Optimizing Application Performance with Cisco Application Control Engine 4710 Appliance Content Compression

The Cisco® ACE 4710 Application Control Engine represents the next generation application switch for maximizing the availability, security, and acceleration of data center applications. This document details the HTTP compression acceleration technology of the Cisco ACE.

Challenge: Performance Bottlenecks

Mission critical business applications such as SAP Enterprise Portal, Siebel CRM, Microsoft Outlook Web Access and custom J2EE applications are increasingly found at the core of today's Enterprise IT infrastructure and are often accessed by an ever more widespread and mobile workforce. However, network latency and bandwidth constraints outside the control of IT staff often present a major hurdle to Organizations that deploy web applications accessed by remote users. Bottlenecks such as these degrade application performance resulting in reduced productivity as well as increased capital and operational costs.

Cisco ACE 4710 HTTP Compression

Compression of HTTP content for web applications improves end user experience for applications accessed over the WAN by reducing the amount of data sent from the web server to the client. Thanks to a reduction in the amount of data sent to the client, the time to transfer the response to the client decreases. This significantly improves performance for web applications accessed via

congested or lossy networks as well as for clients with restricted bandwidth such as 56Kbps dial-up access.

Figure 1. Cisco ACE 4710 accelerates web applications over the WAN

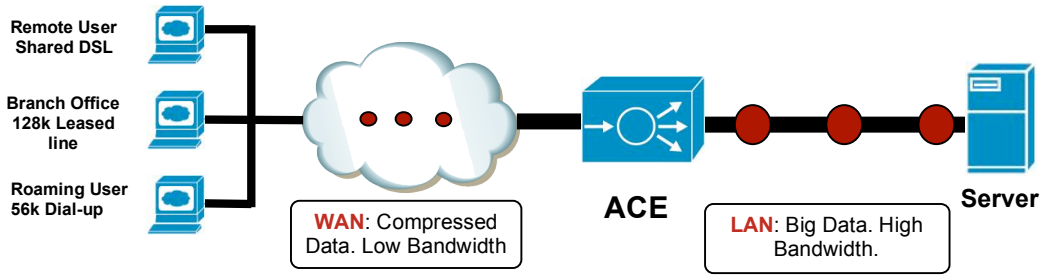

Benefits of Compression on Cisco Application Control Engine (ACE) 4710 Appliance

- **Improved End User Experience:** The reduced number of packets associated with compressed data has a significant impact on reducing application response times over the WAN. By applying compression to web applications clients accessing these applications from a 56 kbps dial-up links experience a 400% improvement in application response times. For example 100 Kilobytes of data requires 14 seconds to download over dial-up versus 3 seconds with compression.

Figure 2. Time in seconds to transfer 100KB for common bandwidth-contraints.

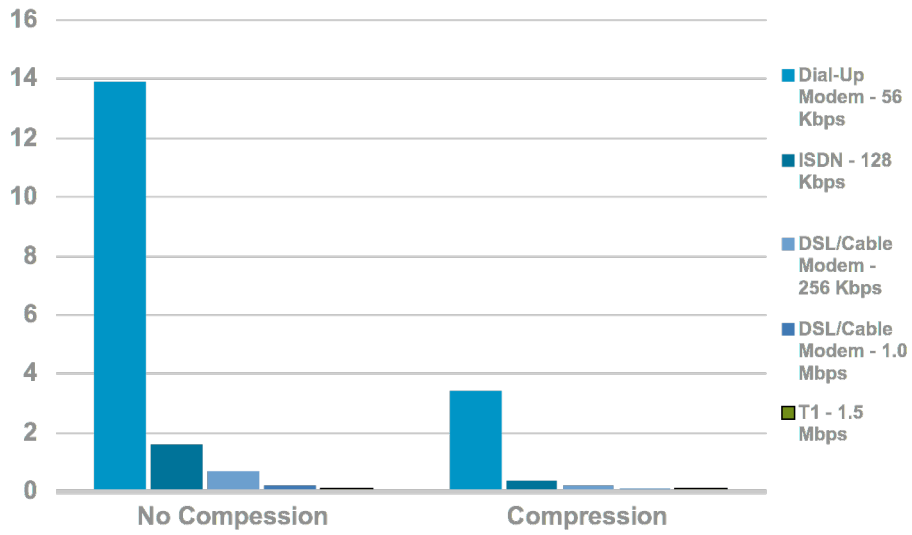

- **Data Reduction**: Compressing web content cuts bandwidth costs by reducing the size of the data delivered to the client by as much as 70 to 90%. Cisco ACE 4710 HTTP

compression reduces the size of the HTTP payload delivered to the web browser and as a result the overall amount of data transferred across WAN links drops by a factor of 3 to 4. This decrease in data sent downstream increases the effective available bandwidth allowing more transactions to be carried across the same links.

Figure 3. ACE 4710 HTTP compression applied to a 100 KB HTML file.

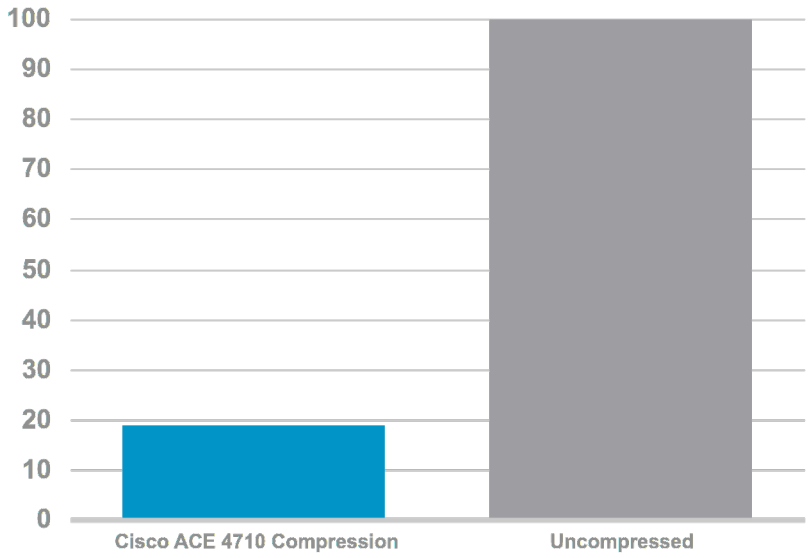

- **Server Offload**: Enabling compression is fairly easy to do on modern web servers and its use has become widespread. However this optimization comes at the cost of additional CPU utilization on the web server. The Cisco ACE 4710 enables web servers to scale to even greater traffic demands by offloading CPU intensive compression operations from the web server onto purpose-built hardware based compression on the Cisco ACE 4710.

  Figure 4 below demonstrates the result of offloading the workload of compressing HTTP responses onto the ACE 4710. With an Apache server using mod_deflate to compress 1,000 requests per second for a 70 KB page, the CPU is 80% utilized. As shown in Figure 4, offloading compression onto the ACE 4710 reduces the server CPU utilization by 70%.

Figure 4. Web server CPU utilization with ACE 4710 compression offload.

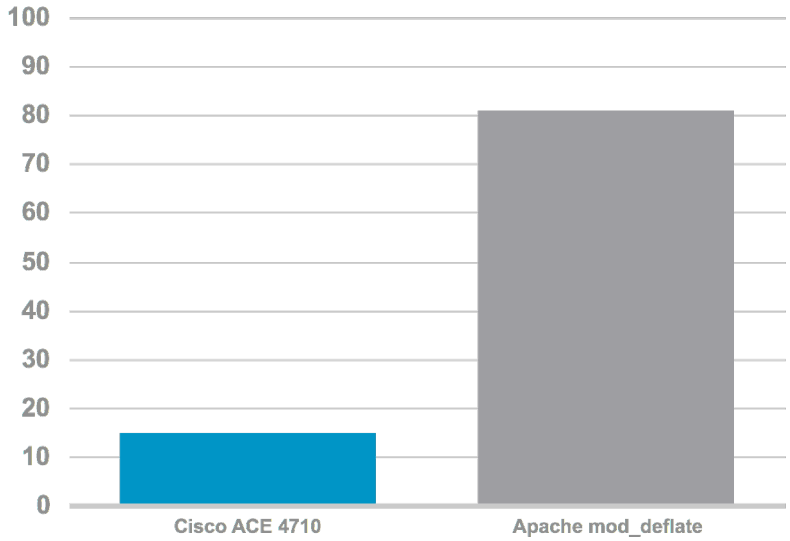

- **Asymmetric Solution:** Cisco ACE 4710 accelerates web applications for users on the WAN without the need for client side or server side changes by leveraging native compression support in web browsers such as Microsoft Internet Explorer and Firefox. ACE does not require a second device to accelerate traffic to WAN users.

- **Scalability**: The Cisco ACE 4710 delivers up to 1 Gbps of hardware accelerated data compression in addition to providing application availability and security for real world deployments.

- **Flexibility**: “Try before you buy” and “pay as you go” licensing. Cisco ACE 4710 provides for 100 Mbps throughput of compressed traffic on the base appliance with an optional upgrade license to increase throughput to 500 or 1000 Mbps of compression.

- **Ease of Use:** HTTP Compression can be enabled with a single click using the Cisco ACE 4710 Device Manager.

How Does Cisco ACE 4710 HTTP Compression Work?

HTTP compression is defined in the protocol specifications of HTTP 1.1 in RFC 2616 (http://tools.ietf.org/html/rfc2616). HTTP 1.1 allows for two different encoding methods of the compressed HTTP payload. The Cisco ACE 4710 provides support for both:

- Deflate, the data format for compression described in RFC1951 (http://tools.ietf.org/html/rfc1951)

- Gzip, the file format for compression described in RFC1952 (http://tools.ietf.org/html/rfc1952)

The majority of modern web browsers on the market today, such as Internet Explorer, Mozilla, Netscape, Firefox and Opera have implemented support for HTTP compression and accept both gzip and deflate compression methods. By compressing web application server responses ACE can deliver optimized web content to a wide variety of end users without having to make any changes to either the clients web browsers or web servers.

Traffic flow with Cisco ACE 4710 Compression

The following figure outlines the process of how the server's response to a request made by a web browser is compressed by the Cisco ACE 4710 appliance. Each step in the diagram is explained below.

Figure 5. ACE 4710 compression traffic flow

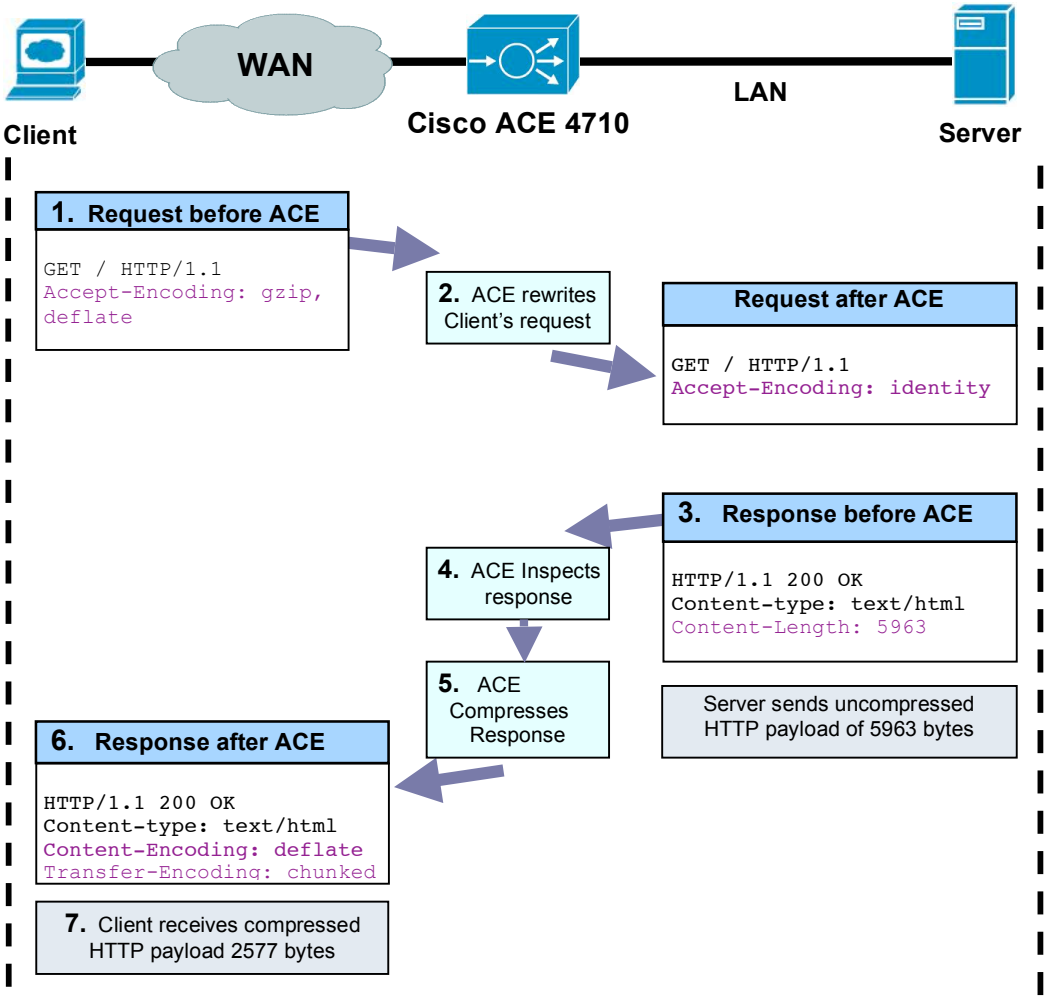

Step 1. An HTTP/1.1 client that supports HTTP compression will alert the Cisco ACE 4710 which method is supported by sending the "Accept-Encoding:" HTTP header in an HTTP/1.1 request for content. The example below demonstrates an HTTP GET request by an

Internet Explorer web browser for a web page served by a web server load balanced by Cisco ACE 4710.

```
GET / HTTP/1.1
Accept: */*
Accept-Language: en-us
Accept-Encoding: gzip, deflate
User-Agent: Mozilla/4.0 (compatible; MSIE 6.0; Windows NT
5.1; SV1; .NET CLR 1.1.4322)
Host: www.foo.com
Connection: Keep-Alive
Cache-Control: no-cache
```

Step 2. The Cisco ACE 4710 receives the HTTP request from the web browser and load balances it to the most available and least loaded web server in the server farm. The client's HTTP GET request is modified by replacing the "Accept-Encoding: gzip, deflate" header with "Accept-Encoding: Identity". This prevents the web server from sending a compressed response to Cisco ACE 4710.

```
GET / HTTP/1.1
Accept: */*
Accept-Language: en-us
Accept-Encoding: identity
User-Agent: Mozilla/4.0 (compatible; MSIE 6.0; Windows NT
5.1; SV1; .NET CLR 1.1.4322)
Host: www.foo.com
Connection: Keep-Alive
Cache-Control: no-cache
```

Step 3. The web server receives the load balanced GET request and sends an uncompressed HTTP response with 5963 bytes of data to Cisco ACE 4710.

```
HTTP/1.1 200 OK
Cache-Control: private
Content-Type: text/html; charset=UTF-8
Server: fws
Content-Length: 5963
Date: Wed, 06 Feb 2008 00:42:19 GMT
```

Step 4. Cisco ACE 4710 parses the HTTP response from the web server to verify the following conditions are true:

1. The response is an "HTTP/1.1 200 OK" response. An HTTP 1.0 response is not compressed by Cisco ACE 4710. ACE does not compress responses for HTTP 1.0 clients.

2. The response is configured for compression by verifying that the "`Content-Type:`" matches the "mimetype" compression parameter configured on Cisco ACE 4710. By default all content types that match "text/.*" are considered eligible for compression.
3. The server response must be larger than the minimum size which is 512 bytes by default. This is configurable by changing the 'minimum-size' compression parameter on Cisco ACE 4710.

Step 5. Once Cisco ACE 4710 determines that the server response is eligible for compression, it can compress the HTTP response using one of the compression methods listed by the "Accept-Encoding:" header sent in the client's GET request. Almost all modern web browsers support Gzip and or Deflate and Cisco ACE 4710 uses Deflate as the default compression method if the client supports both methods.

Step 6. Cisco ACE 4710 compresses the response and sends it to the client. Cisco ACE 4710 notifies the web browser that the data is compressed by sending the "Content-Encoding:" header in the HTTP response to the web browser.

```
HTTP/1.1 200 OK
Cache-Control: private
Content-Type: text/html; charset=UTF-8
Content-Encoding: deflate
Server: fws
Date: Wed, 06 Feb 2008 00:42:19 GMT
Transfer-Encoding: Chunked
```

Step 7. The web browser decompresses the payload of the HTTP response using the method specified by the HTTP response header "Content-Encoding: deflate". The "Content-Length:" header in the web servers HTTP response to Cisco ACE 4710 shows that the response is 5,963 bytes in size. When this value is compared to the actual size of the uncompressed HTML file rendered in Internet Explorer of 2,577 bytes we see that HTTP compression reduced the size of the HTTP payload by roughly 60%.

What type of files should be compressed?

Static & Dynamic HTML: ACE 4710 HTTP compression improves the delivery of static and dynamic HTML files common to all web applications. Dynamic HTML files with extensions such as *.asp, *.aspx, *.cfm, *.dhtml,*.html,*.jsp,*.shtml, *.php, etc as well as other text based responses such as XML (*.xml) benefit from compression with a 70-80% reduction in size.

Objects Embedded in HTML: It can also be beneficial to compress scripts (*.js) and style sheets (*.css). Certain version of Microsoft Internet Explorer can experience difficulty with compression of these file types. To prevent problems with these browsers ACE does not compress these files by default.

Note: Internet Explorer versions 5.5 (SP1 and SP2) and 6.0 (SP1) have several known bugs in this area. Microsoft has published the following related Knowledge Base articles documenting the

problem: http://support.microsoft.com/kb/313712/en-us, http://support.microsoft.com/kb/312496/en-us,http://support.microsoft.com/kb/871205/en-us and http://support.microsoft.com/kb/825057/en-us

Images: Images (*.jpg, *.gif, *.png) should not be compressed with gzip/deflate because they are already compressed. Applying compression to these files types wastes processing resources and can also potentially increase file sizes since they typically do not compress. For content with large image sizes the ACE 4710 can utilize the Smart Image Optimization feature to compress JPEGs (*.jpg) and PNG (.png) image formats. For more information on the Smart Image Optimization please refer to the Cisco ACE 4700 Series Appliance Application Acceleration and Optimization Configuration Guide (http://www.cisco.com/en/US/docs/app_ntwk_services/data_center_app_services/ace_appliances/vA1_7_/configuration/app_acceleration_and_optimization/guide/aloptmze.html#wp1137836).

The following chart summarizes how ACE 4710 handles the various elements of common web applications when compression is enabled from the Device Manager.

Table 1. ACE 4710 configuration for compression of common web page elements.

Object Type	Example File Extensions	Gzip/Deflate Compressible	Content-Type	Device Manager enabled default configuration
Static HTML	◦ *.htm ◦ *.html ◦ *.shtml	◦ Yes	◦ text/html	◦ Compressed by default
Dynamic HTML	◦ *.dhtml ◦ .*php ◦ .*asp ◦ .*aspx ◦ .*jsp ◦ .*cfm	◦ Yes	◦ text/html	◦ Compressed by default
XML	◦ *.xml	◦ Yes	◦ text/xml	◦ Not compressed by default
Images	◦ *.gif ◦ *.jpg ◦ *.png	◦ No	◦ image/gif ◦ image/jpg ◦ image/png	◦ Do not apply gzip/deflate compression Use Smart Image Optimization to reduce image sizes.
Scripts	◦ *.js	◦ Yes	◦ application/x-javascript ◦ text/javascript	◦ Not compressed by default to avoid known browser bugs.
Stylesheet	◦ *.css	◦ Yes	◦ text/css	◦ Not compressed by default to avoid known browser bugs.
Downloaded Files	◦ *.pdf ◦ *.doc ◦ *.xls ◦ *.ppt	◦ Yes	◦ Application/pdf ◦ Application/word ◦ Application/excel ◦ Application/powerpoint	◦ Not compressed by default.

It is also important to note that there is a cost associated to applying compression to HTTP traffic since it takes a few additional milliseconds to carry out the compression on ACE and additional

CPU on the client to decompress the file. To determine if the benefit of compression outweighs the cost to compress and decompress you should consider the following:

- The size of the response: It is advantageous to apply compression to files greater than 1K in size. Server responses smaller than 1K are delivered in a single tcp segment and do not benefit greatly from compression. The `"minimum-size" compression parameter` controls the minimum file size ACE will compress. The default value is 512 bytes.

- The bandwidth of the connection: The benefit of compression is greater for clients with lower bandwidth speeds. Although users with 56 Kbps modems will benefit the most, compression provides a noticeable improvement even for clients with faster connections such as DSL. In contrast users accessing compressed content across high bandwidth LANs do not typically benefit from compression since throughput is not often a bottleneck in this case.

- Physical distance between the client and the server: The greater the distance between the client and server the more round trip latency is experienced. Compression of content over high latency networks improves communication efficiency by reducing the number of packets transmitted. Users accessing content across networks experiencing congestion and packet loss also benefit from the reduced number of packets. Clients accessing applications accross low latency environments such as LANs do not typically benefit from compression.

How to Enable Compression on the Cisco ACE 4710 Appliance

From the Cisco ACE 4710 Device Manager you can begin compressing HTTP traffic on Cisco ACE 4710 by clicking the “Enable Compression” command within the Virtual Server configuration for server farms. A single click enables compression for the load balancing policy configured.

Figure 6. Enabling compression from the Cisco ACE 4710 Device Manager

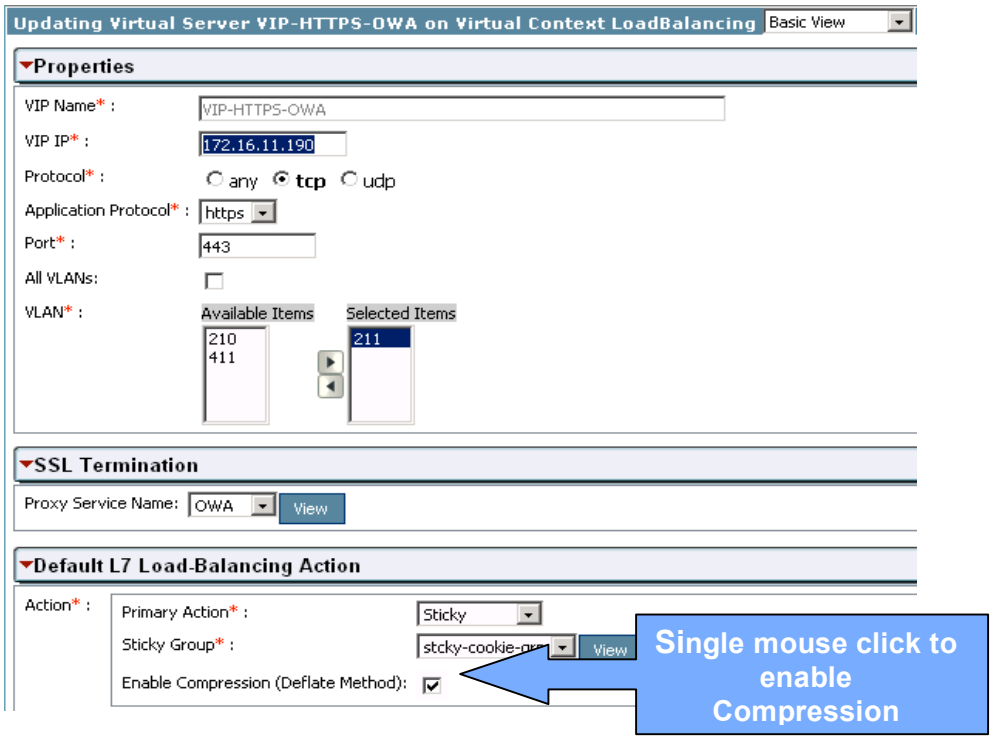

The ACE 4710's load balancing policy in the current running configuration is modified by the Device Manager to include "compress default-method deflate" in the class "class-default". For example the load balancing policy ORACLE-HTTP now contains the following:

- Load balancing policy ***before***:

```
policy-map type loadbalance first-match ORACLE-HTTP
  class class-default
    serverfarm ORACLE
```

- Load balancing policy ***after***:

```
policy-map type loadbalance first-match ORACLE-HTTP
  class default-compression-exclusion-mime-type
    serverfarm ORACLE
  class class-default
    compress default-method deflate
    serverfarm ORACLE
```

Enabling compression from the Device Manager creates an additional class-map, shown below, in the Cisco ACE 4710 running configuration that is used to prevent compression of certain content types based on the file extension. The additional called 'default-compression-exclusion-mime-type" is added to the load balancing policy and prevents compression for responses matching requests for the following URLs:

```
class-map type http loadbalance match-any default-compression-
exclusion-mime-type
description "DM generated classmap for default LB compression
exclusion mime types."
   2 match http url .*gif
   3 match http url .*css
   4 match http url .*js
   5 match http url .*class
   6 match http url .*jar
   7 match http url .*cab
   8 match http url .*txt
   9 match http url .*ps
   10 match http url .*vbs
   11 match http url .*xsl
   12 match http url .*xml
   13 match http url .*pdf
   14 match http url .*swf
   15 match http url .*jpg
   16 match http url .*jpeg
   17 match http url .*jpe
```

```
18 match http url .*png
```

The Cisco ACE 4710 takes a conservative approach to applying compression whereby any response that matches the 'default-compression-exclusion-mime-type will not be compressed by ACE. This class prevents the compression of certain types of objects such as images (*.gif, **.jpg,** *.png) which are already served in a compressed format and do not benefit from gzip/deflate compression. Other objects in the list such as java-script (*.js) and cascading style-sheets (*.css) are known to cause problems for older browsers when compressed.

Verifying Compression is enabled on Cisco ACE 4710

To determine the overall reduction in amount of bytes sent from the web servers the "show service-policy" command can be used from the CLI. In addition the results are diplayed in the Cisco ACE 4710 Device Manager under the Config > Operations > Virtual Servers > Details page. Comparing the "bytes in" versus "bytes out" provides an indication of the overall compression ratio.

Figure 7. Using Cisco ACE 4710 Appliance Device Manager to verify the compression results.

```
Status     : ACTIVE
Description: -
----------------------------------------
Interface: vlan 1 251
  service-policy: client-vips

    class: VIP-APP-WEB
     VIP Address:    Protocol:  Port:
     172.16.51.201   tcp        eq    80
      loadbalance:
        L7 loadbalance policy: APP-POLICY
        VIP ICMP Reply       : ENABLED-WHEN-ACTIVE
        VIP State: INSERVICE
        curr conns       : 1         , hit count        : 18
        dropped conns    : 0
        client pkt count : 93        , client byte count: 7883
        server pkt count : 112       , server byte count: 74310
        L7 Loadbalance policy : APP-POLICY
          class/match : NO-COMPRESS
            LB action :
              -
            hit count        : 8
            dropped conns    : 0
            compression      : off
          class/match : class-default
            LB action :
              -
            hit count        : 3
            dropped conns    : 0
            compression      : on
      compression:
        bytes_in  : 7747
        bytes_out : 3575
      optimization:
        L7 optimization policy: OPTIMIZER
```

FOR MORE INFORMATION

For more information about the Cisco ACE product family, visit

http://www.cisco.com/go/ace

For more information about Application Networking Services, go to:

http://www.cisco.com/en/US/products/hw/contnetw/index.html

or contact your local account representative.

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Europe Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: +31 0 800 020 0791
Fax: +31 0 20 357 1100

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA C78-331727-01 10/06