



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Many Hands Make Light Work

Firewall Load Balancing

Brett Hose

September 7, 2002

Version 2

Introduction

The intention of this paper is to discuss the topic of firewall load balancing appliances.

I will discuss basically how the appliances manage load balancing and some general design considerations. I will also briefly touch on several of the products available in the market today and their features.

Finally I will focus on one specific firewall load balancing product (Cisco CSS Content Switch) and discuss in detail its implementation. The design covered utilises dual Checkpoint Firewall's running on Nokia appliance's in a fully redundancy configuration.

Appliance Load Balancing

There are several ways in which to load balance traffic across two or more firewall devices. These include; statically balancing traffic through the use of IP routing with VRRP (Virtual Router Redundancy Protocol), Clustering software that runs on the firewall and appliance load balancing, as discussed in this paper. For some excellent information on the other load balancing methods the SANS paper "[The New Firewall Design Question](#)" is well worth reading.

Application load balancers are typically used where there is a requirement to balance network traffic across several devices. These appliances are typically layer 4 to 7 aware network switches capable of load balancing traffic flows across Web servers, Web Caches, VPN (Virtual Private Network) Concentrators, IDS's (Intrusion Detection System), Firewalls etc.

Traffic is broken into flows and load balanced across each of the firewalls based on pre-configured balancing parameters. Typically the appliance will also track the state of each firewall, such that if one firewall was to fail, this firewall would be removed from the load balance group and as such no more traffic would be directed to it.

Load balancing appliances whilst costly can provide additional services within your network. As well as firewall load balancing the same device/s can offer load balancing for any of the servers/services mentioned above.

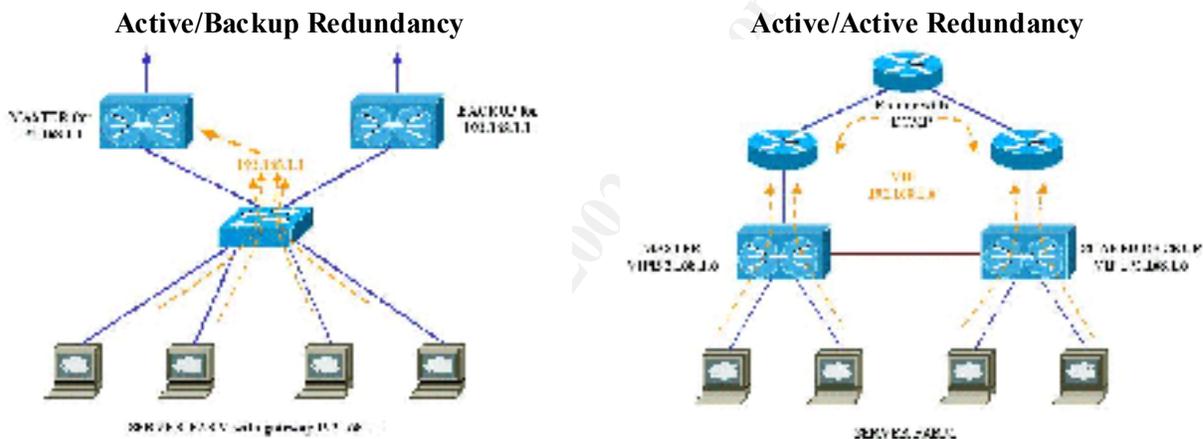
One other benefit of hardware load balancers is their ability to understand the higher layer protocols traversing the device. For example they have the ability to identify specific URL from within web traffic. Through this ability it may be possible to identify undesirable web traffic, say from a CodeRed attack, and block this traffic before it reaches the destination server.

Design Considerations

When used in a firewall load balancing environment, an appliance must exist within the traffic flow into the firewall's, and a separate appliance within the flow exiting the firewalls. The reason for this is due to the requirement for all traffic that enters via a firewall must return via the same firewall.

If a fully redundant configuration is required dual appliances will need to be deployed either side of the firewalls. All appliances offer an Active/Backup mode of redundancy and most vendors offer several ways in which to achieve an Active/Active mode. It is very important to investigate the network redundancy design requirements of the product being used as this may change from product to product.

For example the Cisco CSS content switch if used in an Active/Active, which is referred to as "VIP (Virtual IP Address) VIP redundancy", then there cannot be any layer 2 switches between the border routers and the CSS devices or between the servers and the CSS devices. Refer to the URL listed in the references section labelled "Understanding and Configuring VIP and Interface Redundancy on the CSS 11000" for more information.



A dedicated network segment is required between each of the appliances that are operating in a redundant pair. This network is used for passing appliance-to-appliance heartbeat information as well as flow state information (if supported). This network can typically be achieved via the use of a network crossover cable.

Firewall state synchronization is a process by which each firewall maintains flow state information for every flow through every firewall. The implementation of state synchronization between the firewalls and appliances will allow for device failures to be transparent to the end user. While this would appear as a must, some consideration for the performance implications on the firewalls must be given. State synchronization can seriously reduce a firewalls ability to pass traffic.

If state synchronization is not used, failure of a firewall will cause all sessions via that firewall, to drop and therefore need to be re-established. Whilst this may seem unacceptable, if the applications being used are tolerant to session failures (for example Web) then a configuration without state synchronization may be a better option.

Some appliances require a communication path between the frontend and backend redundant pairs. This would involve a firewall rule allowing this traffic.

An asynchronous traffic flow entering via one firewall and leaving via another must be avoided, as this will cause a heavy load on the firewall synchronisation process. If flow state information is not updated by the time the return packet arrives at the firewall it may be dropped even though it is valid traffic. Some older versions of Checkpoint will allow the flow and await the state update (could be perceived as a security risk) while later versions will drop it.

Most of the appliances available today do support NAT (Network Address Translation) being performed on the firewalls. There are a few that do not support this and in this situation NAT would be performed on the appliance device or the border router.

How they work

Traffic is basically broken into individual flows (For example a TCP session would constitute a flow). Based on a pre-configured balancing algorithm each flow is directed to one of the firewalls.

The basic functions of a firewall load balancer: -

- Monitoring firewall device status
- Mapping (Balancing) flows
- Manage firewall session information
- Maintaining state and session information between redundant boxes

The state of each firewall is periodically checked, typically via ping. If a firewall device fails then it is removed from the available device list. If firewall state synchronisation is implemented then flows associated with the failed firewall will continue to be serviced by another firewall.

In a fully redundant load balancing configuration, status of the peer appliance will be monitored. Some load balancing appliances support stateful failover and as such any flow associated with the failed appliance will continue to be serviced on the surviving appliance. If stateful failover is not supported then a flow must be re-initiated and may result in a disruption to the user.

Load balancers deal with traffic in the form of flows. Each new flow that enters the device has its source and destination IP addresses hashed. Based on the hash value it is forwarded through a specific firewall. While most load balancers use this balancing method as the default, some offer more sophisticated methods (eg Least Connections, Response Time, Round Robin etc).

The returning traffic flows are checked against the session information table and forwarded back through the same firewall that it entered through.

Load Balancing Products

Cisco CSS

The Cisco CSS Content Switch is available in a 2RU fixed configuration or an 8 slot chassis based configuration. Each configuration offers various ports densities providing 10/100/1000 Mbps Ethernet connectivity.

- NAT performed on the Firewall is not supported. Can be done by the CSS
- Load balancing is based on hash of source and destination IP only.
- Device redundancy can be provided in an Active -Backup or Active-Active mode.
- Device management is provided via console, telnet, SSH or Web GUI.
- Supports stateful failover at layer 5 only.

Alteon

The Alteon Stackable Web Switch is one of the content delivery devices offered by Nortel Networks. The Alteon 180 is a 3 RU rackable device that comes in various configurations of 8 or 9 ports, providing 10/100/1000 Mbps Ethernet connectivity.

- Device redundancy can be provided in an Active -Standby or Active-Active mode through the use of VRRP.
- Supports stateful failover (In Active -Standby mode only)
- Device management is provided via console, Web GUI, or telnet.
- NAT performed on the firewall is supported
- Supports true load balancing based on leastconn, response, roundrobin etc

F5 BIG-IP

BIG-IP FireGuard 520 is a firewall specific load balancer offered by F5. It is provided in a 2RU chassis with 2 10/100 Mbps Ethernet ports and optional GB ports.

- Supports a large variety of load balancing methods - Round Robin, Ratio, Least Connections, Fastest, Fastest-connect, Observed.
- Supports stateful failover
- Device redundancy can be provided in an Active-Standby or Active-Active mode
- Device management is provided via console, Web GUI, SSH, or telnet.

Radware

The Radware FireProof is a 1RU rack mount unit that is provided in either a 16 or 7 port configuration with 10/100/1000 Mbps Ethernet ports.

- Supports stateful failover
- Device management is provided via console, Web GUI, SSH, or telnet.
- Supports a large variety of load balancing methods - Cyclic Least number of users, Least amount of packets, Least amount of bytes, Windows NT algorithm.
- Device redundancy can be provided in an Active -Standby or Active-Active mode

Cisco CSS Content Switch Design

Description

Each CSS pair work in an active/backup mode in which the active device manages all traffic and the backup device is essentially down, from an IP perspective. The active device is configured with a VIP (Virtual IP address) for each of the VLAN's configured. The backup device monitors the state of the active device. Upon failure of the active device, the backup device will move to an active state and

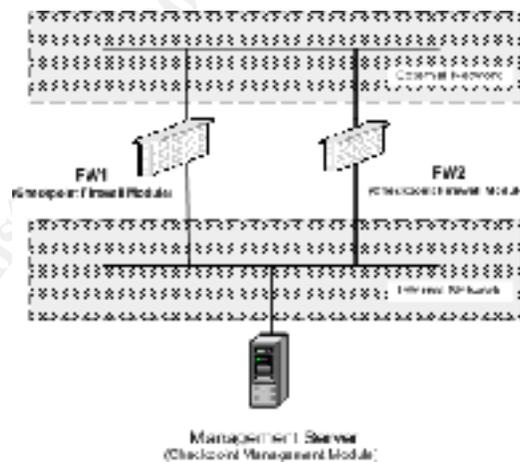
assume responsibility for the all of the VIP's configured. The configuration of the backup device is the same as the active device except for several configuration items associated with the redundancy network.

Frontend CSS: In this design there is a separate VLAN (Virtual LAN) per border router and another VLAN for the outside firewall network. Within the CSS configuration each VLAN will have a single Virtual IP address. The Virtual IP address for each border router VLAN (10.1.1.2 and 10.1.2.2) is the next hop IP address for traffic entering this environment from the border routers.

The Virtual IP address for the outside firewall network (10.1.4.3) will be the next hop IP address for traffic exiting this environment. This address is used by the firewalls as the default route. The CSS is configured with two static routes to direct Internet traffic up via either of the border routers (10.1.1.1 and 10.1.2.1).

The CSS devices have a concept called flows. A flow defines the path a packet will take from the frontend CSS's through one of the firewalls to the backend CSS's. There will be a flow defined for each of the paths that traffic can take through the network (in this case two flows are defined). This flow definition must be the same but reversed on both CSS pairs (frontend and backend). Routing of traffic across the firewalls is performed using static routes that point to these flow definitions.

Firewalls: The two Nokia/Checkpoint firewalls are running state synchronization across a dedicated VLAN (10.1.5.0). The Checkpoint software used here is in a distributed configuration. Each of the firewalls is running the Checkpoint firewall module only. The Checkpoint management module is running on a separate Windows system within the Internet network. Firewall rules are created on the Checkpoint management system and then pushed to both firewalls. The version of Checkpoint used in this configuration was 4.1.



Backend CSS: These CSS devices not only perform firewall load balancing, but also provide load balancing for all web traffic across the four Web servers.

These backend CSS pairs are running in the same redundancy mode as the frontend pair. They have been configured with two VLAN's, one for the inside firewall network and the other for the Web complex network.

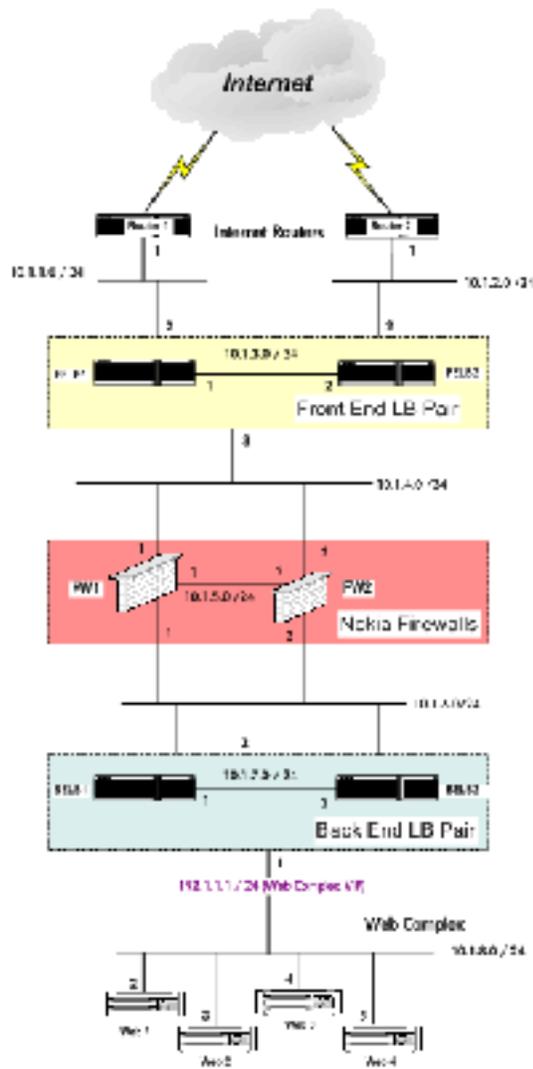
The VIP (Virtual IP address) for the inside firewall network is used by the firewalls as the next hop for traffic entering the environment. The Virtual IP address of the Web complex network (10.1.8.1) is used by the Web servers as the next hop for their default route.

In order to provide load balancing of web traffic across the four web servers, a Virtual IP address is defined (192.1.1.1) within the backend CSS's and advertised to the outside world. Any traffic destined to this address will hit the backend CSS's and be load balanced across the four Web servers. This VIP and the servers included in the load balancing are configured through the owner and content commands (see configuration below). The state of each Web server is monitored via periodic connections to port 80. This is configured through the use of the service command.

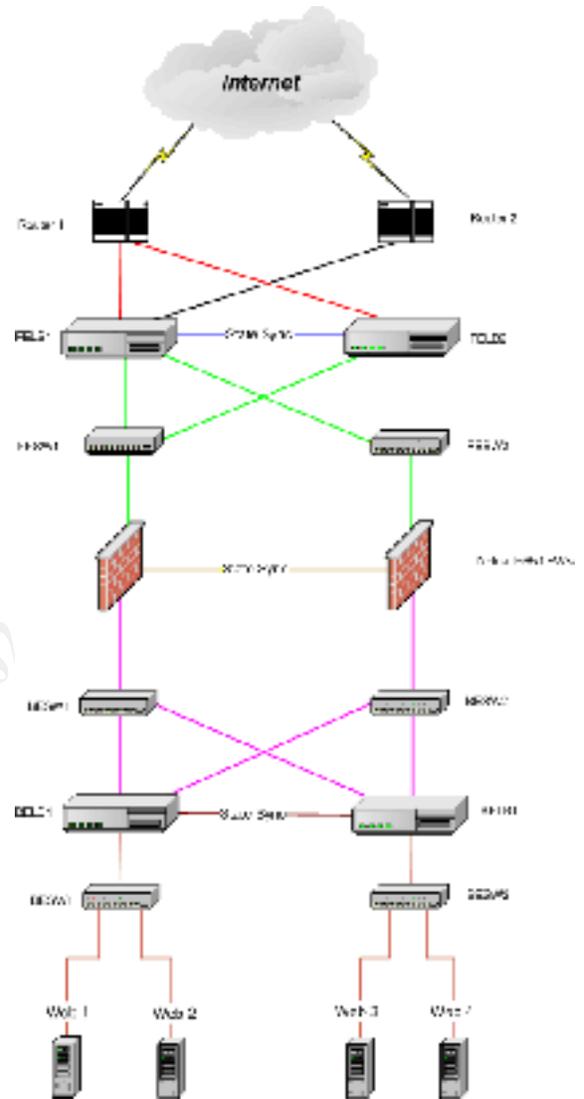
© SANS Institute 2000 - 2002, Author retains full rights.

Design Diagram

Layer 3 Design



Layer 2 Design



Design Considerations

Redundancy: The CSS can support redundancy in several modes either box to box or VIP. Box to box redundancy was chosen as this provided for the easiest implementation and allowed for the establishment of layer 2 redundancy, which was desirable. Box to box is an active/backup type of redundancy.

Spanning Tree: Spanning tree protocol is used to prevent loops within a switched environment. As we have elected to implement layer 2 redundancy, some consideration needed to be given the spanning tree protocol configuration. The configuration discussed later in this paper does not cover any spanning tree configuration as this was performed within the layer 2 switches. Each of the layer 2 switches was configured such that they were unable to become the spanning tree root bridge. In doing this we ensured that the active CSS device was always the root bridge.

State and Redundancy: State synchronization for the firewalls should be established over a dedicated network. Whilst this is not a requirement of the Nokia/Checkpoint firewalls it is certainly a best practice. State synchronization and redundancy for the CSS's must run over a dedicated network.

Border router connection: Avoid designing a network using very high speed links (Gigabit) into the CSS devices from the border routers. In very high volume traffic environments it is fairly important to provide a point within your network where traffic throttling can take place. The location for this function is logically the border routers. High speed border router to CSS links will cause the throttling point to become the firewalls which is not desirable.

NAT: The CSS appliance do not support NAT (Network Address Translation) being performed on the firewalls.

Management: In order to protect access to any of the network components, only SSH access should be permitted. Some of the devices support a Web management interface, which should be disabled.

© SANS Institute 2000 - 2002

Frontend CSS Configuration - FELB1

Redundancy

The following commands setup the type of redundancy being used. The first command turns on redundancy and the app commands are used to establish the redundancy protocol between the active and backup CSS. Note that the app session command is specifying the IP address of the peer device's redundancy interface.

```
ip redundancy
app
app session 10.1.3.2
```

Firewalls Flows

These two commands define the firewall flows as discussed earlier. The configuration below defines flow 1 as being via the first firewall and flow 2 via the second firewall. The IP addresses defined are each of the IP interfaces that traffic would flow through as it traverses the firewalls. For example flow 1 would traverse; Firewall ones outside interface IP, Firewall ones inside interface IP and the CSS's inside firewall network Virtual IP address.

```
ip firewall 1 10.1.4.1 10.1.6.1 10.1.6.3
ip firewall 2 10.1.4.2 10.1.6.2 10.1.6.3
```

IP Routing

The first two routes are for traffic exiting the environment via either of the two border router VLAN's. The second set of routes is for all traffic entering the environment, whose destination is the Web server complex. Note that each of these routes has a gateway definition of firewall 1 or 2. This is indicating that this traffic will use the flows defined earlier as the paths into the environment. The value on the end of the command is the route weighting. As they are the same, traffic will be load balanced across these flows.

```
ip route 0.0.0.0 0.0.0.0 10.1.1.1 1
ip route 0.0.0.0 0.0.0.0 10.1.2.1 1
ip route 192.1.1.0 255.255.255.0 firewall 1 1
ip route 192.1.1.0 255.255.255.0 firewall 2 1
```

Physical Interface configuration (Layer 2)

Each of the physical interfaces is configured for the appropriate speed and duplex settings and placed into the appropriate VLAN with the bridge command.

```
interface ethernet-1
phy 100Mbps-FD
bridge vlan 13
```

```
interface ethernet-2
```

```
phy 100Mbps-FD
bridge vlan 13
```

```
interface ethernet-12
phy 100Mbps-FD
bridge vlan 14
```

```
interface ethernet-13
bridge vlan 11
```

```
interface ethernet-14
bridge vlan 12
```

Virtual IP address configuration (Layer 3)

The following configuration items are for the IP interface parameters. The IP addresses defined are the VIP (Virtual IP) addresses for each of the VLAN's. The redundancy definition indicates that this interface will participate in the redundancy operation. If this option is omitted from the circuit configuration, this interface will be active regardless of the devices state (i.e. Active or Backup)

The redundancy-protocol command under circuit VLAN14 indicates that the redundancy protocol will be running on this interface. This is the interface allocated to the dedicated redundancy network.

```
circuit VLAN13
description " External Firewall Segment "
redundancy
ip address 10.1.4.3 255.255.255.0
```

```
circuit VLAN14
description "CSS Redundancy Segment"
ip address 10.1.3.1 255.255.255.0
redundancy-protocol
```

```
circuit VLAN11
description "Router 1 external Segment "
redundancy
ip address 10.1.1.2 255.255.255.0
```

```
circuit VLAN12
description " Router 2 external Segment "
redundancy
ip address 10.1.2.2 255.255.255.0
```

Frontend CSS Configuration – FELB2

Redundancy

The following commands are the only differences between FELB1 and FELB2. Once these configuration items have been completed a command to update the configuration needs to be run from the **master** CSS, in this case FELB1. The following command will perform this update; script play commit_redundancy –a 10.1.3.2

By checking the configuration on FELB2 with the show command you will see that it has the same configuration as FELB1 except for the commands below.

As per the description above the app session command references the IP address of the redundancy network IP address on FELB1

```
app
app session 10.1.3.1
```

```
circuit VLAN14
description "CSS Redundancy Segment"
ip address 10.1.3.2 255.255.255.0
redundancy-protocol
```

Nokia/Checkpoint Firewalls Configuration

Static Routes

The default static route will have a next hop of the outside firewall network Virtual IP address as advertised by the active frontend CSS device. The second static route is for all traffic destined for the Web server complex network VIP address. This should constitute the majority of traffic entering this environment.

Network	Next Hops	Description
Default	10.1.4.3	Frontend CSS Pair
191.1.1.0/24	10.1.6.3	Web Complex VIP

Firewall Policy

The following two tables are purely examples of how the firewalls could be configured. Obviously this will be environment specific.

Rule	Source	Destination	Services	Action	Track	Comments
1	Admin	Firewall_A Firewall_B	FireWall1 SSH	accept	Long	Allow access to Firewall for Administrators
2	mgtsvr	Firewall_A Firewall_B	Firewall1	accept	Long	Allow access to Firewalls from CP management server
2	Any	Firewall_A Firewall_B	Any	drop	Long	Block any external access to Firewall
3	Any	Web Servers	http https	accept	-	Allow Web access Web Complex VIP Address
4	Web Servers	Any	domain- udp	accept	-	Allow Web Servers to perform external DNS lookup
5	Any	Any	Any	drop	Long	Deny ANY ANY

Network Objects

Name	Type	IPAddress	Netmask	Members	Comment
Firewall A	Host	10.1.4.1			
Firewall B	Host	10.1.4.2			
mgtsvr	Host	10.1.8.100			
Web Servers	Network	192.1.1.0	255.255.255.0	-	Web Servers
Admin	Group	-	-	admim1 admin2	Trusted Administrators

Backend CSS Configuration - BELB1

Redundancy

The app session command will establish a redundancy protocol session with BELB2 on the redundancy network.

```
ip redundancy
app
app session 10.1.7.2
```

Firewalls Flows

These flow definitions are the same as configured above (refer to FELB1 configuration) but in the reverse direction. For example firewall 1; inside firewall network IP address, outside firewall network IP address, outside firewall network CSS VIP address.

```
ip firewall 1 10.1.5.1 10.1.4.1 10.1.4.3
ip firewall 2 10.1.5.2 10.1.4.2 10.1.4.3
```

IP Routing

These routes are the default routes for traffic exiting the environment. Traffic will be load balanced across the flow definitions above.

```
ip route 0.0.0.0 0.0.0.0 firewall 1 1
ip route 0.0.0.0 0.0.0.0 firewall 2 1
```

Physical Interface configuration (Layer 2)

```
interface ethernet-3
phy 100Mbps-FD
bridge vlan 16
```

```
interface ethernet-4
phy 100Mbps-FD
bridge vlan 16
```

```
interface ethernet-12
phy 100Mbps-FD
bridge vlan 17
```

```
interface ethernet-1
phy 100Mbps-FD
bridge vlan 18
```

```
interface ethernet-2
phy 100Mbps-FD
bridge vlan 18
```

```
interface ethernet-12
phy 100Mbps-FD
bridge vlan 17
```

Virtual IP address configuration (Layer 3)

```
circuit VLAN16
description "Firewall Segment Internal"
ip address 10.1.6.3 255.255.255.0
redundancy
```

```
circuit VLAN17
description "CSS Internal VRRP Segment"
ip address 10.1.7.2 255.255.255.0
redundancy-protocol
```

```
circuit VLAN18
description "Web Complex Segment"
ip address 10.1.8.1 255.255.255.0
redundancy
```

Service Configuration

All of the configuration items from this point forward are related to the load balancing of web traffic across the four Web servers.

For each of the web servers there will be a service configuration. Within this configuration is the real IP address (10.1.8.X) of the Web servers and some keepalive parameters. The configuration used below specifies that each of the web servers is to be monitored via a TCP connection to port 80 periodically. Should this keepalive fail, the server will be removed from the active server list and no more traffic will be directed to it. The active command enables this service for use.

```
service web_server1
ip address 10.1.8.2
keepalive port 80
keepalive type tcp
active
```

```
service web_server2
ip address 10.1.8.3
keepalive port 80
keepalive type tcp
active
```

```
service web_server3
ip address 10.1.8.4
keepalive port 80
keepalive type tcp
active
```

```
service web_server4
ip address 10.1.8.5
keepalive port 80
keepalive type tcp
active
```

Owner – (content rule)

The owner command provides the ability to bind several content rules together. This configuration involves a single content rule within the owner www.

Under the content www command are the parameters for the configuration of the load balanced Web service. The load balancing mechanism used is round-robin and the Virtual IP address is defined through the vip address command. This VIP address is the advertised IP address that all traffic uses to access the Web server complex. Each of the previously mentioned services created for the web servers are added into the content rule with the add service command.

Finally we must configure what is called stickiness with the advance-balance command. Stickiness is the way in which we ensure that once traffic is directed to a specific server it continues to be directed to the same server. There are several more sophisticated forms of stickiness available but in this case

we have chosen one of the more basic forms. We are using stickiness based on a hash of the source IP address and the destination port number.

```
owner www
```

```
content www
balance round-robin
vip address 192.168.1.1
add service web_server1
add service web_server2
add service web_server3
add service web_server4
advanced-balance sticky-srcip-dstport
active
```

Backend CSS Configuration – BELB2

Redundancy

The following commands are the only differences between FELB1 and FELB2. Please refer to FELB2 configuration section.

Once again the configuration will need to be updated from the active to the backup CSS. This command must be run from the **active** CSS. The following command will perform this update; script `play commit_redundancy -a 10.1.7.2`

```
app
app session 10.1.7.1
```

```
circuit VLAN18
description "CSS Internal VRRP Segment"
ip address 10.1.7.2 255.255.255.0
redundancy-protocol
```

Observations

The following observations are not presented in any specific order. They are intended as purely a list of points observed during the build and testing of this environment.

Firewall Load: This configuration was deployed into a fairly large Internet data centre that was the host to a very popular web site at the time. During the early days of this site very high traffic loads were experienced. It was found that during these times the firewall state synchronization process caused significant load on the firewalls CPU (Central Processing Unit). It was later decided that the reduction in load on the firewall CPU would be more beneficial than transparent failover offered by firewall state synchronization. As the traffic into the site was all web based the effect of a failure to the end user was believed to be negligible.

Also as a result of the high traffic volumes it was found that increasing each of the firewalls memory capacity and the tuning of the connection table proved to be very beneficial.

Redundancy: If the redundancy link between the two CSS devices were to fail, then both the CSS's would try to assume the master state. This obviously caused some confusion within the network as the Virtual IP addresses for all VLAN's were advertised by both CSS's. The Cisco CSS can be deployed in other redundant modes, which may avoid this issue. The other redundancy modes would however require a different layer 2 network design as discussed earlier (removal of the L2 switch's).

Port Auto-Negotiation: There are many issues with auto -negotiation of the speed and duplex on 10/100 Ethernet ports. It is always a good idea to make sure that all 10/100 Ethernet ports are manually configured to the correct settings for duplex and speed. This avoids auto -negotiation issues that can be difficult to isolate.

In-Built Security: Like other load balancing devices the CSS has built -in basic DOS (Denial-of-Service) detection capabilities. DOS indication received from the CSS's, in our case, was found to be erroneous information.

Load Balancing Effectiveness: Even though the CSS devices use a hash of source and destination IP to balance traffic, the load distribution was found to be reasonably equal.

Server Load Balancing: If server load balancing is to be done then the basic mechanisms for load balancing and server health produced better results. For server health, ICMP was found to be the more reliable option and for load balancing round -robin.

Conclusion

For any secure Internet access environment where high volumes of traffic are experienced firewall load balancing is a valuable addition. As well as load balancing it also provides redundancy.

Load balancing appliances have been around for some time now and most of the major products are mature products with a complete range of features.

They offer a very cost effective solution when used for firewall load balancing in conjunction with load balancing over other services such as Web servers, Cache's or even IDS's. Most appliances available today provide a complete range of effective load balancing mechanisms. The benefit through offloading the load balancing function to a dedicated device cannot be ignored.

Whilst there are several ways in which firewall load balancing can be achieved appliances are a worthy consideration.

References

SANS: The New Firewall Design Question

URL: http://rr.sans.org/firewall/new_design.php

RadWare: FireProof Security Application Switching

URL: <http://www.radware.com/library/pdfs/products/FireProofSAS.pdf>

RadWare: High Availability Security Solutions

URL: <http://www.radware.com/library/whitepapers/highaval.pdf>

Network Computing: Web Server Director Comes Out on Top of the Pile

URL: <http://www.networkcomputing.com/1203/1203f1b1.html>

LBDigest.com: Have Your Layer Cake And Eat It Too

URL: <http://lbdigest.com/article.php?sid=48>

Cisco: Web Network Services White Paper

URL: http://www.cisco.com/warp/public/cc/pd/si/11000/prodlit/csecm_wi.htm

Cisco: Configuring Firewall Load Balancing

URL: <http://www.cisco.com/univercd/cc/td/doc/product/webscale/css/advcfggd/firewall.htm>

Cisco: Understanding and Configuring VIP and Interface Redundancy on the CSS 11000

URL: http://www.cisco.com/warp/public/117/vip_appguide.html

Cisco: Cisco CSS 11500 Series Data Sheet

URL: http://www.cisco.com/warp/public/cc/pd/cxsr/11500si/prodlit/css11_ds.htm

Cisco: Configuring Box-to-Box Redundancy on CSS 11000s

URL: http://www.cisco.com/warp/public/117/b_to_b_redund_app_guide.html

Alteon: Firewall Load Balancing

URL: http://www.nortelnetworks.com/products/library/collateral/intel_int/flb_wp.pdf

Alteon: Web OS Switch Software Application Guide

URL: http://www142.nortelnetworks.com/bvdoc/alteon/webos/webos10.0/212777_A.pdf

F5: High-Availability Load Balancing For Firewalls

URL: <http://www.f5.com/solutions/techbriefs/firewall.doc>

F5: BIG-IP® FireGuard 520 Product Brochure

URL: <http://www.bigip.com/f5products/bigip/fireguard/index.html#Specifications>