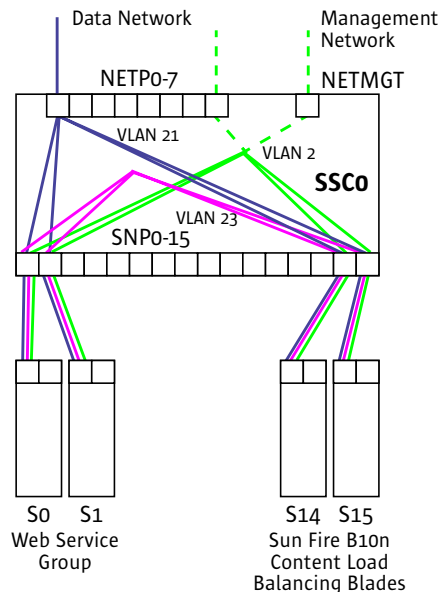


Figure 6: Load balancing VLAN configuration.

Overall Tips & Techniques

The set-up process involves setting up the Solaris OS and a Web server on two (or more) server blades, setting up the Sun Fire B10n Content Load Balancing Blade, and setting up the platform's VLAN configuration. There are several overall tips and techniques that you can use to make the process go more easily — especially if this is the first time you set up load balancing on a blade platform.

- *Use an Incremental Approach*
First set up the OS and the Web server, setting up all of the necessary IP addresses without using VLANs. This way you can first test the Web server and even load balancing using the default switch configuration. Then set up the load balancing blades and test load balancing operation, again without the use of VLANs. Once you're satisfied that you have a working configuration, then add VLANs and a fail-over load balancing blade configuration.
- *Change Configuration Files and Reboot*
Most of the set-up manuals for the blade platform illustrate how to configure settings (like the server-based Content Load

Balancing (CLB) module) on a live system. Since you'll eventually reboot the server blades, don't bother with this step: modify the necessary configuration files and then reboot. That way you'll be sure that your configuration files are correctly set up, and you'll also see any error messages during the reboot process.

- *Use Solaris™ Flash Archives*

If you're setting up load balancing, you'll be setting up two or more server blades with extremely similar configurations. Nothing makes duplicating a configuration easier than using Solaris Flash archives. Using Solaris Flash technology, you can set up one server as a "golden master," use the `flarcreeate` command to create a Flash archive of the configuration and install the same image on as many servers as you like by specifying the archive in your JumpStart™ software configuration. Once you completely understand the configuration steps necessary to set up load balancing, you can write finish scripts to set up each server blade's unique IP addressing automatically. You can read about JumpStart and Flash technology in Issue 11/2003 of the *Journal*.

- *Commit Changes*

When you make changes to the load balancing blade configuration, be sure to execute a `commit` command to save your configuration across power cycles.

Gather Resources

You'll make the configuration go smoothly if you first gather the resources in the "What You'll Need" sidebar and allocate IP addresses according to the "Network Address Checklist" sidebar. Look at the tutorial in the appendix to the Administration Guide for additional set-up guidelines.

You'll need three networks for the configuration described in this issue. The VLAN and network numbering used in the example are summarized in Table 1, and the addressing

What You'll Need

Servers and Software Resources

- Sun Fire B1600 Blade Platform
- Two or more Sun Fire B100s Server Blades
- Two Sun Fire B10n Content Load Balancing Blades
- Up-to-date load balancing blade and SSC firmware (from www.sun.com/software/download/network.html) Download both Solaris OS drivers and load balancing blade firmware.

Useful Documents

- *Sun Fire B10n Content Load Balancing Blade Administration Guide* (from www.sun.com/products/networking/blades/lb, click on Documentation)
- *Sun Fire B10n Content Load Balancing Blade Product Notes* (from the same location)
- *Sun Fire B1600 Blade System Chassis Software Setup Guide* (from the same location)
- *Sun Fire Blade Application Journal, Issue 11/2003* (from www.sun.com/products/blades) Provides information on setting up the Sun Fire Blade Platform and Sun Fire B100s Server Blades)

and VLAN layout is illustrated in Figure 8 on page 6. Addressing used in this example is:

- A data network accessible to the clients your configuration is to serve, for example Web browsers. The network it might be a lab or a DMZ network.
- A management network. This example uses the 192.168.0 network with a netmask of 255.255.255.0. This network must be different than the data network if you use VLANs.
- A services network. Traffic on this network is transferred from MAC address to MAC and IP network addressing is not required.

Network Name	VLAN	Numbering
Data	21 (blue)	192.207.20/24
Management	2 (green)	192.168.0/24
Service	23 (magenta)	0.0.0.0

Table 1: Example network numbering and VLAN assignment.

The Virtual IP (VIP) address is the destination for requests for clients, and is the source address for responses. It is configured on the load balancing blade and on each server blade.

If you follow an incremental procedure to set up your first service group, you don't have to implement VLANs in order to get your server blades and load balancer up and running. Using VLANs is recommended for production environments because they partition traffic on a service group basis. If you plan to add hardware-accelerated SSL encryption using Sun Fire B10p SSL Proxy Blades, you'll especially want to use VLANs because they separate private, plaintext traffic from externally-visible encrypted traffic.

Install Solaris OS and Web Server

The first step to take in setting up a load balanced service group is to configure at least two server blades running the service you wish to provide. For this example, we'll set up Sun Fire B100s server blades running the Solaris Operating System and a Web server.

Understanding the end result will help you understand exactly how load balancing and the server-to-client direct response works. An example server blade OS and Web server configuration is illustrated in Figure 7:

- Services and management network addresses are configured on their respective VLANs. Traffic on the services network is transferred from MAC address to MAC address so no IP addresses are necessary.
- The data network address can be configured using a private address not actually on the data network because it is not used as a source address for server-to-client direct response packets.

Network Address Checklist

Load Balancer Addresses

These addresses will be configured on the content load balancing blade:

- Virtual IP Address for Service (VIP), netmask, and VLAN number
- Management network address, netmask and VLAN number

Server Blade Addresses

If you use VLANs, you'll need one set of addresses for each server blade configured in the service group:

- VIP (same as above)
- Management network address
- Service Network Address (can be 0.0.0.0 for each blade)
- Data network address (can be a private address not on the data network)

- The loopback address, in addition to its default address of 127.0.0.1, also supports the service group's Virtual IP (VIP) address
- The default route sends responses onto the data network.
- The Web server binds to port 80 on the VIP address configured onto the loopback device.

Here's the magic: When the load balancing blade sends a request to the server blade, it cooperates with the content load balancing module (a STREAMS module) on the server to send the request to the VIP configured on the loopback device. The Web server receives the request and sends its response back to the loopback device. This step maintains the VIP as the source address of the response. Finally, the response is sent out onto the data network due to the default route setting.

Note that the data network address never comes into play, so the interface can use a private network address. This has two benefits:

- You don't have to allocate a public IP address for each server blade in the service group.
- It's somewhat more difficult to attempt to attack the server blade because its address is not on the network onto which it is configured.

Update SSC Firmware

Before you begin configuring Web servers and load balancing blades, check to be sure that your SSC is running the current firmware version. Use the `showssc` command to see what version it's currently running, and check the current version on the software download site listed in "What You'll Need" on page 5. Follow the instructions in the *Product Notes* to update your SSC firmware if necessary.

Install Solaris OS and Web Server

Install the operating system software and configure a Web server on each of the two server blades:

- First install the Solaris Operating System on your server blade using a network install server. If you haven't already set up an

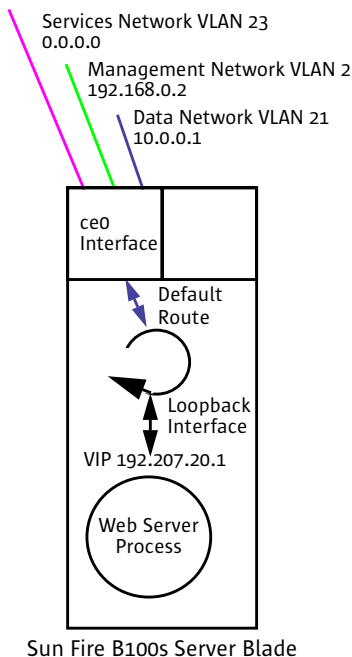


Figure 7: Example server blade and Web server network configuration.

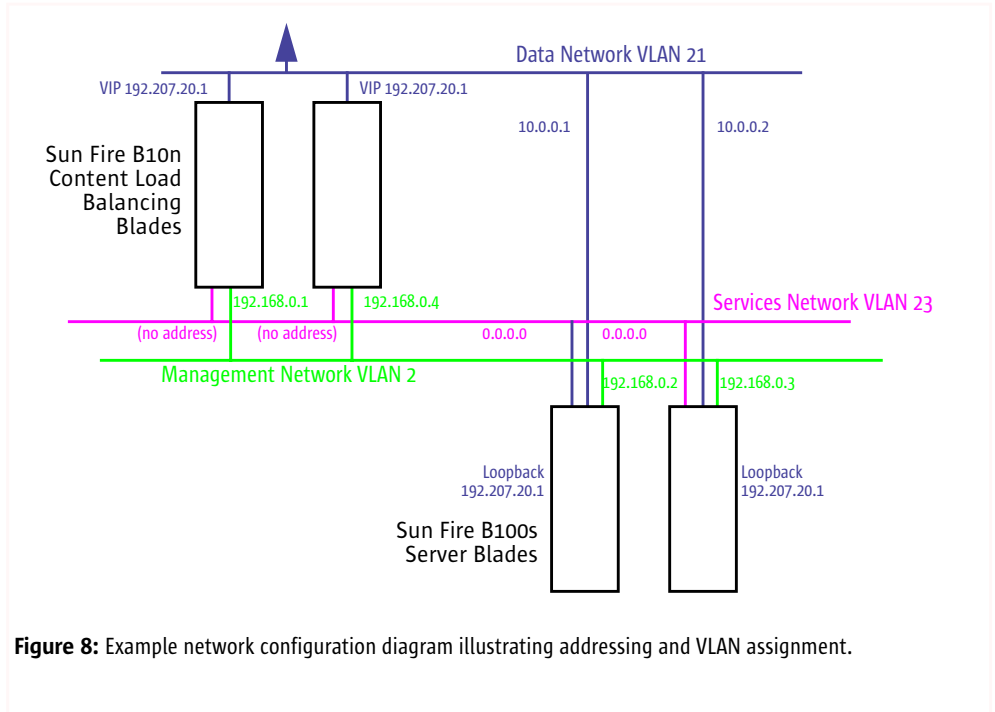


Figure 8: Example network configuration diagram illustrating addressing and VLAN assignment.

install server, refer to issue 11/2003 of the *Journal*. If you plan to use the Apache Web server included with the Solaris OS distribution, be sure to install the cluster SUNWCall. Give the server a host name and address that's accessible on your LAN for testing.

- Install the required Solaris OS patches listed in the *Product Notes*.
- Install the Content Load Balancing (CLB) module as described in the *Product Notes*. The CLB module acts as the load-balancing blade's agent on each server blade. It provides server-related information and exchanges traffic forwarding information with the load balancing blade. Whereas load balancing functions are all handled in hardware on the load-balancing blade, the CLB module extends the software functioning on each server blade.
- Configure your Web server. If you're using the Apache Web server, you'll find the `httpd.conf` file in `/etc/apache` on the server.
- Test to be sure that the Web server responds to requests.

Configure Solaris OS and Web Server

Follow these steps if you want to test out your load balancing configuration without using VLANs. If this is the first time you've set up such a configuration, this is a good incremental step to take. For each server:

- Set up `/etc/hosts` to include entries for the addresses listed in the "Network Address Checklist" on page 5.
- Set up the files:


```

/etc/hostname.ce0
/etc/hostname.ce0:1
/etc/hostname.ce0:2
/etc/hostname.lo0:1
            
```

 to contain the host names (or simply the IP addresses) you've chosen for the three networks and the VIP, respectively.
- Add the three non-loopback interface names to `/etc/opt/SUNWclb/clb.conf`, for example:


```

ce0
ce0:1
ce0:2
            
```
- Create a boot script to create a default route to the data network with zero cost through the private address you've chosen. For the first of the two server blades illustrated in Figure 8, you'd create a file:

```
/etc/rc2.d/S99defaultroute
to contain the command:
route add default 10.0.0.0 0
```

- Delete `/etc/defaultrouter`. If you need to have a route to hosts other than over the data network, you'll need to manually create a route using additional commands in the script created in the previous step.
- Finally, change the address on which your Web server listens to the VIP address.
- Reboot the system and watch for error messages during the reboot process. You'll see errors from `sendmail` if a name server cannot be found.

Configure the Load Balancing Blade

With the Solaris OS and Web server working and ready to interoperate with the Sun Fire B10n Content Load Balancing Blade, the next step is to configure one of the two load balancers. This set of steps will get your network up and running without VLANs.

- Give the load balancing blade a management network IP address for talking to the CLB module on each server and for upgrading firmware. Go into `config` mode and clear out any old IP address with `no ip interface 0`. If you were setting up an interface following the network diagram in Figure 8, you'd type:


```
ip interface 0 192.168.0.1
mask 255.255.255.0
```

 Note that these

It Just Stopped Working!

Some customers encounter this problem after upgrading the firmware on their Sun Fire B10n Content Load Balancing Blade.

What's the problem? The CLB module and load balancing firmware all must be in sync to operate properly. If you upgrade one of these software components, be sure to upgrade both at once and your load-balanced services will operate without a problem.

Configuring The Service Group (no VLANs)

Configure the VIP:

```
config service name svcl4 vip 192.207.20.19:80:tcp interface 0
```

Add the first server (Weight 5 in a weighted round-robin scheme):

```
config service lb-group server svcl4:default server
192.168.0.2:80:tcp:5:1
```

Add the second server (Weight 1 in a weighted round-robin scheme):

```
config service lb-group server svcl4:default server
192.168.0.3:80:tcp:1:1
```

Examine the group that you've set up:

```
show service-lb-group svcl4 default
```

Enable the service group:

```
config enable service name svcl4
```

Commit changes:

```
commit
```

commands only affect the interface connected to SSC 0.

- Make sure that you're running the current version of the content load balancer firmware. Check the software download site listed in "What You'll Need" on page 5, and check the current version of firmware on the load balancing blade using the `show system` command. If necessary, update your load balancing and BSC firmware following the instructions in the *Product Notes*. To download firmware, you'll need to have configured an IP address that can be used to reach a TFTP server.
- Execute the commands illustrated in the "Configuring The Service Group (no VLANs)" sidebar above. You'll first set up the VIP that clients use to access the service and the service group name (`svcl4`). Next you'll add both of the server blades to the service group. The default load balancing scheme is Layer 4 weighted round-robin; the example gives one server a weight of 5, the other a weight of 1. Then you enable the service group and commit the changes.
- Use the `show service` command to see whether any services are currently defined on the load balancing blade. If any are defined (that shouldn't be), remove them using the command:


```
remove service name name
```

- With the management network set up, now be sure that you can ping the two servers, for example `ping 192.168.0.2`. You can monitor the count of heartbeat interactions between the load balancer and the server by executing the command: `kstat clbmod` on a Solaris OS system. Observe that the value of the `hbeat` variable is increasing by executing the command again.
- Now you should be able to ping the VIP from the data network, and you should be able to access the site from a Web browser. You can access several pages from the Web site and examine the Web server logs to verify that both servers are accessed by the load balancing blade.

Your load balanced service now runs using a single load balancing blade and no VLANs. If you don't want to use VLANs, you can skip to "Set Up Load Balancing Blade Failover" on page 9 to configure your second load balancing blade for failover operation.

Set Up VLANs

Configuring VLANs separates the three types of traffic — data, management, and service — into three logical networks. It's a good practice to use VLANs in your configuration because:

- You separate traffic into separate networks that are not accessible from the data network, helping to further secure your network from intrusion
- You can secure your management network using the SSC's built-in packet filter, giving yourself a secure back-door to both the server and the load balancing blades. (See the *Chassis Software Setup Guide* for details)
- You'll especially want to use VLANs to force a separation between plaintext and encrypted traffic if you plan to add Sun Fire B10p SSL Proxy Blades to your configuration.

Adding VLANs to your blade platform involves changing your operating system, switch, and load balancing blade configurations. You need to change your Solaris OS configuration so that its network interfaces pass traffic back and forth on the correct VLANs. You need to change your load balancing blade to add VLAN tags to its network traffic. And most importantly you need to change your switch configuration to enforce VLAN tagging.

Configure VLANs in the Solaris OS

Changing your Solaris OS network configuration to support VLANs is very straightforward.

- Change the names of the three hostname files for the data, management, and service networks to include the each network's VLAN number in the file name. You do this by adding the VLAN number times 1000 to the name. For example: the hostname file for the management network VLAN (2) using interface ce0 would be:

```
/etc/hostname.ce2000.
```

If the ce1 interface is used instead, the file name would be:

```
/etc/hostname.ce2001.
```

When your configuration is complete, you should have file names specifying three host names (or addresses) on VLANs, and one loopback address. The only file name not specifying a VLAN in its name is the

loopback address, for example

```
/etc/hostname.lo0:1
```

- Edit `/etc/opt/SUNWclb/clb.conf` to contain the three interfaces you've configured, namely:


```
ce21000
ce2000
ce23000
```
- If for any reason you've changed the VIP, be sure to change your Web server configuration to reflect the new address.
- Reboot the system and check that the network interfaces come up as expected by issuing the `ifconfig -a` command. Issue a `netstat -rn` command and carefully examine the routes, making sure that the default route is set up as you expect. You can double-check that the CLB module is configured properly with the command:

```
/etc/opt/SUNWclb/bin/clbconfig
list
```

Until you have the load-balancing blade and the two servers configured to use VLANs, your Solaris OS servers will be accessible only from the serial console accessible through the SSC.

Configure Switch VLANs

Before you change your switch configuration, it's a good idea to save your existing configuration so that you have a known working configuration to fall back to if necessary. Create a writable, empty file on the TFTP server, and then use the following command sequence at the switch console:

```
copy running-config file
copy file tftp.
```

The commands prompt you for file and TFTP server names.

Example commands for setting up VLANs on the switch are included in "Setting Up Switch VLANs" on page 9. They illustrate the following set-up steps:

- Create VLAN database names for the networks that you're going to set up (VLAN 2 is set up by default).

- Configure one of the external switch ports (for example NETP0) to carry data traffic to clients on the Internet. The example allows un-tagged traffic to enter the external switch port and be tagged upon entry. (If you want management traffic to exit through another physical port, you can use this set of commands as a model to accomplish this.)
- Configure the internal ports interfacing to the server and load balancing blades as endpoints for the three VLANs.
- Verify that the VLANs have been created properly by looking at the output from the `show interfaces switchport` command.
- Copy the running configuration to the startup configuration, and make a backup of your new configuration on the TFTP server.

Configure VLANs in the Load Balancing Blade

The last step is to configure the load balancing blade to interact with the three networks using the VLANs you've created on the switch.

- For each of the three networks, you'll specify a VLAN number and then enable its VLAN. A command sequence to implement the example configuration is:


```
management vlan 2
enable vlan management
data vlan 21
enable vlan data
service vlan svcl4 vlan 23
enable service vlan svcl4
```
- Once you've configured the VLANs, use the `show vlan` command to double-check the VLAN assignments
- Be sure to `commit` to save the current configuration.

Now you can test the configuration you've created by accessing the site hosted through the VIP. If you have problems, use the suggestions in "Troubleshooting Tips & Techniques."

Setting Up Switch VLANs

Add new VLAN names to the VLAN database:

```
vlan database
vlan 21 name external media ethernet
vlan 23 name loadgrp1 media ethernet
exit
```

Set up an external switch port to carry data network traffic in/out of the blade platform:

```
interface ethernet netp0
no switchport gvrp
switchport allowed vlan add 21 untagged
switchport native vlan 21
switchport allowed vlan remove 1
switchport ingress-filtering
exit
```

Set up internal ports interfacing to server and load balancing blades:

```
interface ethernet snp0
switchport gvrp
switchport allowed vlan add 21 tagged
switchport allowed vlan add 2 tagged
switchport allowed vlan add 23 tagged
switchport native vlan 21
switchport allowed vlan remove 1
switchport ingress-filtering
exit
```

Remember to save the current configuration:

```
copy running-config startup-config
```

Set Up Load Balancing Blade Failover

Now that you have the load balancing configuration up and running using VLANs, or if you've decided to run without VLANs, you can configure the second load balancing blade to take over in case the first one fails.

Under normal operation, one load balancing blade handles all traffic while the other one is in a passive state. The passive load balancer monitors the health of the active one. It also retrieves configuration and state information from the active load balancing blade so that it is prepared to take over if the active blade fails. The pair of load balancers maintain the same *configuration* state, but do not maintain *connection* state of the services they are supporting. If a failover occurs, existing connections to the servers are lost and must be re-tried.

More details on setting up failover are in the *Administration Guide*. The procedure is:

- Enter config mode and set up the VLANs on the passive (second) load balancing blade as discussed in “Configure VLANs in the Load Balancing Blade” on page 8.
- Give the passive load balancing blade an IP address on the management network, for example:


```
ip interface 0 192.168.0.4
mask 255.255.255.0 0.0.0.0
```
- Tell each load balancer who their peer is, for example use:


```
failover peer 192.168.0.1 0.0.0.0
```

 on the second load balancing blade.
- On both load balancers, enable failover:


```
enable failover-monitor
```
- Then start failover service on both blades.


```
On the active load balancing blade:
failover start local
On the passive load balancing blade:
failover start remote
```
- Commit the changes on both blades.

Now you can test the failover configuration. From the SSC console, power off the active load balancing blade and watch the console of the passive blade. You'll see the passive blade discover that the active one is unreachable and then take over the active blade's activities. Use the `show failover` command to examine the failover state of each blade.

Troubleshooting Tips & Techniques

We've suggested taking an incremental approach to setting up load balanced services on your Sun Fire B1600 Blade Platform for the first time. An incremental approach gives you the chance to test your progress along the way. If your configuration doesn't work at some point during the set-up process, use these suggestions to help pinpoint the problem.

Overall Tips and Techniques

The two most important tips are to:

- Be sure that you install the current version of load balancing firmware, BSC firmware, and operating system drivers, and keep the versions in sync.
- Don't forget to commit changes you make to each load balancing blade and to the switch. It's a good idea to save backup copies on a TFTP server.

Load Balancing Blade Troubleshooting

If you can't access the load balanced service that you've set up:

- Try to ping the VIP.
- Try to ping each server's management and service network addresses. Failures here are often indicative of a bad VLAN configuration on the load balancing blade, switch, or server.
- Examine the VLAN configuration on the load balancing blade with the `show config` command. Be sure that you're running the right configuration. Verify that the VIP is entered, and that the servers in the service

Load Balancing the Mars Rover Site

When users began jamming the NASA site, Sun teamed with the San Diego Supercomputer Center (SDSC) to build a mirror site to handle the overwhelming traffic created by users downloading a public version of NASA's Maestro software. Maestro is the command-and-control software used to operate the Mars rovers.

SDSC uses a single Sun Fire B1600 Blade Platform with eight gigabit Ethernet links to the Internet. The site uses a pair of Sun Fire B10n Content Load Balancing Blades to handle the incoming workload. The front-end Web servers are set up on six server blades: three Sun Fire B100s Server Blades running the Solaris OS and three B100x Server Blades running Linux. The site uses an additional three B100s and three B100x server blades for back-end functions and external storage on a network file server.

In the near future SDSC plans to use Sun Fire B10p SSL Proxy Blades to secure certain communications between clients and the site.

For more information, please refer to www.telascience.org.

group are also correct. You don't need to set gateway and DNS addresses.

Server Blade Troubleshooting

There are several steps you can take to debug problems with the server configuration:

- You should be able to ping the load balancing blade's management network addresses. If you can't you should suspect a VLAN configuration problem.
- Test the heartbeat that the load balancing blade maintains over the management network. Use the `kstat clbmod` command and watch the value of `hbeat` increase when you execute the command repeatedly.
- Use `ifconfig -a` to check consistency of your network numbering and VLAN configuration.
- Make sure that the CLB module knows about the interfaces you are using for the management, services, and data networks:


```
/opt/SUNWclb/bin/clbconfig list
```
- Likewise, make sure that each network interface has the CLB module inserted into its stack. For each interface name:


```
ifconfig interface modlist
```

Switch Troubleshooting

If you don't have connectivity between the load balancing blade and each server blade, use the command:

```
show interfaces switchport
```

and examine the output for each switch port and its VLAN configuration.

Conclusion

This issue of the *Sun Fire Blade Application Journal* has focused on setting up a single load balanced service group with a failover pair of

Sun Fire B10n Content Load Balancing Blades. Now that you have set up a basic configuration, you can start to explore the rich set of capabilities the load balancing blade and the blade platform have to offer:

- Explore the details of the failover mechanism by reviewing the tutorial in the *Administration Guide*.
- Configure a Layer 4 or Layer 7 classification scheme that best suits your application.
- Set up multiple service groups, and extend the services networks across multiple blade platforms and to additional servers if necessary.
- Use this issue as a guide for setting up load balancing across single- and dual-processor x86 architecture Sun Fire B100x and Sun Fire B200x Server Blades.
- Integrate the blade platform's management network with your datacenter's management network so that you can monitor and upgrade software on all of your service components.

Acknowledgements

Thanks to Steve Gaede and Bob Gray (Lone Eagle Systems) for researching and preparing this issue, and to Santwona Behera, Jochen Behrens, Robert Cibrario, Mandar Dange, Tom Giles, Dave Killian, David Lawler, Tim Mooney, Fred Smith, and to Muppalla Sridhar for their input, assistance, and comments.