



CSM One-arm Design in the Data Center

This chapter describes the design and configuration of a secure and highly available data center with the Catalyst 6500 Content Switching Module (CSM) in one-arm mode. The data center design with the CSM can be deployed with or without a Cisco Catalyst 6500 Firewall Services Module (FWSM).

This chapter includes the following sections:

- [CSM Design Overview](#)
- [One-arm CSM Architectural Details](#)
- [Configuration Details](#)
- [Configuration Listings](#)

This chapter also provides design and implementation recommendations for using firewall and load balancers in a data center to provide security and load balancing services. These services are important for many types of servers, including web servers, application and database servers (typically used for running web-based transactional applications), and DMZ servers, including DNS servers and SMTP servers.

The FWSM and CSM can be deployed together in several modes, but the following are the two most important:

- FWSM in routed mode combined with the CSM in transparent mode—This design provides an easy-to-implement solution for multi-tier server farms.
- CSM in one-arm mode combined with the FWSM in transparent mode—This design is the topic of this design guide for the use with Supervisor 720.

Either design can be implemented with the Catalyst 6500 Supervisor 2 or with the Catalyst 6500 Supervisor 720. The latter design provides traffic optimization for connections that do not require any load balancing, and it provides increased performance at the cost of a slightly more complex configuration, requiring the use of policy-based routing (PBR).

CSM Design Overview

This section includes the following topics:

- [CSM One-arm Design](#)
- [Designs with FWSM and CSM](#)
- [One-Arm CSM Design with FWSM in Transparent Mode](#)
- [Hardware Requirements](#)

- DoS Protection


Deploying load balancing in a data center typically requires placing the load balancer in the main traffic path between client and servers. This can be achieved by configuring the CSM as the default gateway for the servers, or by placing a load balancer in transparent (Layer 2) mode between the servers and the default gateway.

A different design, called CSM one-arm design, consists in connecting the CSM with a single VLAN to the Multilayer Switch Feature Card (MSFC) with the MSFC providing the default gateway function. This design is the object of this document.

CSM One-arm Design

CSM one-arm design is useful for load-balanced servers requiring high-throughput server-to-server data transfers (such as back-up traffic) and with mainframes that require load balancing. The design in [Figure 5-1](#) provides server-to-server throughput equivalent to the maximum fabric performance of the Catalyst 6500 because no firewall or load balancer is in the path. This design is often used for mainframes connecting at Layer 3 to the Catalyst 6500.

Figure 5-1 Redundant CSM One-arm Design



Traffic that requires load balancing (represented with a continuous line in [Figure 5-1](#)) is directed to the MSFC, where it is intercepted by the Route Health Injection (RHI) route, which is installed dynamically by the CSM when virtual IP (VIP) addressing is active. The traffic then is directed to the CSM for the load balancing decision. The CSM performs the rewriting of the destination IP address to the server IP address and then sends the traffic to the MSFC in the Catalyst 6500 to be routed to the appropriate servers.

PBR is applied to the interfaces indicated with the colored circles in [Figure 5-1](#). An ACL classifies the traffic that needs to return to the CSM through PBR. For example, PBR intercepts return traffic for load-balanced servers and returns it to the CSM.

This CSM one-arm design also simplifies load balancing in a server farm environment with Layer 3 multi-homed servers or with mainframes. In either scenario, the servers or mainframes participate in Open Shortest Path First (OSPF) routing to advertise the IP address of the applications. Front-ending these servers or mainframes with a router is the best choice.

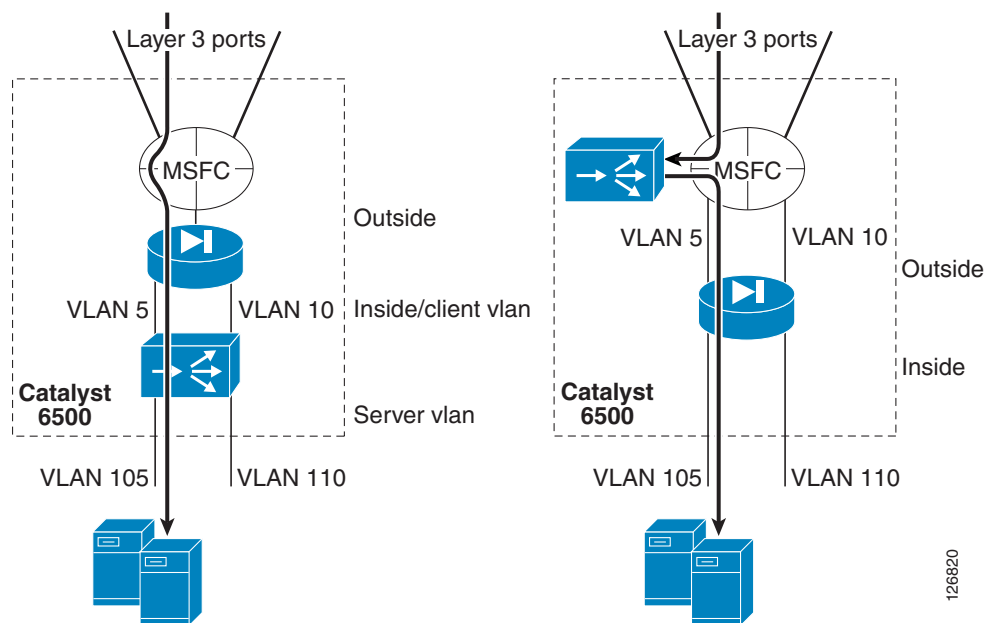
In [Figure 5-1](#), the mainframes connect to the Cisco Catalyst 6500 over Layer 3 links and they participate in OSPF with the router. This is possible when the CSM is deployed in routed or bridge mode. If an application instance on one mainframe needs to communicate with an instance on another mainframe, the traffic they generate is routed directly by the Cisco Catalyst 6500 without involving the CSM.

Designs with FWSM and CSM

The load balancing and firewalling configuration with FWSM and CSM can follow the following two main modes:

- Inline—CSM—MSFC—FWSM—CSM—servers (See [Figure 5-2](#) to the left)
- One-arm—MSFC—FWSM—servers + MSFC—CSM (See [Figure 5-2](#) to the right)

Figure 5-2 *Inline Design versus CSM One-Arm Mode with FWSM Transparent Mode*



The benefit of this design includes the fact that the denial of service (DoS) protection capabilities of the CSM and FWSM are combined, as follows:

- The CSM protects against DoS (SYN flood) attacks directed at the VIP.
- The FWSM protects against DoS (SYN flood) attacks directed at non-load balanced servers.



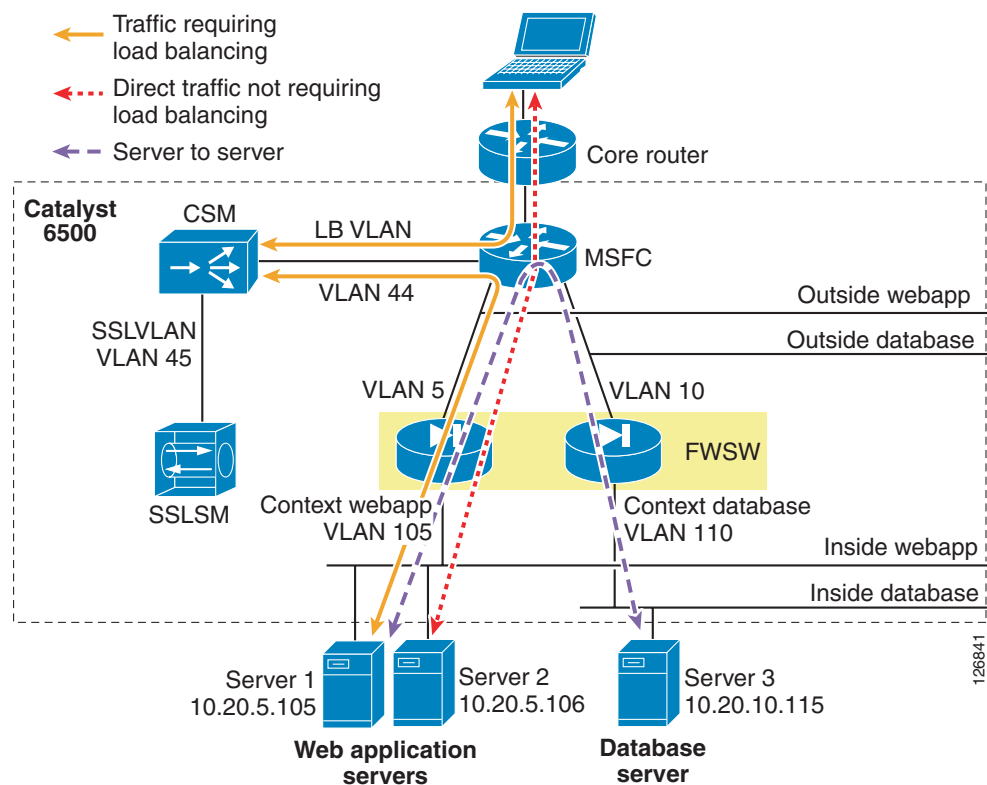
Note

The CSM line card operates in bus mode. When using the CSM in conjunction with the FWSM line card, Cisco recommends forcing the FWSM to operate in bus mode using the **fabric switching-mode force bus** command. When service modules such as the CSM and the FWSM operate in bus mode, traffic from DFC-enabled line cards still use the fabric connection.

One-Arm CSM Design with FWSM in Transparent Mode

Figure 5-3 illustrates the logical topology of the design presented in this chapter, and the VLANs and IP addresses used in the configurations.

Figure 5-3 Logical Topology without Redundancy



The firewall is virtualized in multiple contexts; two in this example protect respectively the presentation/application tier and the database tier.

The Catalyst 6500 is the blue rectangle that includes the FWSM, the CSM, and the Cisco Secure Socket Layer Service Module (SSLSM).

Traffic that requires load balancing (represented with a continuous line in Figure 5-3) hits the MSFC first; then it is intercepted by the RHI route; then it goes to the CSM for the load balancing decision.

The CSM performs the rewriting of the destination IP address to the server IP address and then sends the traffic to the MSFC in the Catalyst 6500 to be routed to the appropriate servers, wherever this server might be located; that is, the CSM can load balance across any application tier or firewall context (it is up to the firewall to prevent unwanted traffic from entering a given segment).

The traffic then enters the appropriate segment through a firewall instance. The firewall is operating in bridge mode; as such, the MSFC simply uses Address Resolution Protocol (ARP) to find the real IP address and then forwards the traffic through the firewall instance.

The return traffic takes the reverse path, and a PBR ACL is configured on the MSFC interface to push the traffic back to the CSM.

Traffic that does not require load balancing is forwarded directly to the servers. This traffic includes client-to-server traffic that is not subject to any load balancing rule on the CSM (dotted line in [Figure 5-3](#)) and server-to-server traffic (dashed line in [Figure 5-3](#)).

**Note**

Whether the service modules are physically in the same Catalyst 6500 or have been placed in a “service switch” is not relevant for the topic of this chapter. This chapter assumes that the CSM and the FWSM are in the same chassis, but it is equally applicable when the FWSM is placed in an aggregation switch and the CSM is placed on a “service switch”; that is, an external Catalyst 6500 used to provide mostly content functions such as load balancing, SSL offloading, and providing connectivity to reverse proxy caches.

Using the FWSM to segregate server farms is useful for servers that belong to different organizations, for applications to which you want to apply different filtering policies, or to tier web/application/database servers to make it more difficult for a hacker to access confidential information.

To segregate servers with different security levels, assign them to different VLANs, with each VLAN trunked to the FWSM and assigned to a different firewall context.

**Note**

Currently, each firewall context provides one outside interface and one inside interface.

The correct placement of the MSFC is a key element for the performance of this design. The traffic hitting the aggregation switches from the core should go to the MSFC first and the FWSM afterwards. This enables the use of Layer 3 links to connect the aggregation switches with the core and the assignment of the MSFC as the default gateway of the servers.

Hardware Requirements

The hardware required to implement the CSM one-arm design described in this guide is as follows:

- Cisco Catalyst 6500s with Supervisor 2
- Cisco Catalyst 6500s with Supervisor 720
- A pair of CSMs installed in the Cisco Catalyst 6500s


The principles described in this chapter also apply to Cisco Content Services Switch (CSS) family implementations with similar topologies.

DoS Protection

Since Release 3.2(1), the CSM protects against DoS attacks using the TCP SYN cookies technology. The CSM with SYN cookies can sustain a DS3 level of DoS attack with no visible impact to user HTTP transactions. The performance degradation is about 10 percent, which means that legitimate transactions


still complete, but the connection setup rate goes down. The performance degradation becomes significant (30–40 percent) at approximately 300,000 SYN/s of SYN flood (see [Figure 5-4](#).) At this point, HTTP transactions still complete, but the setup rate for legitimate transactions is reduced.

Figure 5-4 CSM Transactions per Second During a DoS Attack



The CSM one-arm design can be used to combine the CSM with the FWSM for DoS protection (see [Figure 5-5](#)).

Figure 5-5 CSM and FWSM Combined Protection



In this design, traffic directed to VIPs is intercepted by the CSM, and traffic directed to non-load balanced servers goes directly through the FWSM, bypassing the CSM. If an attacker launches a SYN flood against a VIP, the CSM is hit first, and if the attacker launches a SYN flood against an IP address that does not require any load balancing the FWSM sees the traffic first.

In this topology, the CSMs are one hop away from the servers, and it is essential that load-balanced traffic go through the routers before getting to the CSM. The CSM offers a virtual address as the next hop for the router. This address is called an *alias* on the CSM and is equivalent to an HSRP address on a router.

The horizontal arrows in [Figure 5-6](#) represent redundancy protocol messages, which are transmitted on specific VLANs between the peer routers (HSRP) or between the pair of CSMs (CSRP). These devices are configured to display a common IP address to their clients, which eliminates the single point of failure on the CSMs or the routers. If the master device fails, the backup takes over and is reachable with the same IP address.

Client-to-server traffic is pushed to the CSM with a static route on each MSFC using a VIP destination address that points to the CSM. You can configure a static route manually or the CSM can configure it dynamically based on server availability, using RHI.

Clients use the VIP as the IP address to connect to the services offered by the server farm. The CSM assigns a client connection, using a VIP as the destination address, to a specific server in the server farm (called real). The load balancing algorithms, known as predictors, are defined on the CSM and select a server to which the CSM sends the client request.

The servers send traffic back to the default gateway on the MSFC, which sends traffic to the CSM, using PBR. In this example, PBR is configured on VLAN 5 and 10, and the CSM is the PBR next hop. PBR pushes traffic from VLAN 5 and 10 to VLAN 44. The PBR next hop is on VLAN 44 even when PBR is applied to VLAN 5 or 10.

Policy-Based Routing

PBR supports traffic routing based on policies rather than the destination IP address. On the Cisco Catalyst 6500, PBR is implemented in hardware ASICs when used with Supervisor 2 or Supervisor 720. Examples of policies are the following:

- Incoming VLAN
- Source IP address
- Layer 4 protocol

You can apply PBR on a per-VLAN basis using a route map, which, like a routing table, defines the next hop for traffic that matches the policy. The next hop can be on a different VLAN from where the traffic originated.

Identifying Load-Balanced Servers

The following is an example of PBR configuration:

```
ip access-list extended return-traffic-http
 permit tcp any eq 80 any
 permit tcp any eq 443 any
 deny ip any any
 exit

route-map server-client-http
 match ip address return-traffic-http
 set ip next-hop 10.20.44.44
 exit
interface Vlan5
 ip policy route-map server-client-http
 exit
```


In this configuration, the ACL identifies the traffic to match and send back to the CSM. After the ACL intercepts HTTP (**permit tcp any eq 80 any**) and SSL (**permit tcp any eq 443 any**) traffic arriving on VLAN 5 (**ip policy route-map server-client-http**), it sends the traffic to the CSM (**set ip next-hop 10.20.44.44**).

The assumption in this configuration is that all HTTP and SSL traffic leaving the servers is subject to load balancing by the CSM. However, if some web/application servers are not load balanced by the CSM, two categories need to be differentiated: non-load balanced web-servers and load-balanced web-servers.

You can configure servers to use a different Layer 4 port depending on whether the servers are load balanced or not. For example, load-balanced servers might use port 8080, while non load-balanced servers might use port 80. The CSM translates incoming requests to the VIP on port 80 and rewrites the destination port to 8080 when performing the selection of the real server.

The configuration on the CSM appears as follows:

```
vserver WEBAPPLICATIONS
  virtual 10.20.5.80 tcp www
  vlan 44
  server farm WEBAPP
  persistent rebalance
  inservice
!
server farm WEBAPP
  nat server
  no nat client
  real 10.20.5.100 8080
  inservice
  real 10.20.5.101 8080
  inservice
!
```

The configuration on the Catalyst 6500 appears as follows:

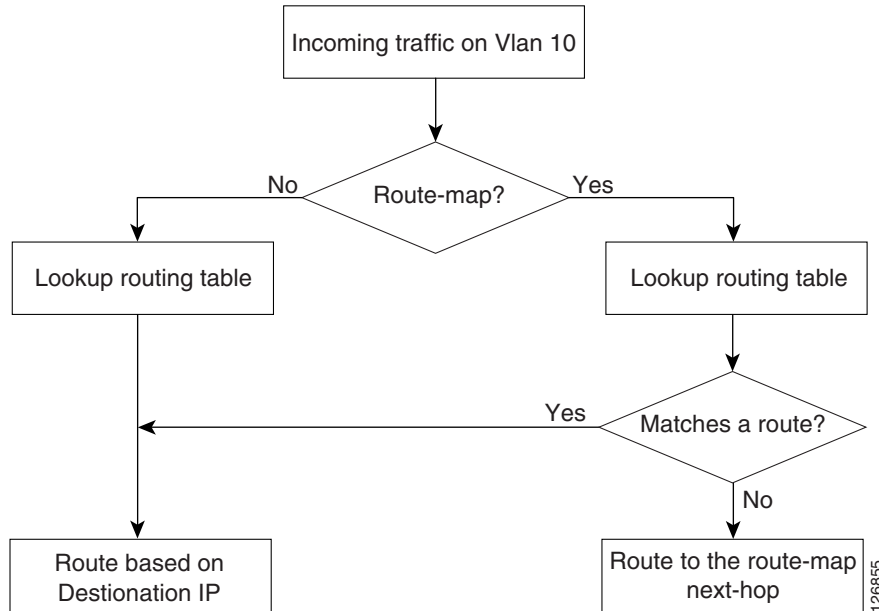
```
ip access-list extended return-traffic-http
  permit tcp any eq 8080 any
  deny ip any any
  exit

route-map server-client-http
  match ip address return-traffic-http
  set ip next-hop 10.20.44.44
  exit
interface Vlan5
  ip policy route-map server-client-http
  exit
```

Default Next-Hop

The PBR default next-hop option provides another way for identifying the traffic to send to the CSM and the traffic to be forwarded according to the routing table. With this option, the MSFC performs routing table lookups on incoming server traffic as normal. If the destination IP address matches a route in the main routing table, the MSFC forwards the traffic accordingly. Otherwise, the Catalyst 6500 forwards the traffic to the CSM. If the destination IP address matches only the default route, and the **default next-hop** command is enabled on the incoming VLAN, the MFSC uses the route map next hop. In this case, the next hop defined by PBR is the CSM (see [Figure 5-7](#)). Alternatively, you can use access lists to explicitly configure the subnets that do or do not require load balancing.

Figure 5-7 IP Default Next-Hop Algorithm



Configuration Details

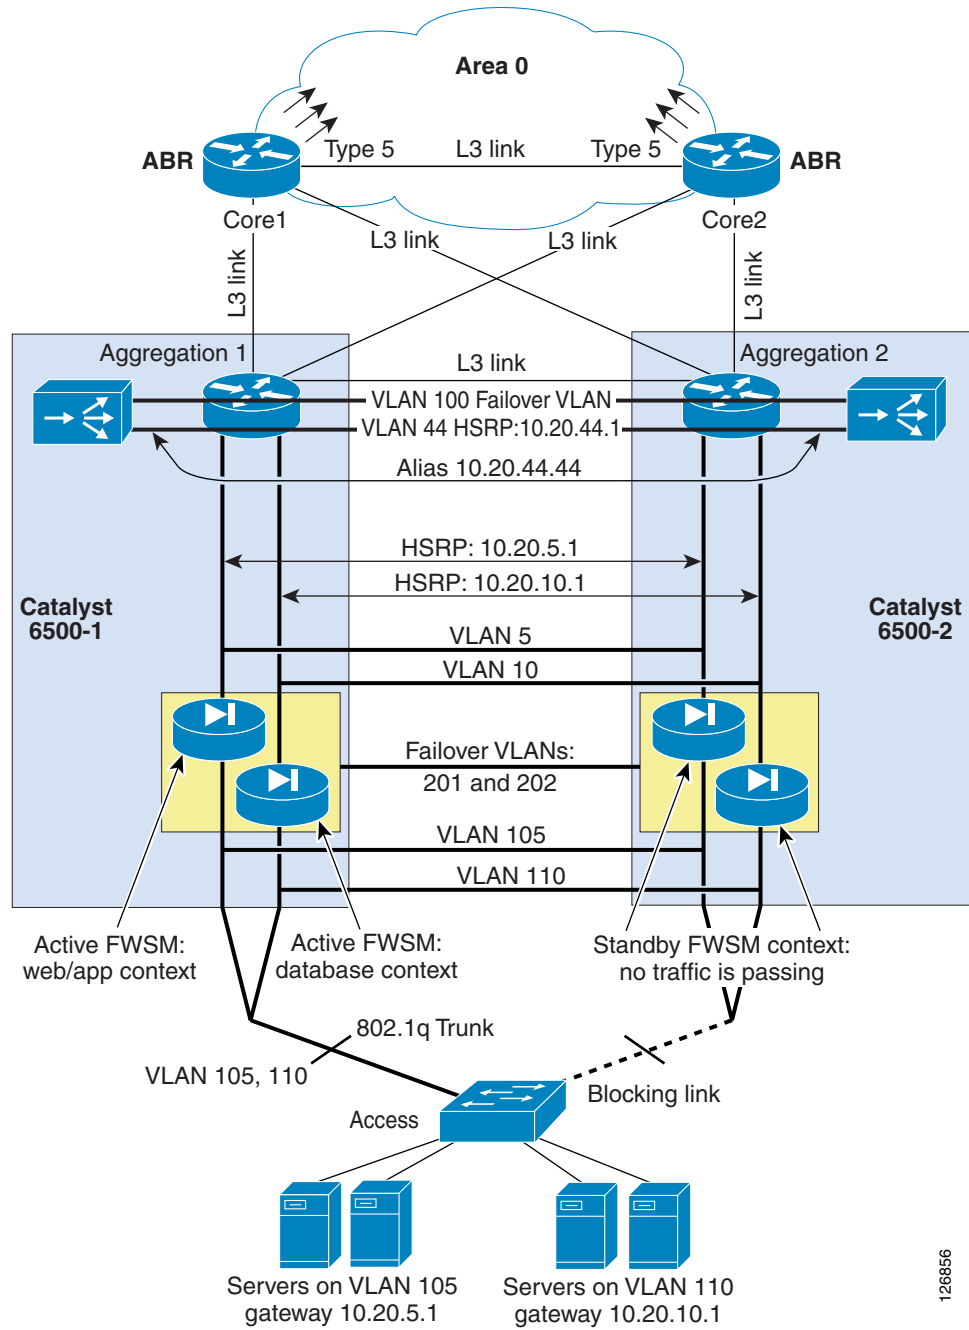
This section provides configuration details for CSM one-arm design and includes the following topics:

- [Topology](#)
- [Server VLANs and Client VLANs](#)
- [Configuration of the Trunk between CSM and Catalyst 6500](#)
- [Server-Originated Connections](#)
- [Configuration Procedure](#)

Topology

This section describes the topology illustrated in [Figure 5-8](#), which uses traditional Cisco multi-layer design.

Figure 5-8 Fully Redundant Data Center Topology with One-Arm CSM



In Figure 5-8, the servers are connected to VLAN 105 and VLAN 110. Some servers in VLAN 105 and 110 are load balanced by the CSM, and other servers are not. The CSM, which is the load-balancing device in this example, uses two VLANs: VLAN 44 and VLAN 100.

The access switches carry VLAN 105 and VLAN 110, respectively. The Spanning Tree algorithm is Rapid PVST+, which uses IEEE 802.1w. When using PVST+, enable UplinkFast and BackboneFast on these switches and enable PortFast on the server ports. The uplinks from the access switches connect to the aggregation switches, 6500-1 and 6500-2. The uplinks can be trunks if needed to carry more than one VLAN.

VLAN 44 provides communication between the routers and the CSM, while VLAN 100 is the fault-tolerant VLAN. The two CSMs use the fault-tolerant VLAN to exchange redundancy information that identifies the active and backup devices.

The aggregation switches (6500-1 and 6500-2) trunk the access VLANs (5, 10, 105, and 110). The CSM uses VLANs (100 and 44) and the FWSM VLANs (201 and 202) on an EtherChannel. 6500-1 and 6500-2 are the root and the secondary root switches respectively for all of the VLANs. When using PVST+, enable BackboneFast.

Because VLAN 44 and VLAN 100 are trunked between 6500-1 and 6500-2, they do not need to be carried to the access layer. The MSFCs use HSRP to provide the default gateway for the servers (10.20.5.1 and 10.20.10.1) and the CSM (10.20.44.1). 6500-1 is the HSRP primary for all groups and 6500-2 is the HSRP secondary for all HSRP groups.

In this example, apply PBR on VLAN 5 and VLAN 10 on both 6500-1 and 6500-2. You can configure a static route on the MSFC to map the VIP address for the server farm to the alias address of the CSM (10.20.44.44) or you can enable RHI on the CSM.

Server VLANs and Client VLANs

When deploying the CSM in one-arm mode, the data path between the MSFC and the CSM uses a single VLAN. This VLAN, configured on the CSM, can be either a server VLAN or a client VLAN, meaning the configuration on the CSM can be either of the following:

```
Agg1(config-module-csm)#vlan 44 ?
  client  client vlan
  server  server vlan
```

The difference between the two VLANs relates to how the CSM rate limits control and slow path traffic, which includes the following:

- FTP control channel
- RTSP control channel and some data channels
- ARP traffic
- ICMP traffic for which the CSM is responsible
- HSRP traffic that the CSM is snooping
- Health monitoring traffic
- Network management traffic (SNMP)

A server VLAN allows four times more packets per second (pps) than a client VLAN. Because this is the only VLAN that the CSM uses, and this VLAN interface generates the health monitoring probes, better scalability is achieved by configuring this as a server VLAN.

Configuration of the Trunk between CSM and Catalyst 6500

Clear the trunk between the CSM and the Catalyst 6500 from unnecessary VLANs.

Use the **show etherchannel summary** command to find out the port channel assigned to the CSM and then use the **range** command from Po255 to the CSM port channel to clear the configuration from unused VLANs:

```
agg1#show etherchannel summary
Flags:  D - down          P - in port-channel
        I - stand-alone  S - suspended
```

```

H - Hot-standby (LACP only)
R - Layer3      S - Layer2
U - in use      f - failed to allocate aggregator

u - unsuitable for bundling
Number of channel-groups in use: 4
Number of aggregators:          4

```

Group	Port-channel	Protocol	Ports			
2	Po2 (SU)	LACP	Gi8/1 (P)	Gi8/2 (P)	Gi8/3 (P)	Gi8/4 (P)
			Gi8/5 (P)	Gi8/6 (P)	Gi8/7 (P)	Gi8/8 (P)
255	Po255 (SD)	-				
260	Po260 (SU)	-	Gi4/1 (P)	Gi4/2 (P)	Gi4/3 (P)	Gi4/4 (P)
272	Po272 (SD)	-	Gi3/1 (D)	Gi3/2 (D)	Gi3/3 (D)	Gi3/4 (D)
			Gi3/5 (D)	Gi3/6 (D)		

In the previous example, you can see that the channel between the Catalyst 6500 and the CSM is Po260.

```

interface range Po255 - 260
  switchport trunk allowed vlan <CSM VLAN list>
!
```

Server-Originated Connections

With server-originated connections, PBR pushes flows to the CSM that were not load balanced by the CSM. These flows are unknown to the CSM, and by default, they are rejected.

By default, the CSM forwards only the traffic that matches either an existing flow or a VIP. Because of the PBR setup, some flows from VLAN 5 might be redirected to the CSM such as direct flows sent to the real server address (not sent to the VIP/80).

There are two possible solutions to this problem:

- With CSM Release 4.2, use the environment variable `ROUTE_UNKNOWN_FLOW_PKTS 2`.
- With a release before Release 4.2, configure a vserver with a server farm that uses the predictor forward.



Note

Before Release 4.2, the environmental variable `ROUTE_UNKNOWN_FLOW_PKTS 1` enables only the forwarding of NON-SYN packets that do not hit any VIP. In Release 4.2, `ROUTE_UNKNOWN_FLOW_PKTS 2` also allows routing for SYN packets that do not hit a VIP, without creating a flow.

Configuration Procedure

You can use the CLI or the CiscoView Device Manager (CVDM) to configure the CSM. If you use CVDM, you need to complete the configuration with the CLI because the current version of CVDM (1.0) does not yet support specific configuration tasks required by the one-arm design.

The following are the key configuration steps to configure the CSM in one-arm mode:

1. Configure the servers to listen on the appropriate port (for example, load-balanced servers to listen to port 8080).
2. Create the data path between the CSM and the MSFC using a VLAN.
3. Define the route map.

4. Apply the route map to the MSFC VLAN interfaces.
5. On the MSFC, define a static route for the VIP pointing to the alias IP address on the CSM or enable RHI on the CSM.
6. Enable DoS protection on the CSM.

CVDM

To use CVDM, start the HTTP server by entering the following commands:

```
! web-based administration requires privilege 15
!
username webadmin privilege 15 secret 0 C1sC0!w3B
!
! Change the web access to use port different from port 80
!
ip http server
ip http port 8768
ip http authentication local
ip http access-class 5
ip http path bootflash:
```



Note

To use HTTP for configuration, be sure to configure authentication and ACLs to limit the devices that are allowed to access this service. Cisco recommends using a special VLAN for management.

CVDM uses the HTTP server on the Catalyst 6500 to download a Java applet that runs on the PC used to configure the Catalyst 6500. If the image on the Catalyst 6500 supports SSH but it has not been enabled, CVDM displays a prompt and enables SSH if you confirm the operation. Subsequently, the applet can use SSH to configure the switch.

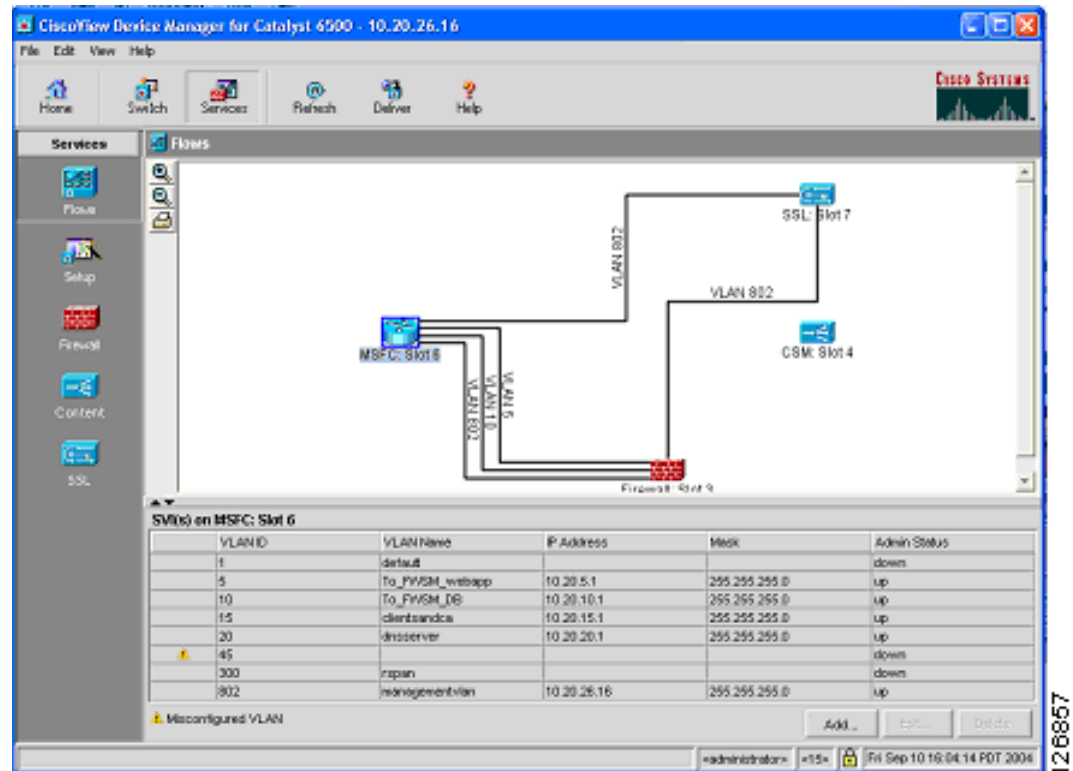


Note

The Java window that prompts for the credentials is often hidden by the CVDM window.

CVDM graphically displays the data path between the service modules inside the Catalyst 6500. For example, [Figure 5-9](#) shows the Flows view for a configuration where the firewall is already connected to the MSFC over VLAN 5 (the outside VLAN of the FWSM context for web/application servers) and VLAN 10 (the outside VLAN of the FWSM context for database servers).

Figure 5-9 Flows View in CVDM



Creating the Data Path between the CSM and the MSFC

The CSM can be configured to use a single VLAN that does not require client and server VLANs. For performance reasons, Cisco recommends configuring this VLAN on the CSM as a server VLAN. Within the server VLAN configuration, define the MSFC as the gateway for the CSM. In the example shown in Figure 5-8, configure VLAN 44 as the server VLAN on the CSMs because that is the VLAN connecting the CSMs.

CLI Configuration

The first CLI configuration step consists in creating the VLAN by entering the following commands:

```
agg1(config)# vlan <vlan number>
agg1(config-vlan)# description msfc-csm
```

On the MSFC, assign an IP address to this VLAN Interface by entering the following commands:

```
interface Vlan44
description CSMVLAN
ip address 10.20.44.2 255.255.255.0
standby 1 ip 10.20.44.1
standby 1 timers 1 3
standby 1 priority 120
standby 1 preempt delay minimum 180
no ip directed-broadcast
no ip unreachable
no ip redirects
no ip proxy-arp
```

```
! >> Disable NTP services <<
ntp disable
no shut
exit
!
```

On the CSM, enter the following commands:

```
module ContentSwitchingModule <module number>
vlan <vlan number> server
ip address <CSM MAIN IP ADDRESS>
gateway <MSFC IP ADDRESS>
alias <CSM IP ADDRESS>
```

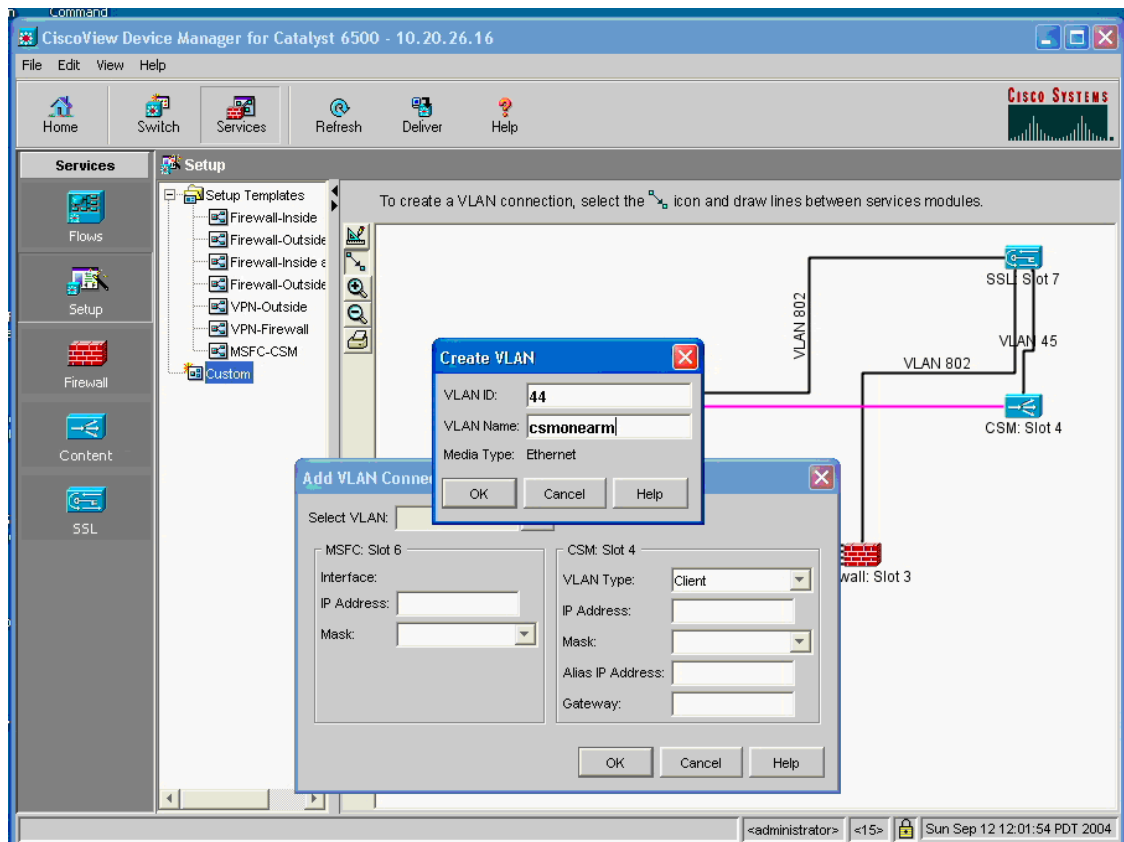
CVDM Configuration

To use CVDM to complete the configuration, connect the CSM to the MSFC over a new VLAN (for example, VLAN number 44), by completing the following steps.

- Step 1** Click Setup on the left side of the window and then click Custom.
- Step 2** Click the Line icon and drag a new line between the MSFC icon and the CSM icon.

A window appears, as shown in [Figure 5-10](#):


Figure 5-10 Configuring the VLAN between CVDM and MSFC with CVDM



- Step 3** On the Select VLAN drop-down list, click **Create VLAN**.

- Step 4** In the Create VLAN dialog box, enter VLAN 44, and give it a name, such as “csmonearm” and click OK.
- Step 5** In the Add VLAN Connection dialog box, configure the HSRP IP address for the MSFC (10.20.44.1) and the IP address for the CSM (10.20.44.43).
- Step 6** Select Server from the VLAN Type list, as shown in [Figure 5-11](#).

Figure 5-11 Configuring the CSM VLAN with CVDM



Note

The MSFC IP address in the CVDM configuration unfortunately is not the HSRP address, so to get the configuration working, use 10.20.44.1. However, when you configure the redundant MSFC change the address to 10.20.44.2 and use 10.20.44.1 for the HSRP configuration

With CVDM, you need to click the **Deliver** button to apply the configuration for it:

- Step 7** Modify the MSFC configuration as follows:

```
interface VLAN44
description CSMVLAN
ip address <MSFC IP ADDRESS>
standby 1 ip 10.20.44.1
standby 1 timers 1 3
standby 1 priority 120
standby 1 preempt delay minimum 180
no ip directed-broadcast
no ip unreachablees
no ip redirects
no ip proxy-arp
! >> Disable NTP services <<
ntp disable
no shut
exit
```

Configuring Policy-Based Routing

The CSM, like the FWSM, remembers the state of flows, and all the traffic sent by the CSM to the real servers needs to flow back through the CSM.

For Layer 5 inspection or higher, the CSM buffers the TCP flows before sending the traffic to a real server until it gets all the upper layer information needed to take a decision (policy, predictor, match, and so on). If the upper layer traffic does not return to the CSM, the action to distribute the traffic to the server farm fails.

You need to apply route maps to the VLAN interface wherever servers are load balanced by the CSM. In example shown in [Figure 5-8](#), the only servers that are load balanced by the CSM are in VLAN 5. The next hop IP address belongs to VLAN 44 regardless of the subnet where the route is applied. In this topology, the next hop is the CSM alias IP, which does not belong to the subnet where the route map is applied. The required configuration is as follows:

```
ip access-list extended return-traffic-http
  permit tcp any eq 8080 any
  permit tcp any eq 443 any
  deny ip any any
exit

route-map server-client-http
  match ip address return-traffic-http
  set ip next-hop <CSM ALIAS>
exit
```

The route map is applied to VLAN 5, as in the following example:

```
interface VLAN5
  ip address 10.20.5.2
  no ip redirects
  ip policy route-map server-client-http
  standby 1 ip 10.20.5.1
  standby 1 timers 1 3
  standby 1 priority 120
  standby 1 preempt delay minimum 180
  no ip unreachable
  no ip redirects
  no ip proxy-arp
  ! >> Disable NTP services <<
  ntp disable
  no shut
  exit
```

In addition to the specific PBR configuration, it is important to configure the CSM to forward segments that do not match either an existing vserver or an existing connection. The required configuration is as follows:

```
module ContentSwitchingModule 4
!
variable ROUTE_UNKNOWN_FLOW_PKTS 1
variable ROUTE_UNKNOWN_FLOW_PKTS 2
!
server farm FORWARD
  no nat server
  no nat client
  predictor forward
exit
```

By default, the CSM creates entries in its flow table for TCP connections that match a vserver and forwards segments for all previously created flows (from outside). PBR sends any return traffic defined in the PBR back to the CSM, and not just traffic for load-balanced connections. Therefore, the CSM needs to be configured to forward the unknown traffic to its gateway (MSFC) as follows:

```
module ContentSwitchingModule 4
variable ROUTE_UNKNOWN_FLOW_PKTS 1
vserver CATCHALL
```

```

virtual 0.0.0.0 0.0.0.0 any
vlan 44
server farm FORWARD
persistent rebalance
inservice
exit

```

**Note**

In Release 4.1, the environmental variable `ROUTE_UNKNOWN_FLOW_PKTS 1` forwards NON-SYN packets that do not hit any VIP. In Release 4.2, `ROUTE_UNKNOWN_FLOW_PKTS 2` also allows routing for SYN packets that do not hit a VIP, without creating a flow.

The configuration (in bold text above) is necessary because of the way `ROUTE_UNKNOWN_FLOW_PKTS` in Release 4.1 works. Starting from Release 4.2, this will not be necessary.

Configuring the CSM Server Farm and Virtual Server

This example assumes that the servers that require load balancing are 10.20.5.105 and 10.20.5.106, while the VIP address to be used for HTTP traffic is 10.20.5.80.

From the CVDM Catalyst 6000 Service window, click **Content Switch** and launch CVDM-CSM. The screen shown in [Figure 5-12](#) appears.

Figure 5-12 CSM Device Manager

The screenshot displays the CSM Device Manager interface with the following sections:

- System Overview:**
 - Model Type: YWS-M5065-SLB-APC
 - Serial Number: S4D075004UP
 - Slot Number: 4
 - Software Version: 4.1(2)
 - Hardware Version: 1.7
 - Overflow Errors: 0
 - Redundancy Status: None
 - FT VLAN ID:
 - Network Processor Utilization:
 - BP1-Session: 0%
 - BP2-TCP: 0%
 - BP3-Layer 7: 0%
 - BP4-Load Balancing: 0%
 - BP5-MAT: 0%
- Service Dashboard:**
 - Virtual Servers:**
 - Inservice: 0
 - Out of Service: 0
 - Policy Associated: 0
 - Default Policy: 0
 - Policies:**
 - Configured Policies: 0
 - Policies Without Conditions: 0
 - Policy Without Actions: 0
 - Server Farms:**
 - Total: 0
 - Available: 0
 - Real Servers:**
 - Needed Rerolls: 0
 - Unrerolled Rerolls: 0
- Connection Dashboard:**
 - Graphs: Absolute Values
 - Current Connections: 0% of 1 Million Connections
 - All values are in units of Kilo connections.
 - Cumulative Connections:
 - Created: 0 K
 - Destroyed: 0 K
 - Failed: 0 K
 - Timed Out: 0 K
 - L4 and L7 Connections:
 - L4 Decisions: 0 K
 - L7 Decisions: 0 K
 - L4 Rejections: 0 K
 - L7 Rejections: 0 K
- Server Dashboard:**
 - My Virtual Servers table with columns: Name, IP Address, Operational Status.

At the bottom, there is an FAQ section with the question: "How do I setup a Virtual Server with Default Policy?" and a "Go" button. The system status bar shows the user is administrator, session ID is 15, and the time is Sat Sep 11 05:12:53 PDT 2004.

Server Farm Configuration

Perform the following procedure to configure the server farm. (See [Figure 5-13](#).)

Figure 5-13 Configuring Real Servers

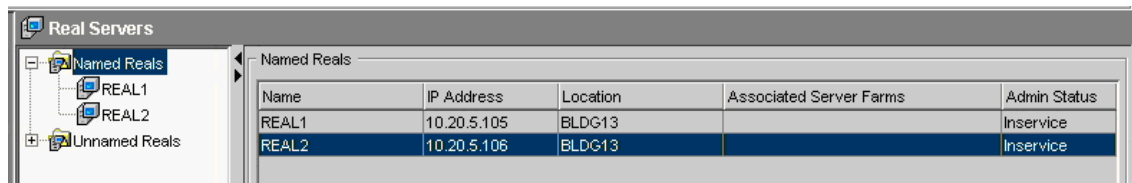

The 'Add Real Server' dialog box contains the following fields and controls:

- Name:
- IP Address:
- Location:
- Status:
- Buttons: OK, Cancel, Help

126861

- Step 1** Click **Setup**.
- Step 2** Click **Real Servers**.
- Step 3** Select **Named Reals**.

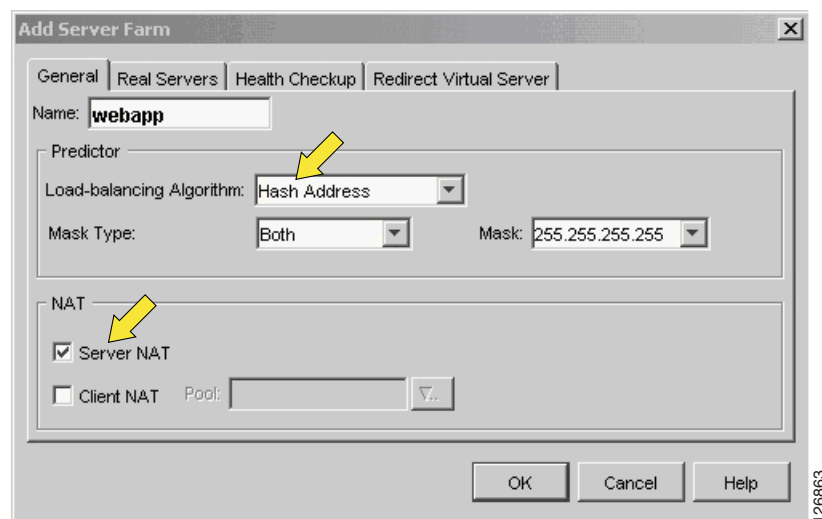
Figure 5-14 is displayed when selecting OK.

Figure 5-14 Named Real Servers


Name	IP Address	Location	Associated Server Farms	Admin Status
REAL1	10.20.5.105	BLDG13		Inservice
REAL2	10.20.5.106	BLDG13		Inservice

126862

- Step 4** Create a server farm and associate a predictor with it. (See Figure 5-15.)

Figure 5-15 Configuring the Server Farm


The 'Add Server Farm' dialog box has the following configuration:

- Name:
- Predictor section:
 - Load-balancing Algorithm: (indicated by a yellow arrow)
 - Mask Type:
 - Mask:
- NAT section:
 - Server NAT (indicated by a yellow arrow)
 - Client NAT Pool:
- Buttons: OK, Cancel, Help

126863

By default, Server NAT, which is also known as Direct Mode, is enabled, and Client NAT is disabled. With Server NAT, the destination IP address and port number in the load-balanced packets are replaced with the IP address of one of the real server defined in the server farm.

**Note**

Client sends a packet to the VIP. The CSM load balances this packet while preserving the original source address and translates the destination IP address from the VIP to one of the real server IP addresses.

- Step 5** Assign the real servers to the server farm by selecting the **Real Servers** tab. (See [Figure 5-16](#).)

Figure 5-16 Assigning the Real to the Server Farm

- Step 6** Select **Add Named Real Servers**.
- Step 7** Select **In Service** from the Status list.
- Step 8** Configure PORT NAT, if Layer 4 port is used by PBR to identify load-balanced traffic.

Virtual Server Configuration

To configure a virtual server, complete the following steps (see [Figure 5-17](#)):

Figure 5-17 Virtual Server Configuration

Step 1 Select **Virtual Server** and add a new virtual server.

Step 2 Select the VLANID.

For example, selecting VLANID 44 makes sure that the virtual server accepts only traffic destined to 10.20.5.80 coming from VLAN44.

Step 3 Configure the protocol to be TCP and the port to be 80.

Step 4 Enable the option **Advertise Virtual IP**.

The CSM allows advertising the IP address of the virtual server as a host route. By enabling the **advertise** command in the virtual server configuration mode, the CSM injects into the router (the MSFC) the VIP as a static route. Also, this keeps the routing table up-to-date as long as the virtual server is operational. If all real servers are down, the static route is removed immediately. This feature is helpful for disaster recovery.

Step 5 Configure a Default Policy and select the server farm that you previously configured.

Configuring DoS Protection

SYN cookies provide greater scalability in withstanding SYN floods. When the number of SYN/s passes a threshold configured by the user on the CSM, the CSM sends a SYN/ACK with an initial sequence number (ISN) calculated according to a cryptographic function of (MSS, source IP address and port, destination IP address and port, and a secret key on the CSM).



Note

SYN cookie technology requires a secret key to generate a cookie. The hashing algorithm uses this key. The CSM generates a new key every three seconds, by default.

The cookie is sent back to the host in the SYN-ACK as the ISN. CSM resources do not maintain the connection request information sent by the host; this information exists in the cookie. If the host responds with an ACK, the cookie is available to the CSM in the acknowledgement number field. The

CSM reconstructs the original SYN information from the cookie (acknowledgement number field -1) by reversing the hash operation. The CSM only initiates a back-end connection to a server when it receives a data packet from the host. The CSM then programs the connection in the fast path.

CLI Configuration

Using CatOS CLI, modify the virtual server for HTTP traffic by entering the following commands:

```
agg(config)# mod contentSwitchingModule 4
agg(config-module-csm) # vserver WEBAPP
agg(config-slb-vserver) #no inservice
agg(config-slb-vserver) # virtual 10.20.5.80 tcp www service termination
agg(config-slb-vserver) # replicate csrp connection
agg(config-slb-vserver) # replicate csrp sticky
agg(config-slb-vserver) #inservice
agg(config-slb-vserver) #end
```

The default embryonic threshold is 5000. To modify this threshold set the `SYN_COOKIE_THRESHOLD` variable to any number between zero and one million. For example, to utilize SYN cookies for all connections requests set the threshold to zero. To modify the threshold, enter the following commands:

```
agg(config)# mod contentSwitchingModule 4
agg(config-module-csm) # variable SYN_COOKIE_THRESHOLD threshold
agg(config-module-csm) # end
```

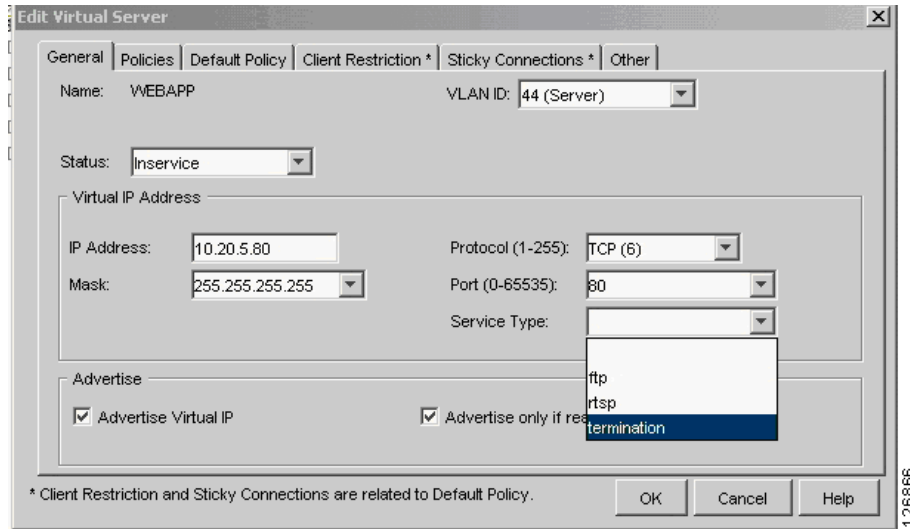
Use the `SYN_COOKIE_INTERVAL` variable to modify the key generation period from 1 to 60 seconds. To modify this variable, enter the following commands:

```
agg(config)# mod contentSwitchingModule 4
agg(config-module-csm) # variable SYN_COOKIE_INTERVAL time
agg(config-module-csm) # end
```

CVDM Configuration

From the CVDM, select **termination** under Service Type from Edit Virtual Server window (see [Figure 5-18](#)).

Figure 5-18 Configure DoS Protection on the CSM



The default embryonic threshold is 5000. To modify this threshold set the `SYN_COOKIE_THRESHOLD` variable to any number between zero and one million. For example, to use SYN cookies for all connection requests set the threshold to zero. To change this variable, enter the following commands:

```
agg(config)# mod contentSwitchingModule 4
agg(config-module-csm) # variable SYN_COOKIE_THRESHOLD threshold
agg(config-module-csm) # end
```

Use the `SYN_COOKIE_INTERVAL` variable to modify the key generation period from 1 to 60 seconds. To change this variable, enter the following commands:

```
agg(config)# mod contentSwitchingModule 4
agg(config-module-csm) # variable SYN_COOKIE_INTERVAL time
agg(config-module-csm) # end
```

Monitoring

If the VIP address is under attack, you can monitor the CSM operations by entering the following commands:

```
agg1# show mod csm 4 tech-support processor 1
agg1# show module csm 4 tech-support utilization
```

To view the information about the IXP engines utilization, enter the following command:

```
agg# show mod csm 4 tech-support util | include IXP
```

```
IXP Engines
  IXP1      34%
  IXP2      23%
  IXP3       0%
  IXP4       0%
  IXP5      16%
```


Configuring Redundancy

The configurations described so far in this guide apply to the Catalyst 6500-1 in [Figure 5-8](#). To complete the configuration, you also need to configure Catalyst 6500-2. The same configurations apply with specific changes to the CSM VLAN configuration, the CSM failover VLAN configuration, and the HSRP groups on the MSFC.

Trunk Configuration

Configure the trunk between the two Catalyst 6500s to carry the following VLANs (the VLANs used in this example are shown in parentheses):

- The VLAN used by the MSFC and the CSM (VLAN 44)
- The outside VLANs of the FWSM (VLAN 5 and 10)
- The inside VLANs of the FWSM (VLAN 105 and 110)
- The CSM fault-tolerant VLAN (VLAN 100)
- The FWSM failover VLANs (VLAN 201 and 202)

PBR

The PBR configuration is the same on 6500-1 and 6500-2, as shown below:

```
ip access-list extended return-traffic-http
 permit tcp any eq 8080 any
 permit tcp any eq 443 any
 deny ip any any
exit
!
route-map server-client-http
 match ip address return-traffic-http
 set ip next-hop 10.20.44.44
exit
!
interface Vlan5
 ip policy route-map server-client-http
exit
```

HSRP

The HSRP configuration on 6500-2 has the same HSRP IP address, different priority and obviously a different Interface VLAN IP address.

```
interface VLAN44
 description msfc_to_csm_VLAN
 ip address <6500-2 IP ADDRESS on VLAN 44>
 standby 1 timers 1 3
 standby 1 priority 110
 standby 1 preempt delay minimum 180
 no ip unreachable
 no ip redirects
 no ip proxy-arp
 ! >> Disable NTP services <<
 ntp disable
 no shut
 exit
!
interface VLAN5
```

```

ip address <6500-2 IP ADDRESS on VLAN 5>
no ip redirects
ip policy route-map server-client-http
standby 1 ip 10.20.5.1
standby 1 timers 1 3
standby 1 priority 110
standby 1 preempt delay minimum 180
no ip unreachable
no ip redirects
no ip proxy-arp
! >> Disable NTP services <<
ntp disable
no shut

```

CSM

Complete the configuration on 6500-1 with the Fault-Tolerant group:

```

ft group 1 vlan 100
priority 20
heartbeat-time 1
failover 3
preempt
!

```

Disable IGMP snooping on the FT VLAN:

```

int vlan 100
no ip igmp snooping
shut
!

```

The CSM configuration on 6500-2 is as follows:

```

module ContentSwitchingModule <module number>
vlan <vlan number> server
ip address <6500-2 CSMIP ADDRESS>
gateway <MSFC IP ADDRESS>
alias <CSM IP ADDRESS>
!
ft group 1 vlan 100
priority 10
heartbeat-time 1
failover 3
preempt
!

```

The rest of the configuration is identical to the CSM on 6500-1:

```

real REAL1
address 10.20.5.105
inservice
!
real REAL2
address 10.20.5.106
inservice
!
vserver WEBAPPLICATIONS
virtual 10.20.5.80 tcp www service termination
vlan 44
serverfarm WEBAPP
advertise active
replicate csrpf connection

```

```

        replicate csrp sticky
        persistent rebalance
        inservice
    !
serverfarm WEBAPP
    nat server
    no nat client
    real REAL1 8080
        inservice
    real REAL2 8080
        inservice
    !

```

Configuration Listings

This section provides sample configuration listings for the different devices in the recommended design. It includes the following topics:

- [CSM1 Configuration](#)
- [CSM2 Configuration](#)
- [MSFC-AGG1 Configuration](#)
- [MSFC-AGG2 Configuration](#)

CSM1 Configuration

```

module ContentSwitchingModule 4
    !
variable ROUTE_UNKNOWN_FLOW_PKTS 1
variable ROUTE_UNKNOWN_FLOW_PKTS 2
    !
vlan 44 server
    ip address 10.20.44.43 255.255.255.0
    gateway 10.20.44.1
    alias 10.20.44.44 255.255.255.0
    !
ft group 1 vlan 100
    priority 20
    heartbeat-time 1
    failover 3
    preempt
    !
probe HTTP-8080 http
    port 8080
    interval 5
    retries 3
    !
real REAL1
    address 10.20.5.105
    inservice
    !
real REAL2
    address 10.20.5.106
    inservice
    !
serverfarm WEBAPP
    nat server
    no nat client

```

```

real name REAL1 8080
  inservice
real name REAL2 8080
  inservice
probe HTTP-8080
!
vserver WEBAPPLICATIONS
  virtual 10.20.5.80 tcp www service termination
  vlan 44
  serverfarm WEBAPP
  advertise active
  replicate csrp connection
  replicate csrp sticky
  persistent rebalance
  inservice
!
serverfarm FORWARD
  no nat server
  no nat client
  predictor forward
!
vserver CATCHALL
  virtual 0.0.0.0 0.0.0.0 any
  vlan 44
  serverfarm FORWARD
  persistent rebalance
  inservice
!

```

CSM2 Configuration

```

module ContentSwitchingModule 4
!
variable ROUTE_UNKNOWN_FLOW_PKTS 1
variable ROUTE_UNKNOWN_FLOW_PKTS 2
!
vlan 44 server
  ip address 10.20.44.45 255.255.255.0
  gateway 10.20.44.1
  alias 10.20.44.44 255.255.255.0
!
ft group 1 vlan 100
  priority 10
  heartbeat-time 1
  failover 3
  preempt
!
probe HTTP-8080 http
  port 8080
  interval 5
  retries 3
!
real REAL1
  address 10.20.5.105
  inservice
!
real REAL2
  address 10.20.5.106
  inservice
!
serverfarm WEBAPP
  nat server

```

```

no nat client
real name REAL1 8080
  inservice
real name REAL2 8080
  inservice
probe HTTP-8080
!
vserver WEBAPPLICATIONS
  virtual 10.20.5.80 tcp www service termination
  vlan 44
  serverfarm WEBAPP
  advertise active
  replicate csrp connection
  replicate csrp sticky
  persistent rebalance
  inservice
!
serverfarm FORWARD
  no nat server
  no nat client
  predictor forward
!
vserver CATCHALL
  virtual 0.0.0.0 0.0.0.0 any
  vlan 44
  serverfarm FORWARD
  persistent rebalance
  inservice
!

```

MSFC-AGG1 Configuration

```

hostname aggl
!
! VTP and Spanning-Tree
! =====
!
vtp domain mydomain
vtp mode transparent
!
power redundancy-mode combined
!
spanning-tree mode rapid-pvst
spanning-tree loopguard default
spanning-tree extend system-id
spanning-tree vlan 1-1000 root primary
spanning-tree pathcost method long
!
! VLAN CONFIGURATION
!
vlan internal allocation policy descending
!
vlan 5
  name webappoutside
!
vlan 10
  name databaseoutside
!
vlan 82
  name networkmgmt
!
vlan 105

```

```

    name webappinside
    !
vlan 110
    name databaseinside
    !
vlan 200
    name fwsm_failover_vlan
    !
vlan 201
    name fwsm_flink!
    !
interface Port-channel2
    no ip address
    switchport
    switchport trunk encapsulation dot1q
    switchport mode trunk
    switchport nonegotiate
    ! >> use a != native VLANs on trunks than on access ports <<
    switchport trunk native vlan 2
    ! >> do not trunk VLAN 13 (13) , 14 (13) , 82 (mgmt) <<
    switchport trunk allowed vlan 5,10,30,44,45,100,105,110,200,201,300
    no shut
    !
    ! SVI CONFIGURATION
    ! =====
    !
    ! ip directed-broadcast often needed
    ! in serverfarms disable it if possible
    !
interface Vlan5
    description webapp
    ip address 10.20.5.2 255.255.255.0
    standby 1 ip 10.20.5.1
    standby 1 timers 1 3
    standby 1 priority 120
    standby 1 preempt delay minimum 180
    ip policy route-map server-client-http
    no ip unreachable
    no ip redirects
    no ip proxy-arp
    ! >> Disable NTP services <<
    ntp disable
    no shut
    !
interface Vlan10
    description database
    ip address 10.20.10.2 255.255.255.0
    standby 1 ip 10.20.10.1
    standby 1 timers 1 3
    standby 1 priority 120
    standby 1 preempt delay minimum 180
    no ip unreachable
    no ip redirects
    no ip proxy-arp
    ! >> Disable NTP services <<
    ntp disable
    no shut
    !
interface VLAN44
    description CSMVLAN
    ip address 10.20.44.2 255.255.255.0
    standby 1 ip 10.20.44.1
    standby 1 timers 1 3
    standby 1 priority 120

```

```

standby 1 preempt delay minimum 180
no ip directed-broadcast
no ip unreachable
no ip redirects
no ip proxy-arp
! >> Disable NTP services <<
ntp disable
no shut
exit
!
route-map server-client-http
  match ip address return-traffic-http
  set ip next-hop 10.20.44.44
!
ip access-list extended return-traffic-http
  permit tcp any eq 8080 any
  permit tcp any eq 443 any
  deny ip any any
!

```

MSFC-AGG2 Configuration

```

hostname agg2
!
! VTP and Spanning-Tree
! =====
!
vtp domain mydomain
vtp mode transparent
!
power redundancy-mode combined
!
spanning-tree mode rapid-pvst
spanning-tree loopguard default
spanning-tree extend system-id
spanning-tree vlan 1-1000 root secondary
spanning-tree pathcost method long
!
! VLAN CONFIGURATION
!
vlan internal allocation policy descending
!
vlan 5
  name webappoutside
!
vlan 10
  name databaseoutside
!
vlan 82
  name networkngmt
!
vlan 105
  name webappinside
!
vlan 110
  name databaseinside
!
vlan 200
  name fwsm_failover_vlan
!
vlan 201
  name fwsm_flink!

```

```

!
interface Port-channel2
  no ip address
  switchport
  switchport trunk encapsulation dot1q
  switchport mode trunk
  switchport nonegotiate
! >> use a != native VLANs on trunks than on access ports <<
  switchport trunk native vlan 2
! >> do not trunk VLAN 13 (13) , 14 (13) , 82 (mgmt) <<
  switchport trunk allowed vlan 5,10,30,44,45,100,105,110,200,201,300
  no shut
!
! SVI CONFIGURATION
! =====
!
! ip directed-broadcast often needed
! in serverfarms disable it if possible
!
interface Vlan5
  description webapp
  ip address 10.20.5.3 255.255.255.0
  standby 1 ip 10.20.5.1
  standby 1 timers 1 3
  standby 1 priority 110
  standby 1 preempt delay minimum 180
  ip policy route-map server-client-http
  no ip unreachable
  no ip redirects
  no ip proxy-arp
  ! >> Disable NTP services <<
  ntp disable
  no shut
!
interface Vlan10
  description database
  ip address 10.20.10.3 255.255.255.0
  standby 1 ip 10.20.10.1
  standby 1 timers 1 3
  standby 1 priority 110
  standby 1 preempt delay minimum 180
  no ip unreachable
  no ip redirects
  no ip proxy-arp
  ! >> Disable NTP services <<
  ntp disable
  no shut
!
interface VLAN44
  description CSMVLAN
  ip address 10.20.44.3 255.255.255.0
  standby 1 ip 10.20.44.1
  standby 1 timers 1 3
  standby 1 priority 110
  standby 1 preempt delay minimum 180
  no ip directed-broadcast
  no ip unreachable
  no ip redirects
  no ip proxy-arp
  ! >> Disable NTP services <<
  ntp disable
  no shut
  exit
!

```



```
route-map server-client-http
  match ip address return-traffic-http
  set ip next-hop 10.20.44.44
!
ip access-list extended return-traffic-http
  permit tcp any eq 8080 any
  permit tcp any eq 443 any
  deny ip any any
!
```

