

# Content Services Switch FAQ

Document ID: 25901

---

## Questions

### Introduction

- Where can I find the MIBs for the CSS?
  - What is the maximum number of scripted keepalives that the CSS supports?
  - How can I clear or remove core files?
  - Where can I find interpretations of log messages?
  - Is there a command that controls how often peers send load reports to each other?
  - Do license keys change with code versions?
  - I lost my license key. What do I do?
  - What is the default time for the retention of an entry in a sticky table?
  - How do I configure sticky mask in order to cover requests from a mega-proxy like America Online (AOL)?
  - Why is there no option for sticky when I use advanced-balance Secure Socket Layer (SSL)?
  - What type of encryption does Content and Application Peering Protocol (CAPP) or Application Peering Protocol (APP) use?
  - What does the "gratuitous arp" message mean?
  - How do I synchronize configurations over the CSS in failover mode?
  - What settings should I use in a terminal program?
  - Is there a way to reprogram the MAC address on a CSS?
  - How do I make a permanent prompt change on the CSS?
  - What is the difference between operational and locked Flash?
  - Why are there different versions of Flash?
  - Why can I not access the management port of the CSS from a remote port?
  - Does Cisco Technical Support support custom-script keepalives that the customer writes?
  - How do I remove core files from the CSS disk?
  - When I authenticate to a RADIUS server with my CSS, I get the "RADIUS-4: RADIUS Authentication failed with reason code 2" error message. What does the message mean?
  - How large is the sticky table, and what causes the removal of entries?
  - How can I take a service out of rotation?
  - Is network proximity part of the enhanced feature set?
  - What details does the show dos command provide?
  - Can I turn off the Denial of Service (DoS) protection feature on the CSS line of switches?
  - Can I turn off the Denial of Service (DoS) protection counters?
  - How do I use port ranges in access lists?
- NetPro Discussion Forums – Featured Conversations
- Related Information
- 

## Introduction

This document addresses the most frequently asked questions (FAQ) about the Cisco Content Services Switch (CSS).

Refer to Cisco Technical Tips Conventions for more information on document conventions.

## Q. Where can I find the MIBs for the CSS?

A. The MIBs are already on the CSS. You can consider the CSS an agent in the Simple Network Management Protocol (SNMP) network scheme. All you need to do is configure the SNMP parameters on the CSS. Refer to the document *Configuring Simple Network Management Protocol (SNMP)* for more information.

## Q. What is the maximum number of scripted keepalives that the CSS supports?

A. The maximum number of scripted keepalives that CSS supports is 255. Refer to the *New Features in Software Version 5.00* section of the Release Note for the Cisco 11000 Series Content Services Switch.

## Q. How can I clear or remove core files?

A. Issue the **clear core** command. The command is available in CSS software version 5.00 and later, in debug mode. The syntax is:

```
css150(debug)#clear core filename CR
```

## Q. Where can I find interpretations of log messages?

A. For interpretations of log messages, refer to the document *Log Messages*.

## Q. Is there a command that controls how often peers send load reports to each other?

A. You can use the **dns-peer interval** command. There are also additional commands that you can configure locally in order to achieve a quicker measure of the local load:

- ◆ **ageout-timer** Sets the time (in seconds) of the ageout of stale load information.
- ◆ **teardown-timer** Sets the maximum time period (in seconds) that the system waits to send a teardown report.

## Q. Do license keys change with code versions?

A. No, license keys do not change with code versions.

## Q. I lost my license key. What do I do?

A. Send an email with the serial number of your CSS to [licensing@cisco.com](mailto:licensing@cisco.com). The **version** command displays the feature pack, but not the license key.

## Q. What is the default time for the retention of an entry in a sticky table?

A. Unless you use the command **sticky-inact-timeout**, there is no default time. The sticky table is kept on a FIFO basis (32,000 or 128,000 entries, according to the device type and memory available), or until the reboot of the CSS.

## Q. How do I configure sticky mask in order to cover requests from a mega-proxy like America Online (AOL)?

A. If an application requires a user to be stuck for the entire life of the session, consider a Layer 3 sticky. A Layer 3 sticky sticks a user to a server on the basis of the user IP address. The CSS has a sticky table of 32,000, which means that when 32,000 simultaneous users are on the site, the table wraps and the first users become "unstuck". However, the volume of your site can be such that you have more than 32,000 users at a time. Or a large percentage of your customers can come to you through a mega-proxy. In these cases, consider either the use of a different sticky method (such as the cookie, cookieurl, or url) or an increase of your sticky mask. The default sticky mask is 255.255.255.255, which means that each entry in the sticky table is an individual IP address. Some of the mega-proxies have a situation in which one user over the life of one session uses several different IP addresses in a range of addresses. This situation causes some of the TCP connections to get stuck to one server, and can cause other connections to get stuck to a different server for the same transaction. A result can be the loss of some items from the grocery cart. If you cannot use one of the more advanced methods of sticking, use the sticky mask of 255.255.240.0 when your client base comes through one of these mega-proxies.

## Q. Why is there no option for sticky when I use advanced-balance Secure Socket Layer (SSL)?

A. Advanced-balance SSL is the same as sticky SSL.

## Q. What type of encryption does Content and Application Peering Protocol (CAPP) or Application Peering Protocol (APP) use?

A. By default, CAPP uses no encryption. You can configure the APP session to use Message Digest 5 (MD5). The encryption type must be the same on both peers in order for the APP session to come up.

## Q. What does the "gratuitous arp" message mean?

A. When the backup switch does not detect a heartbeat from the master switch within 3 seconds, the backup switch transitions to become the master and sends a "gratuitous arp" message. The message indicates an Address Resolution Protocol (ARP) transmittal from the new master switch. The message contains the MAC address of the current master switch. The gratuitous arp is enabled by the **ip gratuitous-arps** command in global configuration mode. It cannot be enabled on a single interface and block it on other interfaces.

## Q. How do I synchronize configurations over the CSS in failover mode?

A. In order to synchronize configurations in software version 4.0, use the **commit config sync** command. In order to synchronize configurations in software version 3.10 code, you must use FTP in order to move the configuration from one switch to another. In order to synchronize configurations in software versions 6.x and 7.x code, use the command **commit\_redundancy** for active/standby or box-to-box redundancy. Or you can use the command **commit\_vip\_redundancy** for virtual IP (VIP)/interface redundancy. You can use the command **show script commit\_redundancy** in order to view in the header of the script the available command-line options for the **commit\_redundancy** script. The same applies to the **commit\_vip\_redundancy** command.

## Q. What settings should I use in a terminal program?

A. Use these settings:

- ◆ 9600 baud
- ◆ 8 bits
- ◆ No parity
- ◆ 1 stop bit
- ◆ No flow control

## Q. Is there a way to reprogram the MAC address on a CSS?

A. Yes, there is a way.

**Note:** You can find the MAC address and serial number on the back of the unit.

Complete these steps in order to reprogram the serial number and MAC address. This example is for a MAC address in the CS800 chassis:

1. Open **Offline Diagnostic Monitor (ODM)**.
2. In the ODM main menu, press **Shift-T** in order to reach the Technician menu.
3. Choose **1** (Configure).
4. Choose **5** (Set Manufacture Information).
5. Choose **2** (Set Backplane Manufacture Information).
6. Follow the prompt and enter the data that corresponds, such as the serial number and MAC address. You can find this data on the top of the CS800 chassis.
7. Reboot the box.

## Q. How do I make a permanent prompt change on the CSS?

A. Log in to the CSS box as user fred, and use your login credentials. In order to make a permanent prompt change, issue this command:

```
Css100#prompt Redsox
<cr>
Redsox#
```

Issue this command to save the change:

```
Redsox#save_profile
```

This command saves the user profile so that each time the user logs in, the CSS uses the same prompt. This action, similar to use of the `?.?` resource files in UNIX, creates a unique profile for each user.

When you go back to the CSS and log in as admin, the prompt does not reflect these changes. The changes are user-specific, so you need to issue the **prompt** and **save\_profile** commands for each user who wants to have the prompt reflect the new change.

## Q. What is the difference between operational and locked Flash?

A. This example shows the different types of Flash that the **show version** command displays:

```
CSS150-2#show version
```

```
Version:                ap0401049s (4.01 Build 49)
Flash (Locked):        3.10 Build 33
```

```
!--- This image is the original image that was installed on the CSS.
!--- The image serves as a backup in the event that the CSS is not able
!--- to boot from the operational Flash because of an image corruption.
```

```
Flash (Operational):   5.00 Build 10-
```

```
!--- This is the image that currently runs on the CSS.
```

```
Type:                  PRIMARY
Licensed Cmd Set(s):   Standard Feature Set
                       Enhanced Feature Set
                       SSH Server
```

## Q. Why are there different versions of Flash?

A. Locked Flash shows the version of software that was originally installed on that CSS. The version remains the same and serves only as a backup. The version in operational Flash is the version that currently runs on that CSS.

## Q. Why can I not access the management port of the CSS from a remote port?

A. In all versions of Cisco WebNS that are earlier than 5.03, the management port is not a routable interface. In version 5.03, you can add a default gateway to the management port in order to make the port a routable interface.

## Q. Does Cisco Technical Support support custom-script keepalives that the customer writes?

A. No, Cisco Technical Support does not support keepalive scripts that a customer writes.

## Q. How do I remove core files from the CSS disk?

A. If, after you issue the **show core** command, you find a list of core files, you can remove the files in one of two ways:

**Note:** The method you use depends on the version of code.

```
◆ CSS50-1(config)#llama
   !--- This command places the CSS in debug mode.
```

```
CSS50-1(debug)#clear core corefilename
```

or

```
◆ CSS50-1(config)#llama
   !--- This command places the CSS in debug mode.
```

```
CSS50-1(debug)#dir c:/Core/?
```

```
!--- This command lists the names of all the core
!--- files in the c:/Core directory.
```

```
CSS50-1(debug)#ap_file delete c:/Core/ corefilename
```

```
!--- This command deletes the specified core file.
```

## Q. When I authenticate to a RADIUS server with my CSS, I get the "RADIUS-4: RADIUS Authentication failed with reason code 2" error message. What does the message mean?

A. This error message indicates that the reply has reached CSS and there is a problem. A failure to set the service-type attribute to administrative on the RADIUS server can be the cause of the problem. Check the RADIUS server and verify the service-type attributes.

## Q. How large is the sticky table, and what causes the removal of entries?

A. The CSS has a 32,000 or 128,000 (which depends on the model type and memory available) sticky table that contains entries for **sticky source-ip** and sticky Secure Socket Layer (SSL). The sticky table does not maintain sticky cookies on the CSS. The removal of entries in the sticky table on the CSS occurs in these situations:

- ◆ By default, with a FIFO method. Entries remain in the table until the 32,000 or 128,000 buffer is full. At this time, any new entries cause the CSS to remove an entry on the basis of FIFO.
- ◆ **sticky-inact-timeout** minutes. In a content rule, you can specify the inactivity timeout by which the CSS removes a sticky entry, as this example shows:

```
owner arrowpoint
; content l5sticky
; vip address 192.1.1.1
; add service test1
; add service test2
; protocol tcp
; port 443
; url "/*"
; advanced-balance ssl
; application ssl
; sticky-inact-timeout 9

!--- Entry removal occurs after 9 minutes.

; active
```

**Note:** The CSS rejects the next sticky request in a case when all these items are true:

- ◇ The **sticky-inact-timeout** parameter is used.
- ◇ The CSS has filled the 32,000 or 128,000 buffer.
- ◇ No entries are about to timeout.
- ◆ Content rule. With the suspension and reactivation of a content rule, the removal of sticky table entries that apply to that rule occurs.

For more information, refer to the document *Configuring Sticky Parameters for Content Rules*.

## Q. How can I take a service out of rotation?

A. With the configuration of the content rule (Layer 3, Layer 4, or Layer 5) as a basis, the CSS behaves differently with the manual suspension of a service, which takes a server out of service. Many times, web developers need to temporarily suspend a service and make administration changes to the web pages. Because these web changes can occur during production hours, you do not want to kill connections that exist to the service or services when the manual service suspension occurs. Perform the updates to a service during the manual service suspension.

This example shows sample Layer 5, Layer 4, and Layer 3 content rules:

```
owner REDSOX
  content layer5
  vip address 200.200.200.200
  add service test
  add service test1
  add service test2
  protocol tcp
  port 80
  url "/*"

!--- This is a Layer 5 rule.

  active
content layer4
  vip address 200.200.200.200
  add service test
  add service test1
  add service test2
  protocol tcp

!--- This is a Layer 4 rule.

  port 80

!--- This is a Layer 4 rule.

  active
content layer3
  vip address 200.200.200.200

!--- This is a Layer 3 rule.

  add service test
  add service test1
  add service test2
  active
```

The CSS diverts the connections that exist when the content rules are either Layer 3 or Layer 4. If the suspension of a service under a Layer 3 or Layer 4 content rule occurs, the CSS diverts any connection that exists and forwards all subsequent TCP requests to the active service under that respective content rule.

With the manual suspension of a service that resides under a Layer 5 content rule, the CSS resets any or all connections that associate with that service.

## Q. Is network proximity part of the enhanced feature set?

A. Network proximity features are not part of the enhanced feature set and require an additional license. If you try to issue **proximity** commands on the CSS without the appropriate license, you receive this error message:

```
CSS50-1(config)#proximity db 0 tier1
                        ^
%% Invalid License to execute command.
This command belongs to the Proximity Database. Refer
to the user manual or contact Cisco Systems, Inc for
further information concerning license keys.
```

In order to purchase a license, see your local Cisco reseller. If you purchased a license and need a replacement, send an email to [licensing@cisco.com](mailto:licensing@cisco.com).

## Q. What details does the show dos command provide?

A. Cisco CSS can display details about the most recent attack events, which include:

- ◆ Source and destination IP addresses
- ◆ The event type
- ◆ Total occurrences

If multiple attacks occur with the same Denial of Service (DoS) type and source and destination address, there is an attempt to merge them as one event. This merge reduces the display of events.

Issue the **show dos** command in order to display:

- ◆ The total number of attacks since the boot of the CSS
- ◆ The types of attacks and the maximum number of these attacks per second
- ◆ The first and last occurrence of an attack

This example shows the output from the **show dos** command:

```
CSS50-1#show dos
Denial of Service Attack Summary:
Total Attacks: 0
SYN Attacks:           0 Maximum per second:           0
LAND Attacks:          0 Maximum per second:           0
Zero Port Attacks:     0 Maximum per second:           0
Illegal Src Attacks:   0 Maximum per second:           0
Illegal Dst Attacks:   0 Maximum per second:           0
Smurf Attacks:         0 Maximum per second:           0

No attacks detected
```

This list provides a brief description of each of the fields that the command displays:

- ◆ **Total Attacks** The total number of DoS attacks that were detected since the boot of the box. You can find a description of the type of attacks that appear in the list, along with the number of occurrences, below.
- ◆ **SYN Attacks** The TCP connections that a source initiates but that are not followed with an acknowledgement frame in order to complete the three-way TCP handshake.
- ◆ **LAND Attacks** Any packets that have identical source and destination addresses. The CSS does not allow internal IP addresses to be the source address of a flow.

Also, the CSS does not allow the source and destination addresses of frames to be equal.

- ◆ **Zero Port Attacks** Frames that contain source or destination TCP or User Datagram Protocol (UDP) ports that are equal to zero.

**Note:** Older SmartBits software can send frames that contain source or destination ports equal to zero. The CSS logs them as DoS attacks and drops these frames.

- ◆ **Illegal Src Attacks** Illegal source addresses.
- ◆ **Illegal Dst Attacks** Illegal destination addresses.
- ◆ **Smurf Attacks** Pings with a broadcast destination address. The CSS does not allow directed broadcasts by default. A **Smurf Attack** uses an Internet Control Message Protocol (ICMP) echo to a broadcast address. The CSS can block access to UDP echo ports via Access Control Lists (ACLs).
- ◆ **Maximum per second** The maximum number of events per second. Use the `maximum-events-per-second` information to set Simple Network Management Protocol (SNMP) trap threshold values.

**Note:** The maximum number of events per second is the maximum per Small Form Factor Pluggable (SFP). For a CSS 11800, for example, which can have up to four SFPs, the maximum rate per second can be as high as four times the number that appears in the display.

**Note:** Another FAQ asks if you can disable DoS protection on the CSS. The answer is no. The DoS protection is part of the flow admission process. The intention of DoS protection is to protect the resources in the CSS as well as the servers behind the CSS. DoS is not a configurable item. The intention is for DoS to be transparent when protocols work correctly. The flow setup process deeply involves the DoS features. The features help the CSS conserve fast path resources and protect devices that the CSS reaches. The features are always present in software version 3.0 and later.

Also consider the setup of certain SNMP traps for the detection of possible DoS attacks. The available traps are:

- ◆ **snmp trap-type enterprise** In order to enable SNMP enterprise traps and configure trap types, issue the **snmp trap-type enterprise** command. Issue the **no snmp trap-type enterprise** command in order to disable all traps. You must enable enterprise traps before you configure an enterprise trap option. You can enable the CSS to generate enterprise traps when DoS attack events occur, a login fails, or a CSS service transitions state.
- ◆ **dos\_attack\_type** Generates SNMP enterprise traps when a DoS attack event occurs. One trap generation occurs each second when the number of attacks during that second exceeds the threshold for the DoS attack-type configuration. The options are:
  - ◇ **dos-illegal-attack** Generates traps for illegal addresses, either source or destination. Illegal addresses are:
    - Loopback source addresses
    - Broadcast source addresses
    - Loopback destination addresses
    - Multicast source addresses
    - Source addresses that you ownThe default trap threshold for this type of attack is one per second.
  - ◇ **dos-land-attack** Generates traps for packets that have identical source and destination addresses. The default trap threshold for this type of attack is one per second.
  - ◇ **dos-ping-attack** Generates traps when the number of pings exceeds the threshold value. The default trap threshold for this type of attack is 30 per

second.

**Note:** This option does not track pings of death DoS attacks.

- ◇ **dos-smurf-attack** Generates traps when the number of pings with a broadcast destination address exceeds the threshold value. The default trap threshold for this type of attack is one per second.
- ◇ **dos-syn-attack** Generates traps when the number of TCP connections that a source initiates but that are not followed with an acknowledgement frame to complete the three-way TCP handshake exceeds the threshold value. The default trap threshold for this type of attack is 10 per second.

## Q. Can I turn off the Denial of Service (DoS) protection feature on the CSS line of switches?

A. In the current line of software for the CSS (Cisco WebNS), there is no option to disable the DoS protection feature.

## Q. Can I turn off the Denial of Service (DoS) protection counters?

A. There is no option to disable the counters that log DoS/SYN attacks.

**Note:** For more information on DoS and SYN attacks, see the response to the FAQ What details does the **show dos** command provide?.

## Q. How do I use port ranges in access lists?

A. The use of port ranges in an Access Control List (ACL) helps simplify the number of ACLs that you configure, given a situation in which you want to block user access for some TCP/User Datagram Protocol (UDP) ports. For example, suppose that you want to block ports 20 through 23 for all users who come into the box from outside your network. First, assume that the outside network or public side of the CSS is in VLAN 2. Also assume that the internal or server side of the network is on VLAN 1. The ACL configuration is:

```
acl 1
  clause 10 deny any any destination range 20 23

  !--- This clause blocks.

  clause 20 permit any any destination any

  !--- This clause allows everything else.

apply circuit-(VLAN2)
acl
  clause 10 permit any any destination any
apply circuit-(VLAN1)
```

## NetPro Discussion Forums – Featured Conversations

Networking Professionals Connection is a forum for networking professionals to share questions, suggestions, and information about networking solutions, products, and technologies. The featured links are some of the most recent conversations available in this technology.

## Related Information

- **End of Sale Announcement for the Cisco CSS 11000 Series**
  - **Cisco CSS 11000 Series Content Services Switches Bulletins**
  - **CSS 11000 Series Content Services Switches Technical Support**
  - **Software Center (Downloads) – Content Networking ( registered customers only)**
  - **Technical Support & Documentation – Cisco Systems**
- 

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2007 – 2008 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

---

Updated: Apr 19, 2006

Document ID: 25901

---