

Cisco ACE GSS 4492R Global Site Selector

Product Overview

The Cisco® GSS 4492R Global Site Selector is part of the Cisco ACE Application Control Engine family and a crucial component of any data center or cloud computing architecture that requires an appliance-based, security-focused, universal global load balancer.

The Cisco GSS 4492R allows businesses to deploy global Internet and intranet applications with the confidence that all application users will be quickly routed to a standby data center if a primary data center outage or overload occurs. Further, they can be assured that global traffic management and load balancing of these same applications will be maintained per their business rules. The Cisco GSS 4492R does this by performing as an intelligent Domain Name Service (DNS) authoritative server, such that when a user's web browser send a request such as "where is www.bxb.com" they receive the best of all the possible global answers for that particular user at that particular time (the best DNS "A record" IP address).

The Cisco GSS 4492R traffic-management process continuously monitors the load and health of any SNMP-capable device (such as a server load balancing device or particular server) within each data center. The Cisco GSS 4492R uses this information in conjunction with customer-controlled load-balancing algorithms to select the best data center, server complex, or server destinations that is available and not overloaded, within user-definable load conditions, in real time. In this manner, the Cisco GSS 4492R intelligently selects best destination to ensure application availability and performance for any device that uses common DNS request for access.

While the Cisco GSS 4492R can monitor health and performance of almost any server or Server Load Balancing/Application Delivery Controller devices using ICMP, TCP, HTTP-header, and SNMP probes, when used with Cisco's load balancing/ADC offerings it has enhanced monitoring capabilities. To retrieve more granular monitoring data, that is securely transmitted in a timely manner, the Cisco 4492R can make use of the special KAL-AP monitoring interface built into the Cisco Application Control Engine (ACE) Modules and Cisco ACE 4710 appliances, the Cisco CSS 11500 Series Content Services Switches (CSS) and the Cisco Content Switching Module (CSM) for Cisco Catalyst® 6500 Series Switches.

For users this means: (1) higher application availability due to Cisco GSS 4492R providing failover across distributed Data Centers, (2) better application performance due to Cisco GSS 4492R optimization of load growth to multiple Data Centers, servers and switches.

For IT operators this means the Cisco GSS 4492R adds agility by automating reactions to changes in Local and Global Networks to ensure application availability & performance. If a network outage occurs, the Cisco GSS 4492R can automatically or under administrative control direct clients to a disaster-recovery site within seconds. The Cisco GSS 4492R also adds security & intelligence to the DNS process by offering cluster resiliency that can be managed as a single entity. The Cisco GSS 4492R also offers a unique use of other Cisco technologies to mitigate the effects of a DNS-based distributed-denial-of-service (DDoS) attack.

Figure 1. Cisco ACE GSS 4492R Appliance



By offloading or replacing the DNS server resolution process of traditional DNS servers for those applications requiring global load balancing and / or business resiliency, the Cisco GSS 4492R adds a new level of DNS self-defense, simplifies the DNS infrastructure, optimizes global site selection, boosts DNS responsiveness, helps ensure data center availability, and increases the scalability of Websites and data centers. The Cisco GSS 4492R is a crucial component for enterprises and service providers deploying globally distributed data centers, installing disaster-recovery solutions, or looking for a way to consolidate and strengthen the DNS architecture of a standalone data center.

Features and Benefits

The Cisco GSS 4492R offers the following benefits:

- Provides a scalable, dedicated hardware platform to help ensure that applications are always available, by detecting site outages or site congestion and performing global traffic management for client requests
- Provides a unique self-defending resilient architecture that is crucial for disaster recovery and for multisite Web application deployment
- Offers site persistence for e-commerce applications
- Scales to support hundreds of data centers or SLB devices
- Offloads, augments and optimizes the DNS infrastructure by taking over the domain resolution process for content requests and delivery for all types of static and dynamic Web content, and processes the request-responses at thousands of requests per second
- Offers flexible, heterogeneous support for all Application Networking Services (L4-7) devices and DNS-capable networking products, including third-party load balancers
- Improves the global data center selection process by offering user-selectable global load-balancing algorithms along with universal SNMP load and health probes
- Tightly integrates with Cisco SLB/ADC devices without sacrificing the ability to work in a heterogeneous environment of DNS-capable networking products
- Offers two unique network proximity features that use Cisco routers and Layer 4 through 7 content switches to allow the Cisco GSS 4492R to direct content consumers to the closest data center in real time
- Runs a hardened, non-DNS caching, GSLB specific code base so the Cisco GSS 4492R is not affected by common caching server or Berkeley Internet Name Domain (BIND) services vulnerabilities
- Provides centralized command and control of the DNS resolution process for direct and precise control of the global load-balancing process
- Supports a Web-based GUI and DNS wizard to simplify GSLB command and control
- Supports role-based access control (RBAC) and operation to limit access to Cisco GSS 4492R functions and features

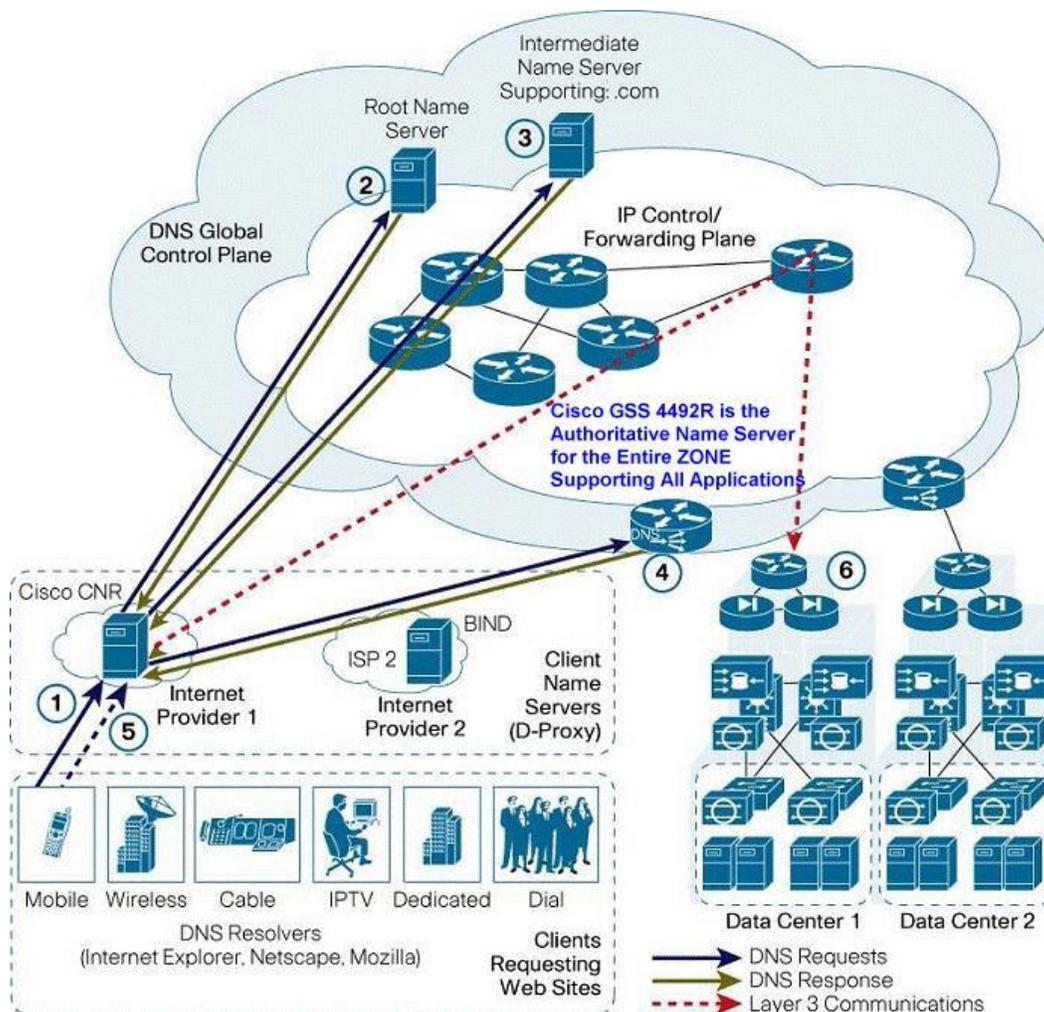
- Supports configuration using a flat text file, the command-line interface (CLI), and the GUI
- Supported by Cisco Application Networking Manager for unified operations management with Cisco ACE, CSS, and CSM

How the Cisco GSS 4492R Performs Its Major Tasks

The Cisco GSS 4492R performs two major functions as part of the global site selection process:

- Takes an active role in the DNS infrastructure to connect the client to the SLB device that supports the requested Website
- Continuously monitors the load and availability of these SLB devices to select the SLB device most capable of supporting the new client

Figure 2. Cisco GSS 4492R Site Selection Process



In Figure 2, the Cisco GSS 4492R provides GSLB services for the entire DNS zone (Cisco.com). The Cisco GSS 4492R can continuously and simultaneously monitor the load and health of over 3,000 devices / IP addresses using ICMP, TCP, HTTP and SNMP, as well as thousands more Virtual IP addresses served by Cisco ACE, CSS and CSM. These devices can be located together or at disparate remote or standalone data centers.

How the Cisco GSS 4492R interacts with the client in the data center selection process is summarized in the following six steps (corresponding to the numbers in Figure 2):

- Step 1. A client wants to access an application at Cisco.com (Web, e-mail, VPN, etc.). The resolver (client) sends a query for Cisco.com to the local client DNS server (D-proxy). In this case, the Cisco Network Registrar client acting as the D-proxy could be a Cisco GSS 4492R running the Cisco Network Registrar software. This could also be a server running Berkeley Internet Name Domain (BIND) services or a DNS server appliance from one of many available vendors.
- Step 2. The local D-proxy does not have an IP address for Cisco.com, so it sends a query to a root name server. The root name server can respond to the request in two ways. The most common way is to send the D-proxy directly to the authoritative name server for Cisco.com. In the other method, iterated querying (shown in Figure 2), the root name server sends the D-proxy to an intermediate name server that knows the address of the authoritative name server for Cisco.com.
- Step 3. The local D-proxy sends a query to the intermediate name server, which responds, referring the D-proxy to the authoritative name server for Cisco.com.
- Step 4. When the local D-proxy sends a query to the authoritative name server for Cisco.com, the name server responds with the IP addresses of the two Cisco GSS 4492R devices, and tells the D-proxy to ask the Cisco GSS 4492R for the IP address for Cisco.com or www.cisco.com.
- Step 5. The local D-proxy sends its final request directly to one of the two Cisco GSS 4492R devices. The Cisco GSS 4492R is authoritative for the Cisco.com sub domain, so it will be responsible for sending the “A record” response, the IP address, to the D-proxy. The Cisco GSS 4492R sends the best IP address for that requester at that time — in this case, the IP address that is a Virtual IP on the server load balancer device at Data Center 1.

In order to send the best IP address to the D-proxy, the Cisco GSS 4492R applies intelligence to selecting its response. Following are examples of this intelligence, which is not supported by generic DNS servers. The Cisco GSS 4492R:

- Intelligently manages client traffic flow to each data center, routing users to the closest, least loaded, or otherwise selected “best answer” based on any of the ten global load balancing algorithms that can be applied; Cisco GSS 4492R will not send an IP address to the D-proxy if the device is overloaded
- Automatically routes users to an alternative data center if the primary data center / device becomes unavailable; Cisco GSS 4492R will not send an IP address to the D-proxy if the device (Website, Application Server, etc.) is unavailable
- Where appropriate, replies with a Virtual IP address, not a real IP address of a back-end server, thus taking advantage of integration with the local load balancing services

- Step 6. The DNS global load balancing process is complete; the client is directed to the server load balancer device at Data Center 1 and the client communicates directly with the Virtual IP on the server load balancer at Data Center 1.

Universal, Security-Focused Advanced DNS Services

Business Resiliency When Combined with Local (Server) Load Balancers

The Cisco GSS 4492R, in combination with local Server Load Balancers (SLBs), provides an outstanding solution for large enterprises and service providers planning to deploy highly reliable distributed data centers. The Cisco GSS 4492R selects the best site based on the global load and availability information supplied by the SLB; the SLB then selects the best local server within the data center based on availability and local load. The Cisco GSS 4492R

simplifies the network deployment architecture with its centralized command and control features. For example, the Cisco GSS 4492R can gracefully take a Cisco application switch out of rotation without affecting ongoing operations.

Global Traffic Management

The Cisco GSS 4492R can be deployed as a standalone global traffic manager that globally load balances client requests across distributed data centers using network performance metrics such as content use, round-trip time (RTT) between client and the closest data center, routing topology, and any device performance values that are available through SNMP.

Global Load-Balancing Algorithms for Complete Site-Selection Control

The Cisco GSS 4492R supports 10 global load-balancing algorithms and gives administrators complete flexibility in selecting the global load-balancing algorithm that meets their needs. For instance, administrators can choose among the following algorithms:

- **Ordered list:** This user-definable list specifies one or a group of IP addresses (corresponding to a virtual IP address or the IP address of a back-end server) that the Cisco GSS 4492R uses to respond to a DNS request for a specific domain. The Cisco GSS 4492R uses the first address in the list until it becomes unavailable or overloaded; it then moves to the next address in the list. This process is repeated for every subsequent entry in the list.
- **Static algorithm based on client's DNS address:** This algorithm is a variation on the ordered list that allows the administrator to map the IP address of the client's DNS server to an available virtual IP address on a specific content switch. This feature is used when the administrator wants to allocate a specific community of users to a specific set of SLB devices or back-end servers.
- **Round robin:** This algorithm cycles through available virtual IP addresses in order. The round-robin balancing method is useful when balancing requests among multiple, active data centers that are hosting identical content—for example, SLB devices at primary and active-standby sites that serve requests.
- **Weighted Round Robin (WRR):** The Cisco GSS 4492R cycles through the list of available virtual IP addresses as requests are received, but send requests to a favored virtual IP address based on a user-assigned weighting value.
- **Least loaded:** The Cisco GSS 4492R can receive load values from the Cisco ACE, CSS, or CSM device via a secure, proprietary interchange or from other devices via SNMP MIB variables. The Cisco GSS 4492R monitors these load values to see if they exceed a threshold that is assigned by the administrator. If the load exceeds the specified threshold, the IP address or Virtual IP (VIP) address on the monitored device is considered offline and unavailable to serve requests. If the load falls below the threshold, the Cisco GSS 4492R automatically starts sending requests to the IP or VIP address. An administrative option to “hold down” a device considered offline is available to protect against undesirable network “flapping”.
- **Geo database:** External geographic databases can be loaded onto the Cisco GSS 4492R to allow the Cisco GSS 4492R to send a client request to the closest data center based on the source IP address of the DNS request. This database can scale up to 500,000 entries.

Self-Defending DNS Protection

Cisco GSS 4492R can also be run with integrated DNS-focused DDoS protection. This optional software uses a subset of the unique Multi-Verification Process (MVP) architecture found on the Cisco Guard DDoS Mitigation Appliances. This optional software handles only DNS-related attacks and does not have the performance or full feature set of the Cisco Guard DDoS Mitigation Appliances, but instead is matched to the operational performance levels normally performed by the Cisco GSS 4492R. The following DDoS mitigation capabilities are included:

- Filter
- Rate limitation per D-proxy with learning during normal operation
- Spoofing prevention through cookie insertion

The filter detects the following:

- Rapid DNS queries for the same domain (replay attack and DoS) from a specific source IP
- Broadcast IP addressing as source IP
- Multicast IP addressing as source IP
- Empty IP addressing as source IP
- Cisco GSS 4492R IP addressing as source IP
- Invalid IP range (209.165.202.128 and 209.165.202.159)
- Malformed DNS packets
- Rapid DNS queries for domains not configured on the Cisco GSS 4492R

The unique rate-limiting software establishes the DNS process rate baseline for each DNS server sending requests to the Cisco GSS 4492R during normal operations and profiles these rates. If any DNS server exceeds these normal rates, the Cisco GSS 4492R, according to the rate-limit policy set by the network administrator, will start to rate limit these DNS requests. Therefore, a compromised DNS server will not be allowed to consume all the DNS processing capabilities of the Cisco GSS 4492R.

The Cisco GSS 4492R also can insert a cookie into TCP-based DNS requests, using TCP port 53. This action allows the Cisco GSS 4492R to mark the various DNS servers communicating with the Cisco GSS 4492R. The challenge-response algorithms are based on pseudo-random information. The Cisco GSS 4492R sends a challenge, also known as cookie, to a client that tries to connect with the Cisco GSS 4492R. If the source IP address in the packet header is the IP address that is assigned to the client, the client will receive the challenge and send back a response.

However, if the source IP address in the packet is spoofed, the client that generated the original traffic to the zone will not receive the Cisco GSS 4492R response and therefore will not answer with the correct challenge. The Cisco GSS 4492R considers clients as authenticated only when they return the correct challenge.

Performance and Scalability

Highly scalable, the Cisco GSS 4492R meets the needs of the most demanding environments.

Table 1 lists performance and scalability metrics information for the Cisco GSS 4492R.

Table 1. Cisco 4492R Performance and Scalability Metrics

Feature	Scalability
DNS Requests per Second	<ul style="list-style-type: none"> • 28,000-30,000 tested maximum sustained for simple, single VIP configuration • 20,000-23,000 tested maximum sustained rate for moderately complex, 1000+ VIPs, configuration • 12,000-13,000 tested maximum sustained rate for most complex, 1000's of VIPs, maximum scale configuration
Name Server Forwarding Requests per Second	1500
Active Server Load Balancers	256
DNS Rules	4000
Hosted Domains	2000 (maximum 1000 per SLB)
Hosted Domains Character Count / Length	128

Feature	Scalability		
Hosted Domain List	2000 (maximum 500 per list)		
Virtual IP Addresses	4000		
Source IP Addresses Configurable for DNS Rules	500		
Source Address Lists	60 (maximum 30 members per list)		
Answer Groups	2000		
Answers per Answer Group	Group Type	Maximum Answers	
	VIP	100	
	Name Server	30	
	CRA	20	
Name Server Addresses for NS Forwarding	100		
KeepAlive (KAL) Limits These are the device monitoring probes that check for health and load. Standard KAL for a minimum polling interval of 40 seconds between checks, whereas Fast can be polled as often as every 4 seconds. KAL-AP, the Cisco-specific advanced probe combines the Load and VIP online status of up to 1,000 VIPs per monitored device, thereby dramatically increasing the scaling and granularity of monitoring when GSS is used with Cisco ACE, CSS or CSM.	KAL Type	Fast	Standard
	TCP	150	1500
	HTTP Head	100	500
	ICMP	150	750
	Scripted (SNMP)	120	384
	KAL-AP	40	128
	CRA (DNS Race)	200	
	Name Server	100	

System Requirements

Table 2. System Specifications

Table 2 describes the Cisco GSS 4492R system specifications

Description	Specification
Processor	Pentium CPU
Memory	2GB RAM (fixed)
Storage	One 80-GB hard drive
Image	Software Image SF-GSS-2.0-K9 or higher
Power	Single Integrated AC power (autosensing 110V/60Hz)

Table 3. Electrical Specifications

Table 3 describes the Cisco GSS 4492R AC electrical specifications

AC Specification	Operational
Input Voltage (V)	100 to 240 VAC
Input Voltage (F)	50 to 60 Hz
Current Draw	3.5 Amps
Power Consumption	345 Watts
Heat Dissipation	1314 BTU/hr

Table 4. Environmental Specifications

Table 4 describes the Cisco GSS 4492R environmental specifications

Specification	Operational	Non-operational
Temperature	32°F to 104°F (0° to 40° C)	-4°F to 140°F (-20° to 60° C)
Humidity	90% at 104°F (40° C) (non-condensing)	90% at 140°F (60° C) (non-condensing)
Shock	31 G halfsine	71 G halfsine, 20 G square
Vibration	0.25 G from 3 to 200 Hz	0.5 G from 3 to 200 Hz
Acoustic Noise	50 dBa (maximum)	Not applicable

Table 5. Physical Specifications

Table 5 describes the Cisco GSS 4492R physical specifications

Description	Specification
Rack Units	One Rack Unit
Chassis Dimensions (H x W x D)	1.50 in x 16.92 in x 20.04 in (42.4 mm x 430 mm x 509 mm)
Chassis Shipping Weight	40.0 lbs. (18.2 kg)
Chassis Weight	30.8 lbs. (14 kg)

Table 6. Port Specifications

Table 6 describes the Cisco GSS 4492R port specifications

Description	Specification
Ports	<ul style="list-style-type: none"> • 1000BASE-T, 100BASE-T, 10BASE-T (dual) Ethernet ports • Console serial port (for out of band management)

Warranty Information

Find warranty information on Cisco.com at the [Product Warranties](#) page.

Ordering Information

The Cisco GSS 4492R supports only AC power.

Orderable part numbers are listed in Table 7

To place an order, visit the [Cisco Ordering Home Page](#).

To download software, visit the [Cisco Software Center](#).

Table 7. Ordering Information

Product Name	Part Number
Cisco GSS 4492R Global Site Selector	GSS-4492R-K9
Cisco Global Site Selector Software 3.0	SF-GSS-V3.0-K9
DDoS Mitigation Software	SF-GSS-DDOSLIC

Cisco Services

Cisco Services make networks, applications, and the people who use them work better together.

Today, the network is a strategic platform in a world that demands better integration between people, information, and ideas. The network works better when services, together with products, create solutions aligned with business needs and opportunities.

The unique Cisco Lifecycle approach to services defines the requisite activities at each phase of the network lifecycle to help ensure service excellence. With a collaborative delivery methodology that joins the forces of Cisco, our skilled network of partners, and our customers, we achieve the best results.

Cisco Services can provide you with guidance and support in the design, deployment and configuration of your DNS infrastructure, Load Balancing, and Business Resiliency including ACE 4400 Series Global Site Selector Appliances.

For More Information

For more information about the Cisco GSS 4400 Series Global Site Selector Appliances, visit <http://www.cisco.com/en/US/products/hw/contnetw/ps4162/index.html> or contact your local Cisco account representative.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

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.

 CCDE, CCENT, CCSI, Cisco Eos, Cisco HealthPresence, Cisco IronPort, the Cisco logo, Cisco Nurse Connect, Cisco Pulse, Cisco SensorBase, Cisco StackPower, Cisco StadiumVision, Cisco TelePresence, Cisco Unified Computing System, Cisco WebEx, DCE, Flip Channels, Flip for Good, Flip Mino, Flipshare (Design), Flip Ultra, Flip Video, Flip Video (Design), Instant Broadband, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn, Cisco Capital, Cisco Capital (Design), Cisco-Financed (Stylized), Cisco Store, Flip Gift Card, and One Million Acts of Green are service marks; and Access Registrar, Aironet, AllTouch, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Lumin, Cisco Nexus, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, Continuum, EtherFast, EtherSwitch, Event Center, Explorer, Follow Me Browsing, GainMaker, iLINUX, IOS, iPhone, IronPort, the IronPort logo, Laser Link, LightStream, Linksys, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, PCNow, PIX, PowerKEY, PowerPanels, PowerTV, PowerTV (Design), PowerVu, Prisma, ProConnect, ROSA, SenderBase, SMARTnet, Spectrum Expert, StackWise, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website 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. (0910R)