



Kemp FIPS LoadMaster

Product Overview

UPDATED: 27 July 2023

© 2022 Progress Software Corporation and/or one of its subsidiaries or affiliates. All rights reserved.

These materials and all Progress® software products are copyrighted and all rights are reserved by Progress Software Corporation. The information in these materials is subject to change without notice, and Progress Software Corporation assumes no responsibility for any errors that may appear therein. The references in these materials to specific platforms supported are subject to change.

#1 Load Balancer in Price/Performance, 360 Central, 360 Vision, Chef, Chef (and design), Chef Habitat, Chef Infra, Code Can (and design), Compliance at Velocity, Corticon, Corticon.js, DataDirect (and design), DataDirect Cloud, DataDirect Connect, DataDirect Connect64, DataDirect XML Converters, DataDirect XQuery, DataRPM, Defrag This, Deliver More Than Expected, DevReach (and design), Driving Network Visibility, Flowmon, Inspec, Ipswitch, iMacros, K (stylized), Kemp, Kemp (and design), Kendo UI, Kinvey, LoadMaster, MessageWay, MOVEit, NativeChat, OpenEdge, Powered by Chef, Powered by Progress, Progress, Progress Software Developers Network, SequeLink, Sitefinity (and Design), Sitefinity, Sitefinity (and design), Sitefinity Insight, SpeedScript, Stylized Design (Arrow/3D Box logo), Stylized Design (C Chef logo), Stylized Design of Samurai, TeamPulse, Telerik, Telerik (and design), Test Studio, WebSpeed, WhatsConfigured, WhatsConnected, WhatsUp, and WS_FTP are registered trademarks of Progress Software Corporation or one of its affiliates or subsidiaries in the U.S. and/or other countries.

Analytics360, AppServer, BusinessEdge, Chef Automate, Chef Compliance, Chef Desktop, Chef Workstation, Corticon Rules, Data Access, DataDirect Autonomous REST Connector, DataDirect Spy, DevCraft, Fiddler, Fiddler Classic, Fiddler Everywhere, Fiddler Jam, FiddlerCap, FiddlerCore, FiddlerScript, Hybrid Data Pipeline, iMail, InstaRelinker, JustAssembly, JustDecompile, JustMock, KendoReact, OpenAccess, PASOE, Pro2, ProDataSet, Progress Results, Progress Software, ProVision, PSE Pro, Push Jobs, SafeSpaceVR, Sitefinity Cloud, Sitefinity CMS, Sitefinity Digital Experience Cloud, Sitefinity Feather, Sitefinity Thunder, SmartBrowser, SmartComponent, SmartDataBrowser, SmartDataObjects, SmartDataView, SmartDialog, SmartFolder, SmartFrame, SmartObjects, SmartPanel, SmartQuery, SmartViewer, SmartWindow, Supermarket, SupportLink, Unite UX, and WebClient are trademarks or service marks of Progress Software Corporation and/or its subsidiaries or affiliates in the U.S. and other countries. Java is a registered trademark of Oracle and/or its affiliates. Any other marks contained herein may be trademarks of their respective owners.

Please refer to the NOTICE.txt or Release Notes – Third-Party Acknowledgements file applicable to a particular Progress product/hosted service offering release for any related required third-party acknowledgements.

Table of Contents

1 LoadMaster Overview	6
1.1 LoadMaster Feature Summary	6
1.2 Supported Platforms	6
2 Deployment Scenarios	7
2.1 High Availability Configuration	7
2.2 Clustering	9
2.3 Geographic Server Load Balancing (GLSB)	10
3 Load Balancing on LoadMaster	11
3.1 Scheduling Traffic	11
3.1.1 Scheduling Overview	11
3.1.2 Scheduling Methods	11
3.2 Rule-based Content Switching	12
3.2.1 Terminology	14
3.2.2 Content Rewriting	14
3.2.3 Using Content Switching	14
3.3 Health Checking	14
3.4 Session Persistence	15
3.4.1 Persistence Overview	15
3.4.2 Persistence Methods	17
3.4.3 Port Following	19
3.5 Application Performance	19

3.5.1 SSL Acceleration	20
3.5.2 Caching	20
3.5.2.1 Flushing Cache	21
3.5.2.2 Maximum Cache Size	21
3.5.3 Compression	21
3.6 Authentication	22
3.6.1 Overview	22
3.6.2 ESP Features	23
3.7 Web Application Firewall	24
3.7.1 WAF Overview	24
3.7.2 WAF Benefits	25
3.8 GEO	25
3.9 FIPS 140-2 Compliance	28
3.9.1 FIPS 140-2 Level 2 Compliance	28
3.9.1.1 FIPS Processor Overheating - Detection and Prevention Measures	29
3.10 Sub-Virtual Services (SubVSs)	29
3.11 Certificates	30
3.12 User Management	30
4 Miscellaneous	31
4.1 SNMP Support	31
4.2 WUI Authentication and Authorization	32
4.3 Bonding and VLAN	32

4.4 IPsec Tunneling	33
4.5 Debugging Utilities	33
Last Updated Date	34

1 LoadMaster Overview

1.1 LoadMaster Feature Summary

Kemp leads the industry in driving the price/performance value proposition for application delivery and load balancing to levels that our customers can afford. Our products' versatile and powerful architecture provide the highest value, while enabling our customers to optimize their businesses that rely on Internet-based infrastructure to conduct business with their customers, employees, and partners.

1.2 Supported Platforms

LoadMaster is available as a physical, bare metal, or virtual appliance with support for VMware, Hyper-V, XEN, KVM, and Oracle VirtualBox. Virtual appliances are also available in the marketplace on Microsoft Azure, Amazon Web Services (AWS), and CenturyLink clouds. LoadMaster licensing options include perpetual, subscription, pay-as-you-go (PAYG), and monthly billed metered subscriptions.

2 Deployment Scenarios

2.1 High Availability Configuration

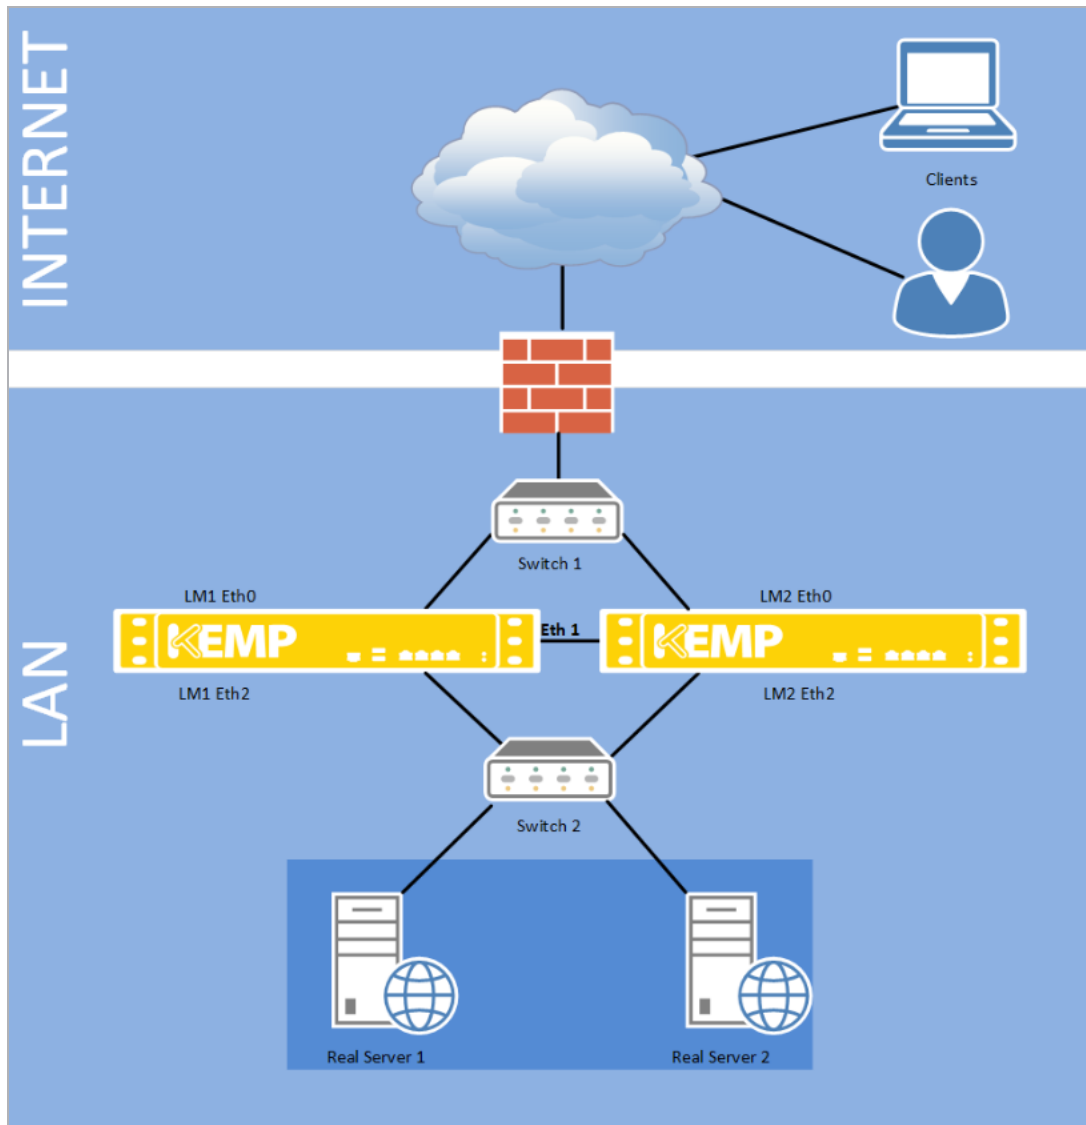
The High Availability (HA) feature of the LoadMaster guarantees the availability of your server farm. HA is achieved by a hot-standby, failover mechanism. Two identical LoadMaster units are integrated into the network as a cluster. One machine serves as the active LoadMaster and the second one remains in a standby, idle state, always prepared to take over the activities from the active server. This cluster appears as a single logical unit to the client side and to the server farm side connections.

With an HA cluster, each network interface in use has an individual IP address and one shared IP address that is shared with the partner unit. The shared IP address is identical for both LoadMaster appliances, though it is associated with only the active LoadMaster at any given time.

If the LoadMaster is set as the default gateway of the servers, remember to use the shared address of the HA pair because this address will always be available.

During normal operation, the active LoadMaster sends HA check messages to the standby peer over connected interfaces. If the active LoadMaster fails, the standby appliance becomes active and takes over the task of balancing.

The topology for HA dual arm looks like this:



Both HA1 and HA2 use eth0 to connect to the network (firewall) and eth2 for connection to the servers. The two eth0 ports have one shared IP address and the two eth2 ports have a different shared IP address. Optionally, eth1 on each unit may be directly connected using a patch cable for added HA health checking though it is quite unnecessary because there are already two health check routes between the HA pair.

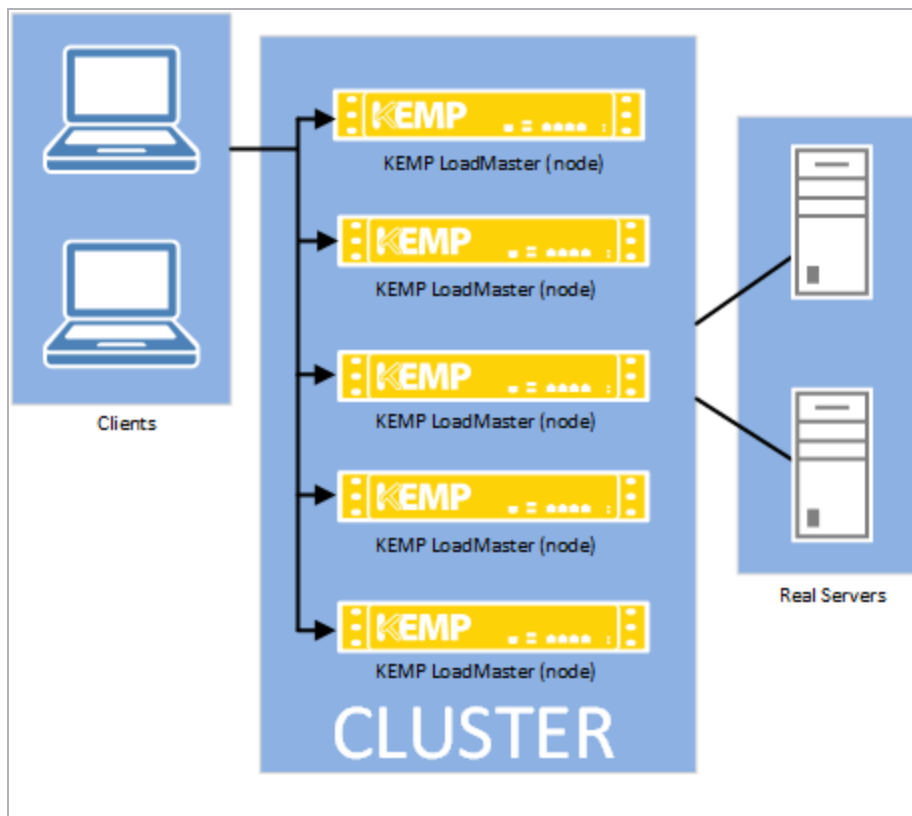
HA1 and HA2 must be on the same subnet with the same default gateway and be located within the same broadcast network. They must not be separated by an intra-site link and must use the same gateway to return traffic.

Running HA spanned across multiple subnets does not provide hardware redundancy in the event of a failure of the link between them. If traffic balancing between multiple sites is required, the GEO LoadMaster is the correct solution. This is Kemp's Domain Name Server (DNS)-based appliance that employs health checking to avoid service outages.

2.2 Clustering

Clustering involves deployment of multiple LoadMaster instances (nodes) as a single management and control domain where each node is actively passing traffic. If any of these nodes fail - the remaining nodes continue to provide service to the remote client, providing resilience.

Clustering supports scaling - LoadMaster nodes can be added to or removed from a cluster so that capacity can be statically adjusted to match business requirements. Each additional LoadMaster node that is added to the cluster can provide a linear increase in throughput and performance. The maximum number of nodes that can be added to a cluster is 16, subject to the license applied.



With LoadMaster clustering, the load-balancing capability can be extended as needed by adding additional LoadMasters for the same Virtual Service. Traffic that is intended for a failed node is either passed on to an existing node or load-balanced across the remaining nodes.

The LoadMasters all work in parallel. If a LoadMaster node fails, traffic is redistributed amongst the remaining nodes. Health checking is performed on each of the LoadMasters.

Most administrative changes should be made on the shared IP address interface. These changes are then propagated to the LoadMasters in less than a second. Local administration of the individual LoadMasters can also be performed by accessing the WUI of the units directly. The WUI of the local LoadMasters contains limited configuration options.

For further information on clustering, refer to the LoadMaster Clustering, Feature Description on the [Kemp Documentation Page](#).

2.3 Geographic Server Load Balancing (GLSB)

Kemp GEO provides GSLB functionality with the ability to move past the single data center, allowing for multi-data center HA. Even when a primary site is down, traffic is directed to the disaster recovery site. Also included in GEO is the ability to ensure clients connect to their fastest performing and geographically closest data center.

GEO can be deployed in a distributed (active/active) HA configuration, or combination of active/active HA, with multiple GEO LoadMasters securely synchronizing information. Introducing GEO into existing Authoritative DNS requires minimal integration work and risk, allowing you to fully leverage the existing DNS investment.

3 Load Balancing on LoadMaster

3.1 Scheduling Traffic

There are several load balancing methods provided by the LoadMaster, which are known as 'Scheduling Methods' or 'algorithms'. These are described in the sections below.

3.1.1 Scheduling Overview

Kemp LoadMaster supports multiple scheduling methods to enable correct operation of applications being load balanced. The simple approach is to distribute client sessions evenly across all servers. However, servers may differ in capacity due to size, existing traffic, or network performance. The scheduling methods available on the LoadMaster can accommodate these scenarios to ensure even distribution of traffic across servers.

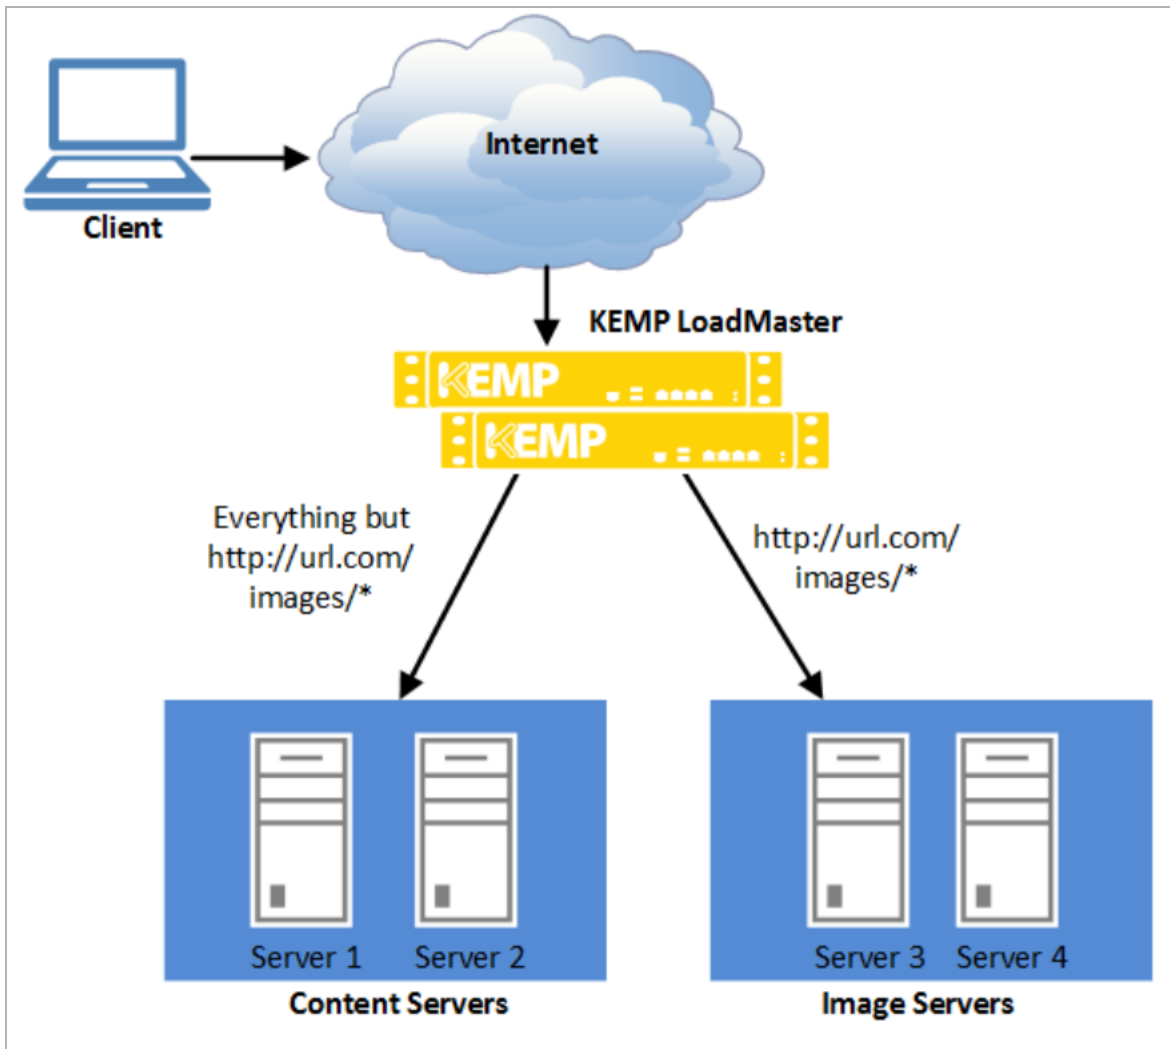
3.1.2 Scheduling Methods

Scheduling Method	Overview
Round Robin	Incoming requests are distributed sequentially across the server farm (cluster), that is, the available servers.
Weighted Round Robin	Incoming requests are distributed across the cluster in a sequential manner, while taking account of a static 'weighting' that can be pre-assigned per server.
Least Connection	Requests are distributed on the basis of the connections that every server is currently handling by the LoadMaster. The server in the cluster with the least number of active connections automatically receives the next request.
Weighted Least Connection	This method uses the ratio of the number of connections and the weight of a server. The server in the pool with the lowest ratio automatically receives the next request.
Agent-Based Adaptive Balancing	The LoadMaster can distribute traffic based on the load on the actual server, which can include CPU load, disk load, and other performance metrics. The LoadMaster Server Agent works by residing on a server and reporting back to the LoadMaster on

how loaded the server is. The metrics used to determine the system's load are entirely configurable by the user.	
Fixed Weighted	The highest weight Real Server is only used when other Real Server(s) are given lower weight values. However, if the highest weighted server fails, the Real Server with the next highest priority number will be available to serve clients. The weight for each Real Server should be assigned based on the priority among the Real Server(s).
Weighted Response Time	The traffic is scheduled using the weighted round robin method. The weights used by the weighted round robin method are calculated using the response times from health check requests.
Source IP Hash	A hash of the source IP is generated and used to find the correct Real Server. This means that the Real Server is always the same from the same host. With this scheduling method, you do not need any source IP persistence.

3.2 Rule-based Content Switching

The LoadMaster series of load balancers support content switching, which is sometimes referred to as URL switching. This allows the LoadMaster to direct specific requests to specific Real Servers based on the contents of the requested URL.



For example, if you have two groups of servers; one group to serve images and the other to serve all other content - you can create content rules to separate these two classes of request as shown in the diagram above.

Any URL that includes /images in it, such as **http://url.com/images/party.jpg** or **http://url.com/images/dogs.jpg** would be directed to server 3 and 4, while anything else would be directed to server 1 and 2.

This can be very useful if you have servers that perform different functions (application servers, static content servers, mapping servers, specialized content generation servers, and so on) that must all be served from the same general hostname (for example, www.website.com).

3.2.1 Terminology

The term content switching does not refer to the process involved with Layer 2 switching. Instead, content switching refers to switching traffic between different servers, depending upon the content requested.

3.2.2 Content Rewriting

Response body string replacement rules can be used to reduce cost and time-to-market when in-transit traffic handling requirements change for applications. For example, you may want to rewrite URLs that are in the BODY of the response and change them from HTTP to HTTPS. This may be needed if there is a third-party application that has a number of hyperlinks embedded. These response body string replacement rules can be used to modify the URLs embedded in the BODY. These rules can also be used to switch domain names, for example, replace test.yourdomain.com with yourdomain.com. Rules can be assigned to process the response body. These rules are run over the response body before it is sent to the client.

For more information, see the **Content Rules** document description on the [Kemp Documentation](#) page.

3.2.3 Using Content Switching

There are two parts to configuring content switching: the content rules, and the Virtual Service configuration. The content rules are configured globally on the LoadMaster and various rules are applied to specific Real Servers operating under a Virtual Service.

For more information, see the **Content Switching** document on the [Kemp Documentation](#) page.

3.3 Health Checking

The LoadMaster utilizes health checks to monitor the availability of the Real Servers and the Virtual Services. If one of the servers does not respond to a health check within a defined time interval, the weighting of this server is reduced to zero. This zero weighting has the effect of removing the Real Server from the Virtual Service configuration until it can be determined that this Real Server is back online. The LoadMaster performs Layer 7 health checks for the following ports

Service	Port	Protocol
FTP	21	TCP
TELNET	23	TCP
SMTP	25	TCP

HTTP	80	TCP
HTTPS	443	TCP
POP3	110	TCP
NNTP	119	TCP
IMAP	143	TCP
DNS	53	UDP
LDAP	389/636	UDP/TCP

For HTTP/HTTPS services, health checks can be based on url availability, HTTP status code, host headers, and by using alternative HTTP methods such as PUT and GET.

For other ports, the LoadMaster uses Layer 4 health checks for TCP services and Layer 3 health checks for UDP services.

Layer 3 health checks utilize ICMP-based echo requests (pings) to test whether a Real Server can be reached over the network. A Layer 3 check is not Virtual Service-specific, for example, when it fails, the corresponding Real Server is removed from all Virtual Services that use it.

In contrast to the Layer 3 health checks, service-based health checking for both the Layer 4 and Layer 7 health checks are Virtual Service-based. When a Real Server fails such a check, it is removed only from the corresponding Virtual Service – all other Virtual Services that use this Real Server are unaffected.


3.4 Session Persistence

3.4.1 Persistence Overview


Persistence – which can also be referred to as 'affinity', 'server affinity', or 'server sticky' - is the property that enables all requests from an individual client to be sent to the same healthy server in a server farm. Persistence is not turned on by default, but it is an option configurable for each Virtual Service.

Without persistence, the LoadMaster directs traffic according to the load balancing scheduling algorithm, such as round-robin, weighted round-robin, and so on (diagram below).

3 Load Balancing on LoadMaster



With persistence, the LoadMaster directs new connections according to the load balancing scheduling algorithm, but returning connections go to the same server (diagram below).



With HTTPS/SSL, there are a few things to consider. If you are not terminating the SSL session at the LoadMaster, then your only options are Source IP Address persistence or SSL Session ID persistence. Since the stream is encrypted in a non-terminated session, the LoadMaster cannot look at the HTTP headers or other Layer 7 information.

If you are terminating the HTTPS/SSL session at the LoadMaster, then any of the LoadMaster persistence options can be used. Since the HTTPS/SSL session is terminated, the LoadMaster sees all of the unencrypted traffic, and is able to look at the HTTP stream. This is true even when you are terminating the HTTPS/SSL session at the LoadMaster and then re-establishing an SSL session with the Real Servers.

3.4.2 Persistence Methods

Persistence Method	Overview
--------------------	----------

Server Cookie Persistence	A Layer 7 feature that uses existing cookies generated from the server to determine which server to send users to.
Active Cookie Persistence	A Layer 7 feature that uses cookies in the same way as Server Cookie Persistence, but with Active Cookie the cookies are generated by the LoadMaster, not the server.
Server Cookie or Source IP Persistence	The Server Cookie or Source IP setting is identical to the Server Cookie setting, but with the additional fall-back method of source IP address.
Active Cookie or Source IP Persistence	The Active Cookie or Source IP setting is identical to the Active Cookie Persistence. If, for any reason, the expected cookies are not present, then the source IP address is used to determine persistence.
Hash All Cookies Persistence	This method creates a hash of the values of all cookies in the HTTP stream. Cookies with the same value are sent to the same server for each request. If the values change, the connection is treated as a new connection. The client is then allocated to a server according to the load balancing algorithm.
Hash All Cookies or Source IP Persistence	This method is identical to Hash All Cookies, with the additional feature that it falls back to Source IP persistence if no cookies are in the HTTP string.
Source IP Address Persistence	This method uses the source IP address of the incoming request to differentiate between users.
Super HTTP	This method creates a unique fingerprint of the client browser and uses that fingerprint to preserve connectivity to the correct Real Server.
URL Hash	With this method, the LoadMaster sends requests with the same URL to the same server.
HTTP Host Header	In this method, the LoadMaster sends all requests that contain the same value in the HTTP Host: header to the same server.
Hash of HTTP Query Item	This method operates that the named item being inspected is a Query Item in the Query String of the URL. All queries with the same Query Item value are sent to the same server.
Selected Header	In this method, the LoadMaster sends all requests that contain the same value in the specified header to the same server.
SSL Session	This is a persistence method that may be used with SSL services even if they are

ID	not offloaded. This relies on the client maintaining the same SSL Session ID for a complete user session.
UDP Session Initiation Protocol (SIP)	This persistence mode is only available in a UDP Virtual Service when Force L4 is disabled. SIP uses request and response transactions, similar to HTTP.

3.4.3 Port Following

When using 'shopping cart'-like services where a user selects items and adds them to a list, any of the previous types of persistency can be used. When the user then decides to pay for the items, this is normally performed using a secure SSL (HTTPS) service. When port following is turned on, the Real Server where the 'shopping cart' connection is active is selected for the SSL session. This selection only occurs when a connection is still open from the same client (as determined by the source IP address), and if the SSL service has the same IP address as the 'shopping cart' service.

For example, if a connection is made to the HTTP service of www.somewebsite.com, and then a new SSL connection is made to the same address, then the SSL session is directed to the same Real Server as the original HTTP service.

Port following can be done between UDP and TCP connections if desired. For more information, see the **Port Following** document on the [Kemp Documentation](#) page.

3.5 Application Performance

Application Front End is a group of features that revolve around web application delivery and network optimization. The introduction of the LoadMaster Application Front-End Services (AFE) solves very core requirements by providing better bandwidth and server utilization. LoadMaster AFE Services include:

- Intrusion Prevention System (Snort-based)
- Caching
- Data Compression

Each feature can be deployed per web Virtual Service.

AFE features are license-based. If you do not have these features and you need them, contact a Kemp Sales Representative.

3.5.1 SSL Acceleration

The LoadMaster series offers SSL termination/acceleration for Virtual Services. With SSL acceleration, the SSL session is terminated at the LoadMaster.

The LoadMaster Supports SSL 3.0, TLS 1.0, TLS 1.2, and TLS 1.3.

Not all LoadMaster platforms support TLS 1.3. Contact your Kemp representative for details.

There are two primary benefits to SSL acceleration:

- The LoadMaster offloads the SSL workload off the Real Servers
- The LoadMaster performs Layer 7 processing: persistence or content switching

Without terminating the SSL session at the LoadMaster, the headers and content cannot be read so persistence cannot be done. The only consistently reliable persistence method available when the SSL session is not terminated at the LoadMaster, is Source IP.

With SSL acceleration, the LoadMaster uses a specialized processor to perform the SSL functions. With this SSL acceleration hardware, the LoadMaster handles SSL connections as easily as it handles non-SSL connections.

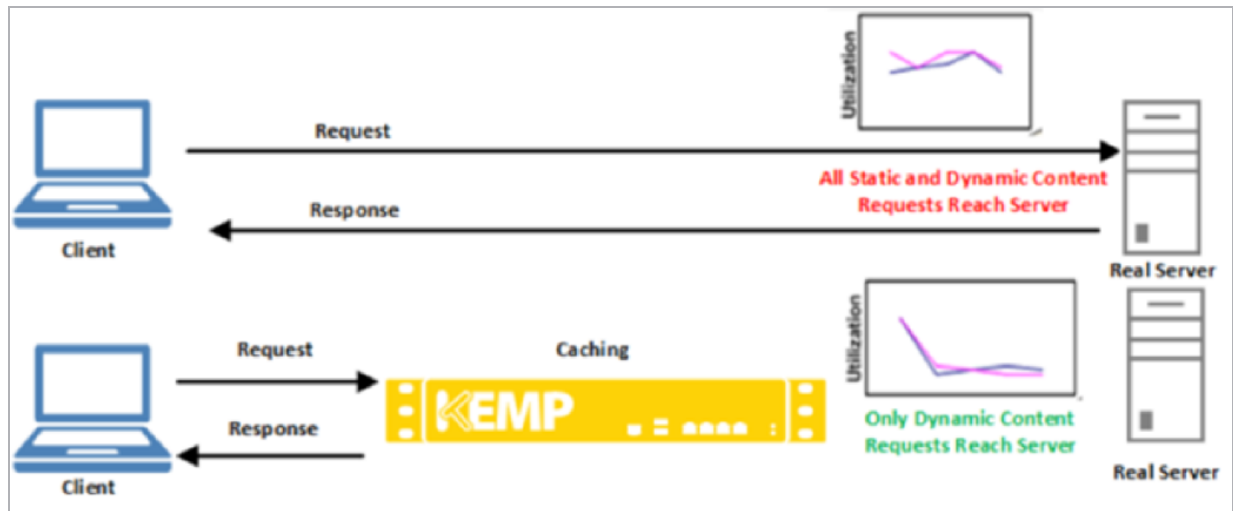
LoadMaster-FIPS supports hardware-based SSL acceleration and offloading. With hardware SSL, the LoadMaster has a separate specialized processor that handles all SSL functions. No matter what level of SSL connections, the LoadMaster's general processor is not burdened. This specialized hardware is purpose-built for SSL and can handle extremely high connection rates (TPS) of SSL traffic.

For more information on SSL, refer to the **SSL Accelerated Services, Feature Description** on the [Kemp Documentation page](#).

3.5.2 Caching

The LoadMaster advanced caching engine saves valuable Real Server processing power and bandwidth, which can be dedicated to performing critical core business application logic. Significant server performance gains can be achieved when implementing caching. Chatty protocols, such as HTTP, require frequent creating and closing of connections for fetching of static resources, creating unnecessary resource utilization on Real Server(s) and the network.

By enabling LoadMaster caching, you can re-purpose connection-related resources for more relevant business logic. By deploying LoadMaster caching, your organization can also greatly reduce web traffic to the Real Server(s), saving on bandwidth in front of your Real Server(s).



Caching is available for HTTP and offloaded HTTPS Virtual Services.

HTTP/HTTPS requests with no-cache headers bypass the cache following RFC 2616. The cache is filled in a delayed manner - allow a few seconds for static content to be cached.

In accordance with RFC 2616, URLs that contain query strings (those containing a question mark symbol (?) in the rel_path part) are not cached.

3.5.2.1 Flushing Cache

LoadMaster will not monitor file changes on the Real Server and auto-reload the cache maintained within the Virtual Service. You can force reload the cache by deselecting and selecting the **Enable Caching** check box. You can also reload a cached object, sending a non-cache request. Most browsers support this by holding the left **Shift** key and clicking reload (or pressing **F5**).

3.5.2.2 Maximum Cache Size

The amount of global memory available for caching can be configured; values have a linear relation to actual memory. To configure this, navigate to **Virtual Services > View/Modify Services > Modify > Advanced Properties** in the LoadMaster WUI.

3.5.3 Compression

The LoadMaster data compression feature reduces the amount of data to be transferred for HTTP objects by utilizing gzip compression, which is available in all modern web browsers. Leveraging

Lempel-Ziv (LZ) compression and HTTP/1.1 GNU zip (gzip) content encoding reduces bandwidth utilization for high compression files such as text files (HTML, CSS, and JavaScript).

Data compression allows the LoadMaster to compress the application payload per request, reducing network bandwidth consumption without degrading content quality and response time. The result is an improvement in the end-users' overall experience. Data compression is supported on all files. Compression ratios vary by file type.

Compression is not recommended for files 100 Mb or greater in size.

The compression feature should be deployed simultaneously with the caching feature to reduce the real-time inline compression requirements. Using only compression can potentially bottleneck the Virtual Service throughput depending on the hardware platform.

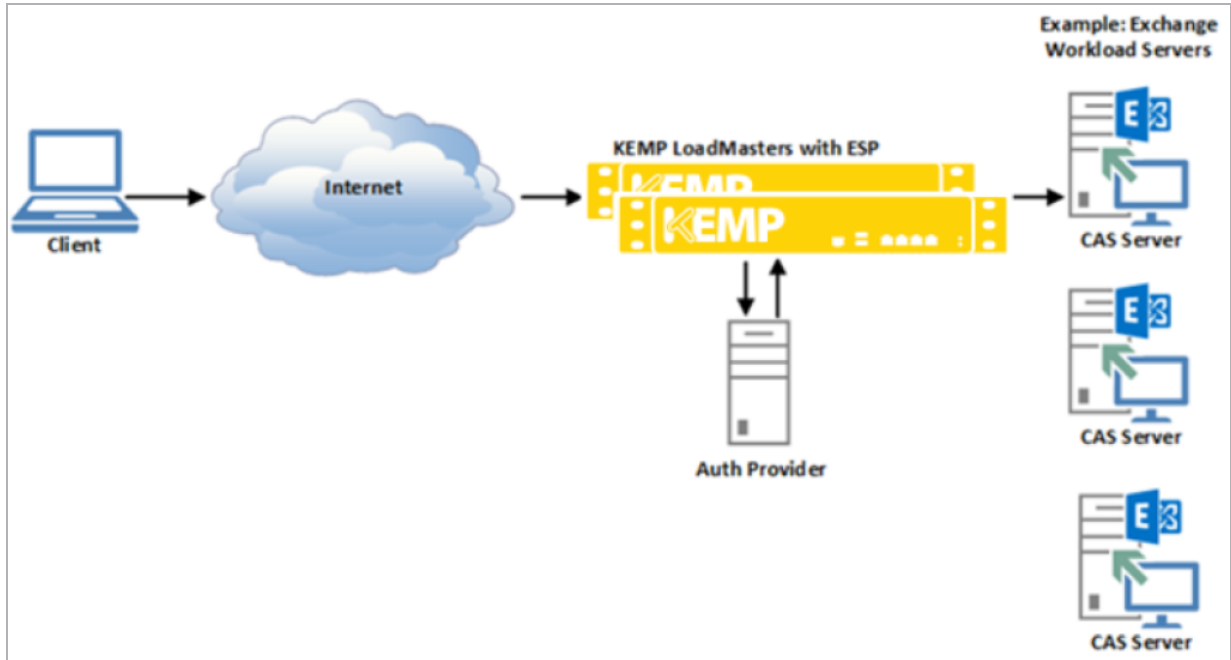
Compression can be enabled per HTTP and per offloaded HTTPS Virtual Services.

Compression depends on clients having gzip support. You can verify that a compressed connection to the Virtual Service exists by tracing the client HTTP traffic. If you can locate the **Content-encoding: gzip** header from the LoadMaster your client communication to the LoadMaster is compressed.

3.6 Authentication

3.6.1 Overview

The Kemp Edge Security Pack (ESP) delivers a solution using the Kemp LoadMaster line of load balancers to customers who would have previously deployed Microsoft's Threat Management Gateway (TMG) to publish their Microsoft applications.



3.6.2 ESP Features

Feature	Description
End Point Authentication for Pre-Auth	Clients who are trying to access Virtual Services on the LoadMaster must provide authentication information which is used by the ESP to validate the clients' right to access the service.
Persistent Logging and Reporting for User Logging	When clients try to access a service, this is logged on the LoadMaster as part of the ESP. This allows monitoring by the administrator.
Single Sign On Across Virtual Services	The LoadMaster is designed to handle multiple virtual services supporting unique workloads. These Virtual Services can be joined together by associating them with the same Single Sign On (SSO) Domain.
Authentication from the LoadMaster to Active Directory	Active Directory is the standard Authentication Provider for Microsoft workloads. LoadMaster supports the key connection types between the LoadMaster and Active Directory.
Basic and Forms-based Authentication	LoadMaster with ESP currently supports basic and form-based authentication including customizable forms.
RADIUS Authentication	A Remote Access Dial In User Service (RADIUS) server can be used to authenticate users who log in to the Kemp LoadMaster.

RSA SecurID Two-Factor Authentication	This scheme authenticates the user on an RSA SecurID Server. When RSA is enabled as the authentication method, during the login process the user is prompted to enter a password that is a combination of two numbers – a Personal Identification Number (PIN) and a token code, which is the number displayed on the RSA SecurID authenticator (dongle).
Kerberos Constrained Delegation (KCD) Authentication	When using KCD as the authentication protocol, the LoadMaster provides seamless access to protected resources in a Kerberos realm even when the credentials provided are not directly valid for such an environment.
Client Certificate Authentication	Using certificates for authentication can be considered more secure because a user cannot gain access to something simply by knowing the username and password. Using certificates prevents key loggers or other malware on a client machine from capturing keystrokes to identify user accounts and passwords.
Dual-factor Authentication	Some authentication mechanisms assume a dual-factor approach where both the Active Directory and a secondary mechanism are used in sequence. For these, the form includes the username, password, and also a passcode which is checked after the username and password.

3.7 Web Application Firewall

Web Application Firewall (WAF) services are natively integrated in the Kemp LoadMaster. This enables secure deployment of web applications, preventing Layer 7 attacks while maintaining core load balancing services, high availability and scalability. The WAF directly augments the LoadMaster's existing security features to create a layered defence for web apps, and therefore enables a safe, compliant, and productive use of web application resources.

3.7.1 WAF Overview

There are many limitations of traditional firewalls, such as the fact that they:


- Typically operate below Layer 4
- Have limited support for protocol inspection (that is, packet filtering)
- Are only aware of 'ports and sockets'

Modern-day hacker attacks present many security challenges due to the fact that they are:

- Becoming increasingly complex
- Frequently changing their attack vectors

- Focusing on application vulnerabilities
- Harder to detect and block on a firewall

3.7.2 WAF Benefits



As depicted in the diagram above, Kemp WAF offers many benefits, such as:

- It processes web-based traffic, that is, HTTP/HTTPS
- It is able to detect modern application attack vectors
- It is deployed between perimeter defences and the application servers
- It works in conjunction with existing security technologies to provide complete application attack prevention

3.8 GEO

GEO assures seamless failover and failback to the best performing and geographically closest datacenter for optimal use of web-based applications. If there is a service disruption, traffic is automatically directed based on a set policy to minimize impact and the need for manual intervention.

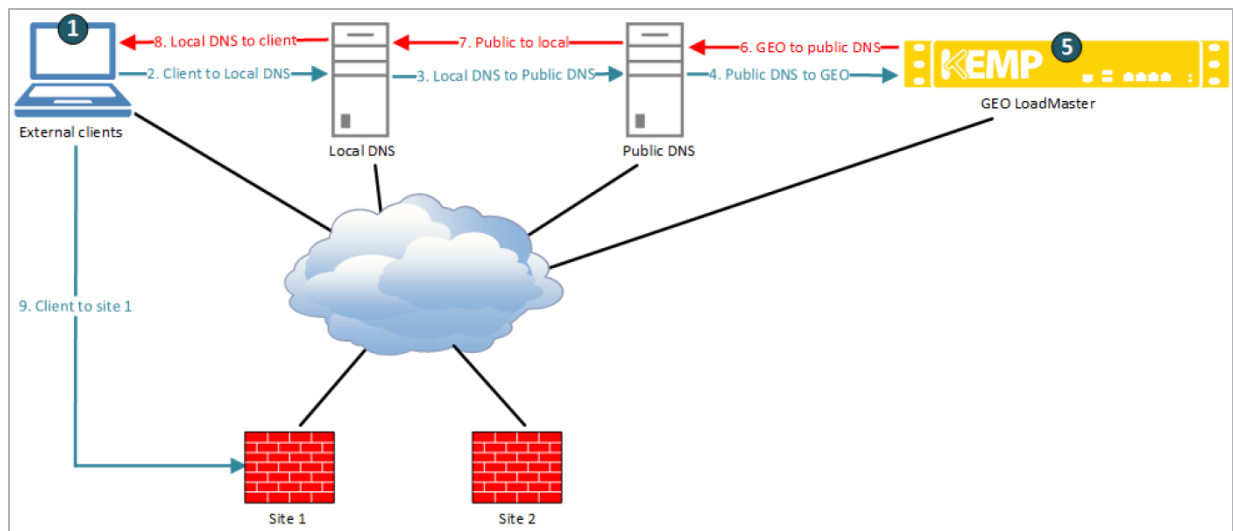
The GEO product is available in two forms:

- A standalone GEO product.
- A Global Server Load Balancing (GSLB) Feature that is part of the Kemp LoadMaster product.

Example GEO Deployment:

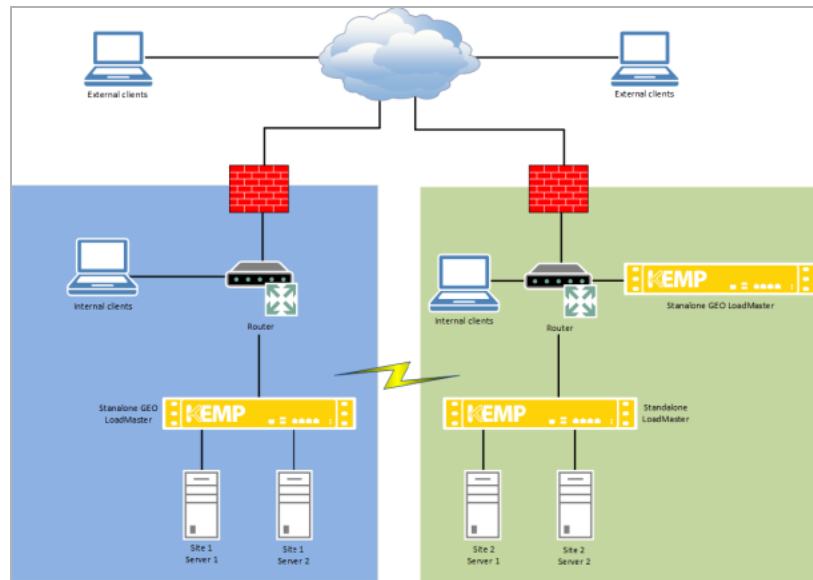
In the following configuration we depict the GEO LoadMaster as being located outside the data centers. Though this can be the case, typically the GEO LoadMaster is located within one or more of the data centers.

3 Load Balancing on LoadMaster



1. An external client tries to connect to FQDN test.domain.com.
2. The external client checks its local DNS.
3. The local DNS forwards the request to the public DNS.
4. The public DNS forwards the request to GEO because it is the authoritative DNS for this record.
5. GEO checks persistence and scheduling and decides which site to return.
6. GEO returns the IP address for the client that made the request (in this case; the public DNS).
7. The public DNS returns the results to the Local DNS.
8. The local DNS returns the result to the client.
9. The client connects directly to the site.

3 Load Balancing on LoadMaster



The above diagram could depict several scenarios:

- Two active/active sites with round robin/location/proximity scheduling
- Two sites – one active and one disaster recovery with fixed weight scheduling
- GEO LoadMasters set up as partners in active/active mode. Refer to the Distributed LoadMaster Partners section of the GEO Feature Description on the [KEMP Documentation page](#) for further information on GEO partners.

By default, if the client is internal, they are given an internal address. If the client is external, they are given an external address.

GEO offers many load balancing algorithms including **round robin**, **weighted round robin**, **fixed weighting**, **real server load**, **location based**, and **proximity**. **Round Robin** load balancing can be used for all active data centers, which includes support for weights and a chained failover option for disaster recovery. **Location Based** load balancing allows GEO to direct a client to a data center based on the client's country or continent, as defined by the created policies. **Proximity** takes **Location Based** one step further and allows for longitude and latitude granularity for definition of proximity. GEO securely and seamlessly integrates with LoadMaster to offer **Real Server Load** load balancing, in which GEO uses local data center metrics provided by LoadMaster, allowing clients to connect to the least busy data center.

GEO can be deployed in a distributed (Active/Active) HA configuration, with multiple appliances securely synchronizing information. Introducing GEO in your existing Authoritative Domain Name Services (DNS) requires minimal integration work and risk, allowing you to fully leverage your existing DNS investment.

For further information on GEO, please refer to the **GEO, Feature Description** on the [Kemp Documentation Page](#).

3.9 FIPS 140-2 Compliance

The National Institute of Standards and Technology (NIST) issued the Federal Information Processing Standard (FIPS) Publication 140-2 to coordinate the requirements and standards for cryptography modules that include both hardware and software components. Federal agencies and departments can validate that the module in use is covered by an existing FIPS 140-2 certificate that specifies the exact module name, hardware, software, firmware, and/or applet version numbers. The cryptographic modules are produced by the private sector or open source communities for use by the U.S. government and other regulated industries (such as financial and health-care institutions) that collect, store, transfer, share and disseminate sensitive but unclassified (SBU) information. A commercial cryptographic module is also commonly referred to as a Hardware Security Module (HSM).

FIPS 140-2 defines four levels of security, 'Level 1' to 'Level 4'. This document describes the compliance of the Kemp LoadMaster products with Level 1 and Level 2 of the FIPS 140-2 standard.

3.9.1 FIPS 140-2 Level 2 Compliance

The FIPS LoadMaster is compliant with FIPS 140-2 Level 2.

Security Level 2 improves upon the levels of security provided for in Level 1. The cryptographic module must protect against unauthorized physical access and include the use of tamper-evident coatings or seals which must be broken to attain physical access to the module. Role-based authentication is also required.

The cryptographic module, or Hardware Security Module (HSM), included in the FIPS products are as follows:

- 1x Cavium CN1620-NFBE3-3.0-FW2.2-G (LM-8020-FIPS product)
- 1x Cavium FIPS CN1610-NFBE1-3.0-FW2.1-G (LM-5305-FIPS product)

These are produced by Cavium Networks. These HSMs are fully compliant with FIPS 140-2 Level 2 requirements as certified by the National Institute of Standards and Technology (NIST), certificate number 1369. The HSM has a secure, tamper-proof barrier and all critical security parameters (CSPs), for example certificate keys, are encompassed within the cryptographic boundary. All keys are cleared if the module is breached.

3.9.1.1 FIPS Processor Overheating - Detection and Prevention Measures

The FIPS processor can be damaged if it is allowed to overheat. To prevent this, the temperature of the card is monitored and the following remedial action is taken if the temperature rises too high.

- If the temperature is above or equal to 60C, but below 70C, a warning is sent to the syslog.
- If the temperature is above or equal to 70C in two consecutive checks, a critical error is sent to the syslog and then the LoadMaster shuts down.
- If the temperature is above 70C but the next check is below 70C, a warning is sent to the syslog but the LoadMaster is not shut down.

When the temperature goes below 60C, the syslog warning messages cease.

3.10 Sub-Virtual Services (SubVSs)

From within a Virtual Service you can create one or more 'Sub-Virtual Services' (SubVSs). A SubVS is linked to, and uses the IP address of, the 'parent' Virtual Service. The SubVSs may have different settings (such as health check methods, content rules, and so on) to the parent Virtual Service and to each other. This allows the grouping of related Virtual Services, all using the same IP address. This could be useful when there are complex configurations such as Exchange or Lync which typically are comprised of a number of Virtual Services.

There are many advantages to using SubVSs, such as:

- SubVSs are linked to, and use the IP address of, the 'parent' Virtual Service
- Using SubVSs reduces the number of IP addresses required by applications such as Lync or Exchange
- SubVSs may have different settings (such as content rules) to the parent Virtual Service and to each other
- Using a SubVS provides the ability to have content switching and persistency on the same Virtual Service - it also allows content switching to a SubVS and then persistence to a Real Server
- Using a SubVS gives the ability to perform multiple health checks on the same Virtual Service
- SubVSs work well with ESP, but ESP is not required

3.11 Certificates

An SSL certificate is required for all SSL transactions, and as such is required for all SSL-enabled Virtual Services. With the LoadMaster, there are two types of SSL certificates:

- Self-signed certificates generated by the LoadMaster itself
- Certificates that are signed by a CA (Certificate Authority) such as Verisign or Thawte
- When an SSL-enabled Virtual Service is configured on the LoadMaster, a self-signed certificate is installed automatically.
- Generally, self-signed certificates should not be used for public-facing production websites.
- They may be acceptable for use in some other scenarios, such as:
 - Intranet sites
 - QA sites, where web sites are tested but not presented to the general public

3.12 User Management

The LoadMaster supports multiple user logins with varying levels of access that can be managed by navigating to **System Configuration > System Administration > User Management**. Each username must be a minimum of three characters and a maximum of 14. Passwords must be a minimum of six characters long. Users created here can only access the WUI, remote access using SSH is not supported.

The LoadMaster can be configured to utilize RADIUS server or client certificates in the authentication process.

4 Miscellaneous

4.1 SNMP Support

Simple Network Management Protocol (SNMP) is a protocol that enables the management of many network devices over the network from a remote management station (SNMP manager).

The manager station can request data from the managed stations (SNMP agents) or it can change the value of data on the agents.

The managed stations (SNMP agents) can also be set up to alert the manager when some predefined events occur, such as a unit failover. The alerting mechanism uses event traps.

The current version is SNMPv3. The two previous revisions in use are SNMPv1 and SNMPv2c (community-based SNMPv2).

The SNMP support of the LoadMaster is based on SNMPv3 and is backwards-compatible such that all three of the above versions can be used. However, since SNMPv1 does not support 64bit-values (as used in the LoadMaster Management Information Base (MIB)), it is recommended to use SNMPv2c or SNMPv3. MsgSecurity is supported only with SNMP v1 and v2c.

When monitoring the LoadMaster in HA mode please monitor individual appliances by the appropriate ethernet address.

The information regarding all LoadMaster-specific data objects is stored in the enterprise-specific MIBs listed below.

ONE4NET-MIB.txt	enterprise id
IPVS-MIB.txt	Virtual Server stats
B-100-MIB.txt	LoadMaster configuration data
CERTS-MIB.txt	SSL certificate information

These MIBs, available for download from www.kemptechnologies.com/documentation, must be installed on the SNMP manager machine to be able to request the performance/config-data of the LoadMaster using SNMP.

SNMP is supported on both IPv4 and IPv6 Virtual Services at both Layer 4 and Layer 7.

The SNMP support is disabled by default.

You can retrieve SSL certificate information including the file name, certificate Subject Name, certificate serial number, certificate start date, certificate end date, and certificate issuer information using SNMP. SNMP can display this information for up to 256 SSL certificates.

You can also retrieve disk space usage details using SNMP. The /var/log and /var/log/userlog partition information is available.

4.2 WUI Authentication and Authorization

The **WUI Authentication and Authorization** screen enables the administration of the available authentication (login) and authorization (allowed permissions) options.

Authentication

Users must be authenticated before logging on to the LoadMaster. The LoadMaster allows authentication of users to be performed using the LDAP authentication method as well as Local User authentication.

When all authentication methods are selected, the LoadMaster attempts to authenticate users using the authentication methods in the following order:

1. LDAP
2. Local Users

If the LDAP server is also not available, then Local User authentication methods are used.

If the LDAP authentication method is not selected, then the Local User authentication method is selected by default.

Authorization

LoadMaster allows the users to be authorized using Local User authorization. The user's authorization decides what level of permissions the user has and what functions on the LoadMaster they are allowed to perform.

Authorization using LDAP is not supported.

4.3 Bonding and VLAN

LoadMaster bonding/VLAN tagging can be easily set up and configured using the WUI. Successful deployment requires that the pre-requisites have been satisfied. This guide is designed to introduce interface bonding and VLAN configuration on the LoadMaster. Bonding support is available with all network modules.

Prerequisite

A list of prerequisites are below:

- VLAN Tagging
- IEEE 802.1Q
- Bonding/Teaming (802.3ad/Active-Backup)
- IEEE 802.1AX/IEEE 802.3ad/LACP

4.4 IPsec Tunneling

Internet Protocol Security (IPsec) is designed and used to provide secure connections between nodes and networks throughout the internet. IPsec has become the standard for most of the IP Virtual Private Network (VPN) technology.

4.5 Debugging Utilities

Debugging utilities can be executed in the WUI at **System Configuration > Logging Options > System Log Files > Debug Options**. These are best utilized when working with the Kemp Support Team.

Last Updated Date

This document was last updated on 27 July 2023.