

Cisco Application and Content Networking System (ACNS) Software *Performance Considerations of ACNS 4.2*

This product bulletin outlines basic performance considerations when deploying ACNS 4.2. Customers should work closely with their systems engineer for proper scaling and network design of their caching and ECDN solution.

Within Cisco's content-networking solutions portfolio, the Cisco Application and Content Networking System (ACNS) Software enables a variety of services that optimize delivery of Web applications and content from the network edge to ensure enhanced speed, availability, and performance for users. ACNS combines the technologies of http caching (Transparent Proxy Caching, Manual Proxy Configuration and Reverse Proxy Caching) and enterprise content-delivery network (ECDN) for accelerated delivery of Web objects, files, and streaming media from a single intelligent edge appliance, the Cisco Content Engine (CE).

ACNS is available for many different types of devices. In the enterprise environment, ACNS provides the software foundation that enables the caching and ECDN solution components to work as a cohesive system:

- Central distribution and management capabilities are provided by the Cisco Content Distribution Manager (CDM) appliance.
- Content routing capabilities are provided by either the CDM or the Cisco Content Router (CR) appliance.
- The CE provides content edge delivery. CEs use HTTP caching technology to store content at the edge based on users'

requests, and they use ECDN technology to prepopulate rich media or large files in the CEs ahead of users' requests.

ACNS also offers flexible http caching services for both Service Providers and Enterprise organizations for Transparent Proxy Caching (TPC) and Manual Proxy Configuration services at the network edge as well as Reverse Proxy Caching (RPC) in the data centers. In both these cases, the CE is deployed stand-alone or as a cluster of CEs and not normally with either CRs or a CDM.

The performance characteristics of ACNS vary depending on the number of devices and types of devices deployed as a single cohesive system. Likewise, the performance characteristics of each CE vary depending on the features/ functionality enabled on that CE. For example, a CE serving streaming-media and HTTP content will have a lower streaming-media aggregate throughput than a CE only serving streaming-media.

Note: If more than one of the three applications (HTTP Caching, ECDN, Real or WMT) is enabled, a CE-560, CE-590 or CE-7320 should be used. The CE-507 is targeted for low-end small branch offices which plan to only run one of the three applications at a time. If more than one application is enabled on the CE-507, consult your system engineer for sizing.



Disk Configuration Guidelines

ACNS 4.1 allows customers to configure four different file systems: SYSFS, ECDNFS, CFS, MEDIAFS.

CE uses some disk space to store software in use and for future upgrade and swapping purpose. This usually takes up to 5GB. So user configurable capacity is the hardware capacity minus 5GB. The amount of user-configurable disk capacity for particular features (HTTP Caching, ECDN, WMT, Real) is dependent on which features are turned on concurrently and which features to optimize on the CE.

- **SYSFS:** There must be at least one SYSFS file system. This is used by the system to store temporary files, logs, debugging information. SYSFS needs to be at least 1GB. If you have lots of HTTP caching traffic and you want to log every transaction, you need a bigger SYSFS. Typically, SYSFS needs at least 10% of “user-configurable capacity”.
- **CFS:** CFS is used by HTTP caching and Cache Preload software to store cached HTTP objects. Its size depends on your HTTP traffic working set. CFS is bounded by both user-configurable capacity and memory limitations of the CE device.
- **ECDNFS:** ECDNFS is used by ECDN to store pre-positioned streaming media content. Content in ECDNFS can be served by HTTP streaming, WMT Server or Real Server Subscriber. ECDN requires at least 2GB of ECDNFS space to run correctly. ECDNFS can be zero if ECDN is not enabled.
- **MEDIAFS:** MEDIAFS is used by Real Proxy or WMT proxy to cache streaming media content. At least 1GB should be configured for MEDIAFS. MEDIAFS can be zero if both Real Proxy and WMT Proxy are not enabled.

Note: On a CE-507 with only one disk drive, if eCDN is enabled, the maximum amount of CFS disk space that can be mounted is around 8GB. This is limited by the amount of total disk space available and the amount of memory required to run WMT, Real, and ECDN. On a CE-507 with two disk drives, the maximum CFS disk space is 19GB.

Performance Considerations of HTTP Caching

Characteristics of caching performance

There are three important parameters required when sizing a CE installation. These are:

1. Transactions-per-second required (“how much http traffic is there”)
2. Number of simultaneous connections that the cache is required to handle (*concurrent connections*)
3. Amount of disk storage to keep the objects for sufficient time to ensure a maximized hit-rate

Ensuring all three parameters are sized appropriately is the key to ensuring that the CE provides the maximum bandwidth savings (optimal hit-rate) and maximum reduction in end-to-end latency.

Some detail about each of these parameters and what they mean is below:

Transactions per Second

Transactions per Second (TPS) is sometimes referred to as requests/sec or URLs/sec. It is a measure of the number of new HTTP Transactions per Second that a cache is capable of dealing with in a second. TPS is the most important metric in sizing a cache installation. This number is directly proportional to the traffic level (mbit/sec) that a Cache is capable of handling.



TPS is sometimes measured in terms of HTTP requests per second or in terms of megabits per second. These two are directly proportional to each other:

$$\text{mbit/sec} = \frac{\text{Transactions Per Second (TPS)} \times \text{Average HTTP object size (bytes)} \times \text{bits in a byte (8)}}{\text{bits in a megabit (1,000,000)}}$$

As can be observed from this formula, the higher the average object size, the “higher” the HTTP performance in mbit/s. This formula can be useful for determining whether a cache is of an appropriate size prior to deployment. The formula requires a few additional figures—today on the internet, we observe an average object-size on the internet of around 8.5 kbytes. If some overhead is included for IP packetization (IP headers), TCP framing (TCP headers) along with the HTTP headers associated with a request, we end up with an average around 10 kbytes, so we end up with every megabit of HTTP traffic is approximately equal to 10 TPS of sustained HTTP traffic.

Concurrent Connections

Put simply, the total number of “Concurrent Connections” is the total number of HTTP flows the CE can service at any single point in time.

This number can be extrapolated from the flow hold-time. That is, how long it takes to download the average HTTP object:

$$\text{Maximum Concurrent Connections Required} = \text{Transactions Per Second (TPS)} \times \text{Average HTTP Flow Hold time (sec)}$$

As an example, if the average http *flow hold-time* is 3 seconds (typically what we observe on the Internet today), and we’re servicing 150 TPS, we end up with the CE servicing an average total of 450 concurrent connections at any point in time:

$$\text{Maximum Concurrent Connections Required} = 150 \text{ (TPS)} \times 3 \text{ (seconds per request)} = 450 \text{ Concurrent Connection}$$

Generally speaking, the maximum number of concurrent connections that a CE supports is significantly higher than required.

Disk space

In order to allow a CE to function—that is, “cache” content—objects need to spend some period of time in the CE. The minimum cache storage time should be around 24 hours, preferably up to 72 hours, to maximize cache savings. Put another way, a CE can only provide cache hits when it can store objects for any period of time. The longer you can let it keep objects on disk, the larger the hit-rate it is capable of.



Sizing the disk requirements is a function of TPS, the average http object size and the expected cache byte hit rate:

Cache Storage required for 24 hours	=	$\frac{\text{Average TPS} \times \text{Average object size (bytes)} \times \text{seconds in 24hrs} \times (1 - \text{anticipated byte hit rate})}{\text{Bytes in a gigabyte}}$
-------------------------------------	---	--

Working on the assumption that we're going to see a "daily average*" of 100 TPS (~10 mbit/sec) an average HTTP object size of 9.5 kbytes and an anticipated cache hit-rate of 35%, we end up with:

*note that this calculation is based on the "daily average TPS" rate; the maximum TPS experienced in peak times may be substantially higher than this:

Cache Storage required for 24 hours	=	$\frac{100 \times 9500 \times 86400 \times (1 - 0.35)}{1,000,000,000}$	=	53.352 gigabytes
-------------------------------------	---	--	---	------------------

Thus, we'll need over 50 gigabytes of storage to sustain an 'optimal' hitrate.

This will be an over-approximation, but is preferable to under-provisioning disk-space. It is an over-approximation due to the fact that:

Object Popularity will tend to follow a Zipf distribution model (a small percentage of objects are extremely popular and result in the majority of the hit-rate), and the CE itself will perform object expiration based on a hybrid Least-Recently-Used (LRU) and Least-Frequently-Used (LFU) algorithm meaning that unpopular objects are expired before popular objects.

Effects of Turning Caching Features ON

When caching features such as transaction logging, rules, content filtering (SmartFilter, Websense, N2H2), authentication, ECDN or streaming are turned on in the same CE, performance will be impacted. The performance of Websense and N2H2 is more dependent on the latency of the Websense and N2H2 Server responding to the CE (i.e. link speed and network topology between Websense/N2H2 Server and the CE and the performance of the Websense/N2H2 Server). Authentication is also dependent on the latency that the AAA Server responds to the CE, especially for NTLM. For RADIUS, LDAP, and TACACS+, authentication can be stored in the CE and reduces the dependence on the AAA response, therefore, the performance will be close to the TPC performance of the CE.

ACNS 4.1 also supports Secure Computing's SmartFilter V3.0.2 Server and Client code running on the CE. There will be performance impact in the CE-560, CE-590, and CE-7320 when SmartFilter is turned on. The performance of the CE-507 is more limited by the available disk space than SmartFilter.

A CE can provide both TPC and RPC services simultaneously. In such a deployment, the maximum achievable throughput will be somewhere between that of individual TPC and RPC services. The rationale behind RPC performance exceeding TPC is that the typical hit-rate is *higher*—that is, more content can be serviced from the cache rather than having to contact the origin web-server. A CE can service cache-hits more efficiently than serving a cache-miss.



Performance Characteristics of ECDN Solution Components

Enterprise Content Delivery Networks (ECDNs) are made up of two or three types of devices:

- Central distribution and management capabilities are provided by the Cisco Content Distribution Manager (CDM) appliance. The CDM may also provide Content Routing (CR) functionality.
- The CE provides the actual content edge delivery. CEs use transparent caching technology to store content at the edge based on users' requests, and they use ECDN technology to prepopulate rich media or large files in the CEs ahead of users' requests.

In larger deployments, Content Routing functionality may be provided by dedicated Cisco Content Router (CR) appliance(s). CRs are optional; if no CRs are deployed, the CDM will provide this functionality.

There are a large number of permutations in terms of the types of devices required to build an ECDN, a large number of topologies on which an ECDN is expected to function and different types of content and request-patterns for content.

With all these variables, sizing an ECDN solution can become a complex task. The following sections should be used as a guide for determining if (a) the ECDN will meet the customer's performance expectations, (b) the ECDN deployment is realistic given the customer's network topology and (c) how many devices will be required (eg. for Content Routing) to handle the customer's traffic-load / request-patterns.

(1) Maximum number of CEs

The maximum number of CEs in an ECDN deployment is dependent on the network configuration. This is a function of the number of distinct content *channels*, whether Self-Organizing Distributed Architecture (SODA) is used, whether there are dedicated CRs, the number of CEs behind firewalls relative to the CDM/CRs, etc.

A CDM-4650 with ACNS 4.1.1 has been tested in a non-SODA environment (all devices were direct 'children' of the CDM), also known as the STAR topology, with slow 56 kbps links. This is the worst case scenario for the number of CEs that the CDM can support. All the CEs remained up over a period of days. At the time of testing, no files were replicated to the CEs. The number of CEs tested with the STAR topology is 350 with 1 channel, no replication turned on the CDM-4650 and no CRs. This number is expected to be impacted when the number of channels and number of assets increase.

The method that the CDM keeps the CEs in-sync is to send an update of the list of other CEs in the network, a list of all channels, and if subscribed to a channel, a list of all files in that channel. (This is regardless of whether SODA is enabled or not). The message is initiated from the CE by sending, to the CDM, the same list as known by the CE. The CDM compares the CE's list, and if different, will send a new list. Therefore, when there are 300 CEs, and one more CE is added, every CE will send a list to the CDM of the existing 300 CEs. The CDM will compare and send an update of the 301 CEs. In this example, the CDM will issue 300 messages in a 5 minute window.

If the user adds or deletes a channel to the CDM, every CE in the network will receive a new update of the lists of channels. For files within each channel, only the CEs subscribed will receive the list. If there are 1000 files in a channel, and one file is deleted, all the CEs subscribed will receive a new list of 999 files.



Therefore, how many CEs can be subscribed in a network depends on how many channels are in the network, how many files per channel, how many CEs are subscribed to each channel, and how often the network changes (such as files being added and deleted).

When multiple CR4430s are used for Content Routing, the maximum number of CEs that can be supported does not change substantially. The positive impact is that Content Routing can be offloaded from the CDM, allowing the CDM GUI to have better response time, and also this is a high availability and load-balanced configuration for http redirection. The negative impact is that more communication between CDMs, CRs and CEs occur.

The CDM-4630 has been tested to support 200 CEs in a STAR topology with 1 channel, no replication and no CR. The CDM-4630 does not have RAID controller or tape backup support, less CPU processing power (lower http redirection performance—see results in the next section), less memory (256MB in CDM-4630 vs. 1GB in CDM-4650), less internal disks (2x36GB drives in CDM-4630 vs. 8x18GB drives in CDM-4650) and less external disk expandability (SA-6 only for CDM-4630 vs. SA-12 for CDM-4650). Therefore, the CDM-4630 is recommended mainly for small to medium ECDN deployments.

Note: The maximum number of CEs supported is only a limitation for ECDN deployments for ACNS 4.1. For HTTP caching deployments, no CDM is required and there is no limit to the number of CEs supported.

(2) Maximum number of content channels

No more than 100 content channels can be configured in ACNS 4.1. Performance impacts will incur based on the number of channels defined. Refer to the section “Content Routing Performance Characteristics” for more details.

A general rule is to allocate channels sparingly. Each channel has its own administration policies, storage-space, and replication bandwidth settings so each additional channel results in additional administrative network traffic during CE-to-CDM communications.

(3) Maximum number of objects

No more than 10,000 objects in total should be prepopulated into ACNS 4.1 via the CDM. No more than 6,000 objects should exist in a single channel with ACNS 4.1. The implication of the number of files in a single channel is that the size of the packets sent to the CEs is increased. The greater the number of files in a single channel, the greater the file list sent to every CE subscribed to that channel. When only one file is deleted from the large list, it is again sent out to all the CEs subscribed, consuming network bandwidth. No more than 3,000 objects should be imported into the CDM at once.



(4) Topology constraints

Given the type of deployment, we have two choices for how to enable content-distribution within a topology: either to use STAR (manual override) or SODA (automatic topology discovery).

(4.1) STAR Topology

A “STAR” topology is where SODA has been disabled and all CEs use the CDM as a parent. This is probably the most typical configuration that will be seen for ECDN deployments. One of the many factors for implementing the ECDN is to push the content to the edge and save valuable WAN bandwidth, especially when the connection is slow, e.g. 64 kbps/ISDN or 56kbps frams relay. Any customer with low bandwidth networks would prosper from the Cisco ECDN solution. An example of such a topology is shown in Figure 1 below.

(4.2) SODA Topology

The ideal network topology for SODA to work is where there is a hierarchy (tree) of CEs, of which only some CEs parent off the CDM and other CEs naturally discover upstream CEs to parent off. An example of such a topology is shown in Figure 2 below. The decision to which CE becomes the child of another device is based upon round trip time (RTT). The distributed content-distribution enabled by SODA allows it to scale to higher numbers of edge-nodes versus that of a STAR topology.

Figure 1. STAR topology—note that all CEs parent off the centralized CDM


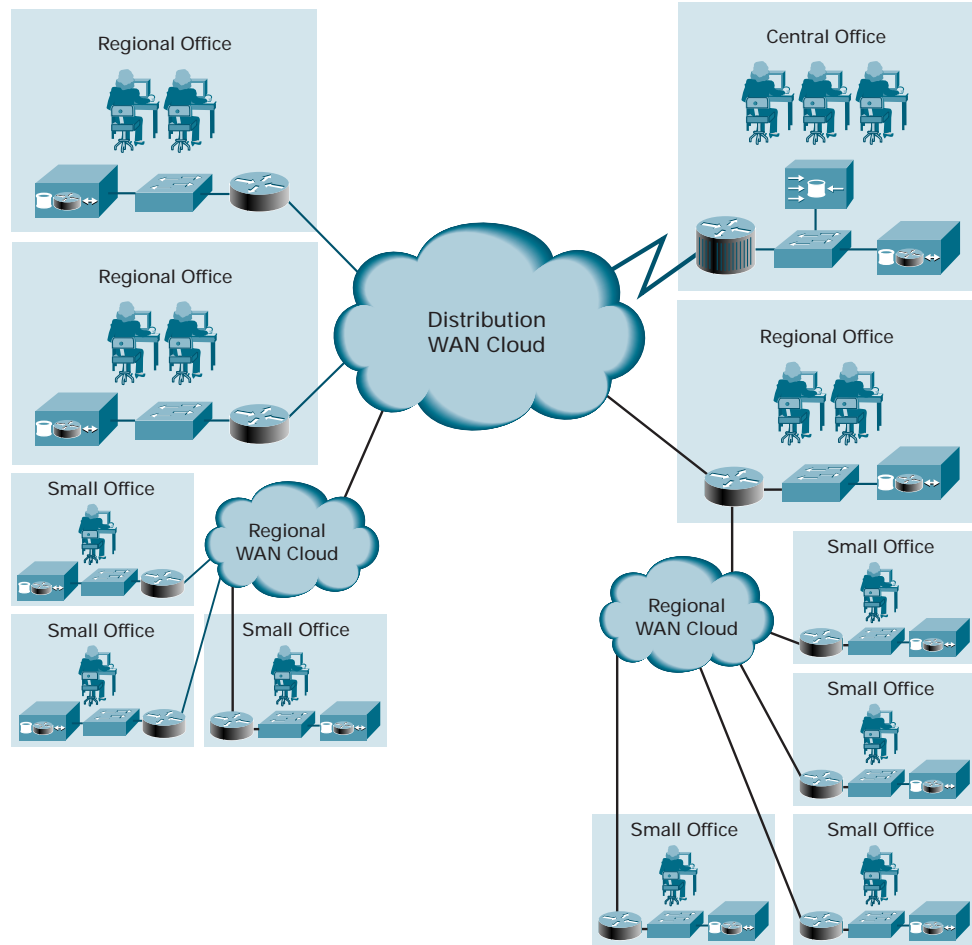




Figure 2. Ideal topology for SODA—a natural hierarchy exists between central, regional and small branch offices



The following performance characteristics apply when using a STAR topology:

Table 1 Performance characteristics when using a STAR topology in ACNS 4.1

CDM Model	Link Speed characteristics	Recommended Maximum number of nodes
CDM-4630	56 kbps	200
CDM-4650	56 kbps	350

Note: Test configuration: CDM is connected to CEs via 56kbps links and only 1 channel is configured per CE and no CR is included. Performance will be impacted with a larger number of channels and objects in the channel.

If there are some CE nodes behind firewalls/NAT, performance will be impacted since additional administrative network traffic may result for a CE to discover the network topology between different CEs and CDM.



(5) Content Routing Performance Characteristics

Content Routing is an integral part of any ECDN deployment. It is the basis for how requests for content are serviced. The CDM and CR make intelligent redirections based on the source-IP address of the client request and the content channel information in the URL. When content is imported into the CDM for distribution, a URL is created to point to the CDM or CR. The CDM or CR then directs the client to the best edge CE for fulfilling the request.

The number of redirections per second (RPS) will depend on several factors: number of channels, Content Services Switch (CSS) configuration, number of devices in the network, and network latency. The testing results detailed below are the maximum limits reachable within a flat network with little network latency.

To achieve maximum performance the CSS must be changed from default settings. The configuration of the CSS will directly impact the results received. When configuring the CSS, the load balance option "Least-Connections" is recommended.

The number of channels defined on the CDM, directly impact the performance results. Highest performance is seen when the number of channels is less than 10. The performance of request per second will decrease as the number of channels increases. The testing showed that with "least-connections" configured on the CSS, one CDM and CR, and one channel, the number of RPS is 41. When the only variable changed is the number of channels, from one channel to 100 channels, the performance decreases to 14 RPS.

What this testing does not take into account is the performance impact of including multiple CEs, asset replication (sending content to the CEs), etc. This testing was using one CE-507. The type of CE is irrelevant to the testing because the generation script counts the number of HTTP 302 message received and does not contact the CE for content.

The CE was configured with a coverage zone of 0.0.0.0. This method ensured that all the UNIX machines would be redirected to the CE, regardless of their IP address. The CE is also required to be subscribed to the channel that is referenced during the test. The channel used is 'Default_Channel'.

(6) AV TV-Out Playlist Constraints

A maximum of 500 entries are allowed in a single playlist. (A playlist itself may consist of both MPEG-1/MPEG-2 video/audio clips and On-Screen-Display overlay bitmap files). A maximum of 100 entries may be added to the playlist at a single time. A video/audio clip is limited to 2 gigabytes in size. To play a video longer than this, it needs to be split into 2GB chunks first.

(7) Miscellaneous Constraints and Characteristics

The maximum length of a path and filename is 256 characters. Note that this includes the name of the content channel.

Note that some web-browsers and FTP clients impose smaller limits on this length, particularly when transferring multiple times at the same time. The maximum size of an object imported into the CDM with ACNS 4.1 is (2 gigabytes - 1 byte), or 2,147,483,647 bytes.

CISCO SYSTEMS



Corporate 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: 408 526-4100

European Headquarters
Cisco Systems Europe
11, Rue Camille Desmoulins
92782 Issy-les-Moulineaux
Cedex 9
France
www-europe.cisco.com
Tel: 33 1 58 04 60 00
Fax: 33 1 58 04 61 00

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: 65 317 7777
Fax: 65 317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the

Cisco Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia
Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland
Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland
Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden
Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2002, Cisco Systems, Inc. All rights reserved. CCIP, the Cisco Powered Network mark, the Cisco Systems Verified logo, Cisco Unity, Follow Me Browsing, FormShare, Internet Quotient, iQ Breakthrough, iQ Expertise, iQ FastTrack, the iQ logo, iQ Net Readiness Scorecard, Networking Academy, ScriptShare, SMARTnet, TransPath, and Voice LAN are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, Discover All That's Possible, The Fastest Way to Increase Your Internet Quotient, and iQuick Study are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, Fast Step, GigaStack, IOS, IP/TV, LightStream, MGX, MICA, the Networkers logo, Network Registrar, Packet, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, SlideCast, StrataView Plus, Stratm, SwitchProbe, TeleRouter, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0206R)

Printed in the USA