
LineRate

- [F5 LineRate Manager Guide](#)
- [About This Guide](#)
- [LineRate Overview](#)
- [Technical Support](#)
- [Accessing LineRate Manager](#)
- [Configuring Licensing](#)
- [Manually Installing a Purchased License](#)
- [Working with Bases](#)
- [Configuring Management Interfaces](#)
- [Configuring Proxies](#)
- [Configuring a Forward Proxy](#)
- [Configuring Data Interfaces](#)
- [Configuring the Virtual IP](#)
- [Configuring the Forward Proxy](#)
- [Configuring a Reverse Proxy](#)
- [Configuring Data Interfaces](#)
- [Configuring Load Balancing](#)
- [Monitoring and Troubleshooting Load Balancing](#)
- [Configuring SSL](#)
- [Configuring A/B Testing](#)
- [LineRate Manager Reference](#)
- [Configuration Tables](#)
- [Dashboards](#)
- [DNS](#)
- [Forward Proxies](#)
- [Health Monitors](#)
- [Interfaces](#)
- [IP Routes](#)
- [Licenses](#)
- [Line and Area Charts](#)
- [NTP](#)
- [Phone Home](#)
- [Real Servers](#)
- [Real Server Groups](#)
- [REST Server](#)
- [Scripts](#)

- [SSH](#)
- [SSL](#)
- [SSL Profiles](#)
- [System](#)
- [Users](#)
- [Virtual IPs](#)
- [Virtual Servers](#)
- [Whiteboards](#)



F5 LineRate Manager Guide

Overview

This guide is your starting point to learn about F5[®] LineRate Manager, the F5[®] LineRate[®] graphical user interface. We'll walk you through the basics that you need to know to use F5[®] LineRate Manager. The guide then continues with configuring basic examples, including management access, configuring a load balancer (reverse proxy), including SSL setup, and configuring a forward proxy.

About the Examples

We provide a detailed architecture example, including all naming, IP addresses, and other settings, so you can focus on understanding how to use the software, not on what to name things.

At the end of the guide, we have a complete annotated example of everything you configured for you to refer to.

After completing this example configuration, you will be better prepared to plan for your F5[®] LineRate[®] implementation.

Contents

The guide is broken into the following sections:

- [About This Guide](#)
- [LineRate Overview](#)
- [Accessing LineRate Manager](#)
- [Configuring Licensing](#)
- [Working with Bases](#)
- [Configuring Management Interfaces](#)
- [Configuring Proxies](#)
- [LineRate Manager Reference](#)

About This Guide

1. [Overview](#)
2. [Audience](#)
3. [Conventions](#)
4. [Example IP Addresses](#)
5. [Searching the Guide](#)
 1. [Relevance Level](#)
 2. [Limiting a Search to Specific Tree](#)
 3. [Term Modifiers](#)
 1. [Wildcard Searches](#)
 2. [Fuzzy Searches](#)
 3. [Proximity Searches](#)
 4. [Boosting a Term](#)
 4. [Boolean Operators](#)
 1. [OR](#)
 2. [AND](#)
 3. [+](#)
 4. [NOT](#)
 5. [Grouping](#)
 6. [Escaping Special Characters](#)
6. [Legal Notices](#)
 1. [Copyright](#)
 2. [Trademarks](#)

Overview

This About page contains general information about this guide, including the audience, typographic conventions, and how to search the content.

Audience

This guide is intended for experienced network administrators and network architects who understand your organization's existing TCP/IP network and who need to configure or manage proxy services, such as traffic steering and SSL offload, using F5[®] LineRate Manager, the graphical user interface for F5[®] LineRate[®].

Conventions

This guide uses the following symbols and typographic conventions.

Convention	Definition
Monospaced bold	Text in a monospaced bold font represents commands or other text that you type exactly as you see it.
<angle bracket>	Text in a monospaced bold font inside angle brackets represents a placeholder that describes what you must type.
[square brackets]	Text in a monospaced bold font inside square brackets represents an optional command or option.
Monospaced	Text in a monospaced font represents output or results the system displays.
Bold	Text in bold shows keys to press and items to select or click, such as menu items or buttons.
	Shows the beginning of a procedure.
 Caution	Cautions contain critical information about configuring your system or data.
 Note	Notes contain important information that may affect how you install or configure your system.
 Tip	Tips contain best practices or useful information to help you when configuring your system.
	Shows that the content is for advanced users.

Example IP Addresses

Throughout this guide, we use example IP addresses for both internal (private) and external (public) uses.

For private addresses, we use the IP addresses designated in [RFC 1918](#):

- 10.0.0.0 - 10.255.255.255 (10/8 prefix)
- 172.16.0.0 - 172.31.255.255 (172.16/12 prefix)
- 192.168.0.0 - 192.168.255.255 (192.168/16 prefix)

For public addresses, we use the IP addresses designated for documentation in [RFC 5737](#):

- 192.0.2.0/24 (TEST-NET-1)
- 198.51.100.0/24 (TEST-NET-2)
- 203.0.113.0/24 (TEST-NET-3)

Searching the Guide

The search box at the top-right of each page lets you enter a term or phrase to search for. By default, the system searches all pages in the F5[®] LineRate[®] content. Searches are not case sensitive. By default, searches find plurals and other matches from word stems, such as tests, testing, tested, and tester if you search for test.

You can search for a single term such as:

```
interface
```

Or

```
certificate
```

You can also search for an exact phrase surrounded by double quotes such as:

```
"real server"
```

Or

```
"IP address"
```

Relevance Level

By default, the system sorts the search results by relevance. The relevance is determined by a weighting algorithm that takes into consideration the page title, content, tags, and attachments. The relevance is also affected by the page rating (thumb up or down) and by how often other users select a page to view from similar searches.

Searches can return a large number of results. You can narrow your searches a number of ways by:

- Limiting your search to a specific tree
- Using term modifiers
- Using Boolean operators

Limiting a Search to Specific Tree

If you only want to search one area or tree of a guide, you can limit your search to that tree. For example, if you only want to search the Configure Command tree of the 2.6 Release of the CLI Reference Guide for the term "interface," you can enter your search like this:

```
+(path:087Release_2.6/200CLI_Reference_Guide/Configure_Commands/*) AND interface
```

You can further narrow the search using the term modifiers and Boolean operators (described below):

```
+(path:087Release_2.6/200CLI_Reference_Guide/Configure_Commands/*) AND interface  
AND CARP
```

```
+(path:087Release_2.6/*) AND load AND balancer
```



Note: For a tree-specific search, words in quotes are not treated as a specific phrase. The search does an OR search for any words in quotes, so you may not want to use quotes and use AND instead, as shown in the example above.

A few steps to help with this type of search:

1. Navigate to the tree you want to search.
2. In your browser's address bar, copy the address of the page.
 - You only need the part after the "<https://docs.lineratesystems.com/>".
3. Using the syntax example above, type in your search and paste in the path of the page you want to search.

Term Modifiers

The search supports modifying query terms to provide a wide range of searching options.

Wildcard Searches

The guides support single- and multiple-character wildcard searches with single terms (not within phrase queries).

To perform a single-character wildcard search, use the ? symbol.

To perform a multiple-character wildcard search, use the * symbol.

The single-character wildcard search looks for terms that match that with the single character replaced. For example, to search for "text" or "test" you can use the search:

```
te?t
```

The multiple-character wildcard search looks for 0 or more characters. For example, to search for test, tests or tester, you can use the search:

```
test*
```

You can also use the wildcard searches in the middle of a term.

```
te*t
```



Note: You cannot use a * or ? symbol as the first character of a search.

Fuzzy Searches

The guide supports fuzzy searches based on the Levenshtein Distance or Edit Distance algorithm. To do a fuzzy, search use the tilde ~ symbol at the end of a single word. Fuzzy searches work for multiple characters. For example, to search for a term similar in spelling to "roam" use the fuzzy search:

```
roam~
```

This search will find terms like foam and roams.

You can add an optional parameter to specify the required similarity. The value is between 0 and 1. With a value closer to 1, only terms with a higher similarity will be matched. For example:

```
roam~0.6
```

The default is 0.5.

Proximity Searches

The guide supports finding words that are within a specific distance from each other. To do a proximity search, use the tilde ~ symbol at the end of a phrase. For example, to search for a "feature" and "standard" within 10 words of each other in a document use the search:

```
"feature standard"~10
```

Boosting a Term

The guide provides the relevance level of matching documents based on the terms found. To boost a term, use the caret ^ symbol with a boost factor (a number) at the end of the term you are searching. The higher the boost factor, the more relevant the term will be.

Boosting allows you to control the relevance of a document by boosting its term. For example, if you are searching for:

```
mindtouch search
```

and you want the term "mindtouch" to be more relevant boost it using the ^ symbol along with the boost factor next to the term. You would type:

```
mindtouch^4 search
```

This will make documents with the term mindtouch appear more relevant. You can also boost phrases as in the example:

```
"mindtouch search"^4 "Apache"
```

By default, the boost factor is 1. Although the boost factor must be positive, it can be less than 1 (e.g. 0.2)

Boolean Operators

Boolean operators allow terms to be combined through logic operators. MindTouch supports AND, +, OR, NOT, and - as Boolean operators.



Note: Boolean operators must be ALL CAPS.

OR

The OR operator is the default conjunction operator. This means that if there is no Boolean operator between two terms, the OR operator is used. The OR operator links two terms and finds a matching document if either of the terms exist in a document. This is equivalent to a union using sets. The symbol || can be used in place of the word OR.

To search for documents that contain either "mindtouch search" or just "mindtouch" use the query:

```
"mindtouch search" mindtouch
```

or

```
"mindtouch search" OR mindtouch
```

AND

The AND operator matches documents where both terms exist anywhere in the text of a single document. This is equivalent to an intersection using sets. You can use the symbol && in place of the word AND.

To search for documents that contain "mindtouch search" and "Advanced" use the query:

```
"mindtouch search" AND "Advanced"
```

+

The + (required operator) requires that the term after the + symbol exist somewhere in a document.

To search for documents that must contain "search" and may contain "advanced," use the query:

```
+search advanced
```

NOT

The NOT operator excludes documents that contain the term after NOT. This is equivalent to a difference using sets. You can use the symbol ! in place of the word NOT.

To search for documents that contain "mindtouch search" but not "Advanced" use the query:

```
"mindtouch search" NOT "Advanced"
```



Note: The NOT operator cannot be used with just one term. For example, the following search will return no results:

```
NOT "mindtouch search"
```

Grouping

The guide supports using parentheses to group clauses to form sub queries. This can be very useful if you want to control the Boolean logic for a query.

To search for either "mindtouch" or "search" and "advanced" use the query:

```
(mindtouch OR search) AND advanced
```

This eliminates any confusion and makes sure you that website must exist and either term mindtouch or search may exist.

Escaping Special Characters

The Guide supports escaping special characters that are part of the query syntax. The current list of special characters is:

```
+ - && || ! ( ) { } [ ] ^ " ~ * ? : \
```

To escape these character use the \ before the character. For example, to search for (1+1):2 use the query:

```
\(1\+1\)\:2
```

Legal Notices

Copyright

Copyright © 2015, F5 Networks, Inc. All rights reserved.

F5 Networks, Inc. (F5) believes the information it furnishes to be accurate and reliable. However, F5 assumes no responsibility for the use of this information, nor any infringement of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent, copyright, or other intellectual property right of F5 except as specifically described by applicable user licenses. F5 reserves the right to change specifications at any time without notice.

Trademarks

AAM, Access Policy Manager, Advanced Client Authentication, Advanced Firewall Manager, Advanced Routing, AFM, APM, Application Acceleration Manager, Application Security Manager, ARX, AskF5, ASM, BIG-IP, BIG-IQ, Cloud Extender, CloudFucious, Cloud Manager, Clustered Multiprocessing, CMP, COHESION, Data Manager, DevCentral, DevCentral [DESIGN], DNS Express, DSC, DSI, Edge Client, Edge Gateway, Edge Portal, ELEVATE, EM, Enterprise Manager, ENGAGE, F5, F5 [DESIGN], F5 Certified [DESIGN], F5 Networks, F5 SalesXchange [DESIGN], F5 Synthesis, f5 Synthesis, F5 Synthesis [DESIGN], F5 TechXchange [DESIGN], Fast Application Proxy, Fast Cache, FirePass, Global Traffic Manager, GTM, GUARDIAN, iApps, IBR, Intelligent Browser Referencing, Intelligent Compression, IPv6 Gateway, iControl, iHealth, iQuery, iRules, iRules OnDemand, iSession, L7 Rate Shaping, LC, Link Controller, Local Traffic Manager, LineRate, LineRate Precision, LineRate Systems [DESIGN], LROS, LTM, Message Security Manager, MSM, OneConnect, Packet Velocity, PEM, Policy Enforcement Manager, Point Load Balancer, Protocol Security Manager, PSM, Real Traffic Policy Builder, SalesXchange, ScaleN, Signalling Delivery Controller, SDC, SSL Acceleration, software designed applications services, SDAC (except in Japan), StrongBox, SuperVIP, SYN Check, TCP Express, TDR, TechXchange, TMOS, TotALL, Traffic Management Operating System, Traffix Systems, Traffix Systems (DESIGN), Transparent Data Reduction, UNITY, VAULT, vCMP, VE F5 [DESIGN], Versafe, Versafe [DESIGN], VIPRION, Virtual Clustered Multiprocessing, WebSafe, and ZoneRunner, are trademarks or service marks of F5 Networks, Inc., in the U.S. and other countries, and may not be used without F5's express written consent.

All other product and company names herein may be trademarks of their respective owners.



LineRate Overview

1. [Overview](#)
2. [Reverse Proxy](#)
3. [Forward Proxy](#)
4. [Security](#)
 1. [IP-based Protocols](#)
 1. [UDP](#)
 2. [TCP](#)
 3. [CARP](#)
 4. [ICMP](#)
 5. [Other](#)
 2. [Non-IP-based Protocols](#)
 1. [ARP, LACP, and VLAN](#)
 2. [Other](#)
5. [What's Next](#)

Overview

This section includes introductory topics to help you understand how to use F5[®] LineRate Manager, the F5[®] LineRate[®] graphical user interface.

F5[®] LineRate[®] is the industry's first purely software-defined network appliance delivering advanced performance, availability, and scale to web and secure web services. You install F5[®] LineRate[®] natively on a server platform or as a virtual appliance on a hypervisor to create a highly scalable network appliance.

Reverse Proxy

The F5[®] LineRate[®] reverse proxy capability provides a full-proxy front-end to the web and secure web services on back-end servers of the application. To successfully manage load, health, and availability of the application, F5[®] LineRate[®] constantly monitors the traffic flows from the client to the back-end servers to ensure the delivery of the application services. F5[®] LineRate[®] provides several discrete services including health monitoring, transaction statistics, and traffic management to ensure balanced usage of back-end servers and rapid client service.

Forward Proxy

The F5[®] LineRate[®] forward proxy capability provides a proxy function from one network to another. A common use case for a forward proxy is for connections from your private network to the Internet.

A forward proxy lets you insert custom logic created with scripts. Scripts can perform a variety of functions, including gathering usage statistics, redirecting requests to your own cache, blocking of access to specific sites, managing cookies, and much more.

To use a forward proxy effectively, be sure to create and attach a virtual IP that includes the range of Internet IP addresses you want to go through the forward proxy.



Caution: When attaching a virtual IP to a forward proxy, the virtual IP must not include any of the system's own IP addresses. For a virtual IP with a single IP address, do not set the virtual IP's IP address to one of the system's own IP addresses. For a virtual IP with a range of addresses, you must ensure that the IP address range does not contain any of the system's own IP addresses. This may mean you need to break the virtual IP into multiple virtual IPs. See [Configuring a range for a virtual IP with forward proxy](#) for more detail and an example.

Security

F5[®] LineRate[®] uses several security measures to mitigate Denial of Service (DoS) attacks. These measures greatly reduce the system's vulnerability to DoS attacks. In general, the system ignores ("black-holes") requests to IP addresses and ports that are not configured on an interface or virtual IP.

The sections below provide additional details.

IP-based Protocols

The sections below describe how the system handles IP-based protocols.

UDP

The system discards most UDP requests, with the exception of DNS request responses.

The system accepts the following traffic:

- MOUNT (dynamic port number)
- NFS (port 2049)
- RPCBIND (port 111)
- STAT (dynamic port number)

Aside from the RPCBIND service, these Sun RPC services have dynamically assigned port numbers which are then discovered by the client via RPCBIND. Client traffic returning on ephemeral ports will also be accepted and processed.

All other traffic on ports without bound sockets in the LISTEN state is discarded without processing or ICMP generation.

These services are all NFS-related. These ports will be open and active only when using the PXE boot feature in F5[®] LineRate[®].

TCP

The system accepts TCP traffic only on IPs and ports configured on virtual IPs, for the REST server, for the SSH server, and for active TCP connections initiated by the F5[®] LineRate[®], such as those to real servers. TCP traffic directed to other IPs and ports on the system not in those categories will be dropped without acknowledgement or ICMP generation.

CARP

The system accepts and processes all inbound CARP traffic.

ICMP

The system accepts ICMPv4 messages with the following [Type, Code] pairs:

- [0x00, 0x00] - Echo Reply
- [0x08, 0x00] - Echo Request*
- [0x03, 0x03] - Destination Port Unreachable
- [0x03, 0x04] - Destination Requires Fragmentation
- [0x0B, 0x00] - Time Exceeded

The system accepts ICMPv6 messages with the following [Type, Code] pairs:

- [0x81, 0x00] - Echo Reply
- [0x80, 0x00] - Echo Request*
- [0x01, 0x04] - Destination Port Unreachable
- [0x02, 0x00] - Packet Too Big
- [0x03, 0x00] - Time Exceeded
- [0x87, 0x00] - Neighbor Solicitation
- [0x88, 0x00] - Neighbor Advertisement

* - When not in gateway (IP forwarding) mode, the following criteria must be met:

- Destination IP address for the incoming request must be an address assigned to the ingress interface, or a CARP IP address on the ingress interface or an IP address on a loopback interface.
- The ICMP response must route out the ingress interface (reverse path filtering).
- The source IP address must be within the same IP subnet as the destination. This means that ICMP "pings" are only responded to when initiated from a host local to one of the F5[®] LineRate[®] system's interfaces.

All other ICMP traffic is ignored without processing.

Other

Any inbound IP traffic not explicitly referenced above is dropped without any processing.

Non-IP-based Protocols

The sections below describe how the system handles non-IP-based protocols.

ARP, LACP, and VLAN

The system accepts and processes all ARP, LACP, and VLAN traffic normally and responds appropriately.

Other

The system drops, without processing, any non-IP traffic not explicitly mentioned above.

What's Next

The rest of this guide describes how to install F5[®] LineRate[®] and takes you through an example configuration using F5[®] LineRate Manager. Continue with [Accessing LineRate Manager](#).



Technical Support

Support tools are available to help you answer your questions whenever and wherever you need help. From the documentation to the global technical community you can collaborate with on DevCentral, F5[®] LineRate[®] self-service tools help you solve issues quickly and proactively.

The [LineRate Support page](#) can help you find the resources you need.



Accessing LineRate Manager

1. [Overview](#)
2. [Operating Systems and Browsers](#)
3. [Logging in to F5® LineRate Manager](#)
4. [Understanding the F5® LineRate Manager Page](#)
5. [Logging Out](#)
6. [What's Next](#)

Overview

This section describes how to access F5® LineRate Manager and what you see in the default dashboard that displays.

For the system requirements and installation information, see the *Getting Started Guide* ([System Requirements](#) and [Installing LineRate](#)).

Operating Systems and Browsers

F5® LineRate Manager can be accessed using an HTML-compatible browser. The following operating systems and browsers have been tested with F5® LineRate Manager.

Windows 7 and Windows 8	Chrome Firefox Internet Explorer 9 and later
Macintosh OS X	Chrome Firefox Safari
Linux	Chrome Firefox

Logging in to F5[®] LineRate Manager

You must log in to configure and manage your system.

Use the default settings the first time:

- Username—admin
- Password—changeme

We highly recommend that you change the password for the admin login, after you have logged in.



Note: For Amazon EC2, the default username "admin" has no password. Instead, the admin user has an SSH key configured from the EC2 key pair specified when you launched the EC2 instance. To access F5[®] LineRate Manager, you must set a password either via the EC2 [User Data](#) or the F5[®] LineRate[®] [CLI](#) over SSH.

If you want to change the password, you must do so using the CLI. See [Logging in to the CLI](#).




To log in to F5[®] LineRate Manager:

1. Configure the management IP address.
 - See [Configuring Management Interfaces](#).
2. From your browser, go to: `https://<management_ip_address>:8443`.
 - The login window displays
3. Enter your User Name and Password.
4. Click **Log In**.
 - Port 8443 also provides access to the REST server. For information about accessing the REST server, the default settings, and how to change the settings, see [Accessing the REST Server](#).
 - The F5[®] LineRate Manager page displays.
 - After one hour of inactivity, the F5[®] LineRate Manager session times out.

Understanding the F5[®] LineRate Manager Page

When you first access F5[®] LineRate Manager for a new installation, the F5[®] LineRate Manager page displays the default dashboard, which includes a whiteboard and a configuration table. These are blank for any new installation, as shown below. You can also add charts to the dashboard. See [Charts](#).



Item	Name	Description
1	N/A	IP address of F5 [®] LineRate [®] system.
2	N/A	Menus to configure dashboards and charts. See Dashboards .
3	Get system information	Click to access system user guides and system version information.
4	Configure system	Click to save to the startup config, to view interface and IP route configuration, and to configure SSL and purchased licenses.
5	Configuration table	Display most configured objects and create, edit, and delete objects. You cannot configure all objects using F5 [®] LineRate Manager. See Configuration Tables .
6	Whiteboard	Display a diagram of the selected object, for example, virtual servers. See Whiteboards .

Logging Out

You log out of an existing F5[®] LineRate Manager session by clicking **Log Out** at the top-right of the page.

What's Next

After you access F5[®] LineRate Manager, you are ready to configure licensing. See [Configuring Licensing](#).

Configuring Licensing

1. [Overview](#)
2. [Activating a Starter Edition License](#)
3. [Activating a Purchased License](#)
 1. [Downloading Point™ Load Balancer](#)
4. [Upgrading or Renewing a License](#)
5. [Confirming the License](#)
6. [What's Next](#)

Overview

This section describes how to configure licensing. After you install F5[®] LineRate[®], you can configure the system. To run traffic on the system, you must activate licensing.



Note: If you subscribed to an hourly product (not BYOL) in the [AWS Marketplace](#), you can skip the licensing procedures below. Your instance is automatically licensed. If you subscribed to a BYOL product in the AWS Marketplace, you must follow the steps for the appropriate version in the sections below.

Activating a Starter Edition License

A Starter Edition license defines the rate limits for HTTP requests, TCP connections, and Mb per second. For information about how the limits work, see [About Licensing](#).

To acquire a Starter Edition license, the system must be able to access the license and phone home servers and will use the phone home feature to send usage and configuration data to F5.

For the system to access the Internet and use phone home, the system must have the following:

- IP address on an interface
- Default route to reach the Internet
- Phone home configuration with a F5 username and password

You should not need to configure DNS, because the system defaults will work in most cases. For information about configuring DNS, see [IP Mode Commands](#).

If you have already configured the Starter Edition license, but something is not working properly, see [Troubleshooting Licensing](#).



To enable a starter edition license:

1. Go to linerate.f5.com/try to select the Starter Edition, register, and download F5[®] LineRate[®].
 - If you already have an F5 account, click **LOGIN**.
 - Be sure to keep the email that contains your registration key. You will need the registration key to enable the Starter Edition license.
2. Install or launch F5[®] LineRate[®].
 - See [Installing LineRate](#). For Amazon EC2, see [Launching F5[®] LineRate[®]](#).
3. Log in to F5[®] LineRate[®].
 - See [Accessing LineRate Manager](#).
4. Make sure the F5[®] LineRate[®] system is properly configured to access the Internet.
 - For virtual environments, DHCP is enabled by default, and you should not need to configure anything.
 - For other environments, do one of the following:

To use DHCP	To manually configure a default route and IP address
<ul style="list-style-type: none">• Use the following commands to enable DHCP on one interface using the CLI: <pre>config interface <interface_name> ip address dhcp</pre>	<ol style="list-style-type: none">1. Click Configure system.2. Click Interfaces.3. Assign a valid IP address to an interface.4. Click IP Routes.5. Click Add static route  and add a default route.<ul style="list-style-type: none">• Enter one of the following in the Destination field:<ul style="list-style-type: none">• Default-IPV4• Default-IPV6• Be sure that the Destination and Gateway address types match, that is, they are both either IPv4 or IPv6.

5. Check network connectivity to ensure that phone home works properly.
 - Using bash mode, you can use ping or telnet to make sure you can both resolve and reach api.f5.com, asb.f5.com, and activate.f5.com before configuring phone home. See [Bash Mode Commands](#) for instructions on using system tools available in the bash shell.
6. Click **Configure system**.
7. Click **Phone Home**.
8. Enter your F5 user name and password.

9. For Amazon EC2, select the "I agree to phone home terms" option.
10. Click **Save**.
11. Click **System** and click **Save to Startup**.
 - Once configured, the system accesses the phone home server using your F5 credentials and automatically installs the license. If you need help with your F5 credentials, go to <https://devcentral.f5.com/login> to retrieve your username or password. For more information about what phone home does, see [Phone Home Mode Commands](#).
 - We recommend that you check network connectivity after configuring your default route and IP address to ensure that phone home works properly. Using bash mode, you can use ping or telnet to make sure you can both resolve and reach api.f5.com and asb.f5.com before configuring phone home. See [Bash Mode Commands](#) for instructions on using system tools available in the bash shell.
12. From F5[®] LineRate Manager, click **Licenses**.
13. Click **Activate License**.
14. In the Registration Key field, paste the registration key.
 - You received the registration key in an email after downloading the F5[®] LineRate[®] installation file.
15. Do one of the following:

To automatically install the license	To manually install the license
<ul style="list-style-type: none"> • Click Automatic Activation. <ul style="list-style-type: none"> • The system accesses the licensing server and automatically installs the license. 	<ol style="list-style-type: none"> 1. Click Manual Activation. 2. In the Create License area, copy the dossier text. 3. Go to https://activate.f5.com. 4. Paste in your dossier and follow the directions on the page. 5. Either download the license file or copy the license that displays. 6. In F5[®] LineRate Manager, click Next. 7. Install the license using one of the following methods: <ul style="list-style-type: none"> • If you downloaded the license file, click Choose File and select the license file. • If you copied the license file, paste the text in text box. 8. Click Next.

- After you configure licensing, the Unlicensed System message across the top of F5[®] LineRate Manager no longer displays. If the message remains, confirm the license as described in [Confirming the License](#).

Activating a Purchased License

To purchase a subscription license, go to linerate.f5.com/try. For other licensing options, [contact sales](#).

After you install F5[®] LineRate[®], you can install the license.

The following are required for your purchased license to automatically renew:

- The credit card on file must be active and must process the purchase or renewal.
- The F5[®] LineRate[®] hardware must have Internet connectivity.
- A F5[®] LineRate[®] interface must be configured with either DHCP or a default route and IP address.
- F5[®] LineRate[®] must have licensing set to auto.



Note: If your F5[®] LineRate[®] system does not meet one or more of the criteria, you can manually install and renew the license before the license period expires. See [Manually Installing a Purchased License](#).

The rate limits for HTTP requests, TCP connections, and Mb per second are based on the license you purchased. For more information about how licenses work, see [About Licensing](#).

Currently, the only feature you can license is called base.



To enable a purchased license:

1. Go to linerate.f5.com/try to select the Starter Edition, register, and download F5[®] LineRate[®].
 - If you already have an F5 account, click **LOGIN**.
 - Be sure to keep the email that contains your registration key. You will need the registration key to enable the Starter Edition license.
2. Install or launch F5[®] LineRate[®].
 - See [Installing LineRate](#). For Amazon EC2, see [Launching F5[®] LineRate[®]](#).
3. Log in to F5[®] LineRate[®].
 - See [Accessing LineRate Manager](#).
4. Make sure the F5[®] LineRate[®] system is properly configured to access the Internet.
 - For virtual environments, DHCP is enabled by default, and you should not need to configure anything.
 - For other environments, do one of the following:

To use DHCP	To manually configure a default route and IP address
<ul style="list-style-type: none">• Use the following commands to enable DHCP on one interface using the CLI: <code>config</code> <code>interface <interface_name></code> <code>ip address dhcp</code>	<ol style="list-style-type: none">1. Click Configure system.2. Click Interfaces.3. Assign a valid IP address to an interface.4. Click IP Routes.

To use DHCP	To manually configure a default route and IP address
	<p>5. Click Add static route  and add a default route.</p> <ul style="list-style-type: none"> • Enter one of the following in the Destination field: <ul style="list-style-type: none"> • Default-IPV4 • Default-IPV6 • Be sure that the Destination and Gateway address types match, that is, they are both either IPv4 or IPv6.

5. Check network connectivity to ensure that phone home works properly.
 - Using bash mode, you can use ping or telnet to make sure you can both resolve and reach api.f5.com, asb.f5.com, and activate.f5.com before configuring phone home. See [Bash Mode Commands](#) for instructions on using system tools available in the bash shell.
6. Click **Licenses**.
7. Click **Activate License**.
8. In the Registration Key field, paste the registration key.
 - You received the registration key in an email after downloading the F5[®] LineRate[®] installation file.
 - If you need help with your F5 credentials, go to linerate.f5.com/try and click **LOGIN** to retrieve your username or password. For more information about what phone home does, see [Phone Home Mode Commands](#)
9. Click **Automatic Activation**.
 - The system accesses the licensing server and automatically installs the license.
 - After you configure licensing, the Unlicensed System message across the top of F5[®] LineRate Manager no longer displays. If the message remains, confirm the license as described in [Confirming the License](#).
10. (Optional) Configure phone home credentials.
 - If your license expires, which happens only if the renewal payment is not made, the system starts using a Starter Edition license after expiration. However, the system can only activate the Starter Edition license if valid phone home credentials are configured and phone home is working.
 - If you need help with your F5 credentials, go to linerate.f5.com/try and click **LOGIN** to retrieve your username or password. For more information about what phone home does, see [Phone Home Mode Commands](#).
 1. Click **Configure system**.
 2. Click **Phone Home**.
 3. Enter your F5 username and password.
 4. For Amazon EC2, select the "I agree to phone home terms" option.
 5. Click **Save**.

6. Click **System** and click **Save to Startup**.

---->section break<----

Activating an Offering License (for Volume Licensing)

To purchase volume licensing for Point™ Load Balancer, contact [sales](#). After you install a Point™ Load Balancer instance, you can activate the offering license. The offering license is the license you activate on each Point™ Load Balancer instance.



Note: For offering licenses, we recommend that you configure your orchestration system to start up and configure licensing using the REST API.

The following are required to activate the offering license and for it to automatically renew:

- Two or more Volume License Manager (VLM) license managers must be configured, licensed, and running. See [Configuring VLM](#).
- The offering registration key must be available. See [Configuring VLM](#).
- The IP addresses for the three license managers must be available.
- The Point™ Load Balancer hardware must have access to the license managers.
- Point™ Load Balancer must have licensing set to auto.

The rate limits for HTTP requests, TCP connections, and Mb per second are based on the license you purchased. For more information about how licenses work, see [About Licensing](#).

Currently, the only feature you can license is called base.

If your Point™ Load Balancer hardware does not have access to the license managers, you can manually activate an offering license. See [Manually Activating an Offering License](#).

To enable a volume offering license on each F5® LineRate® instance:

- Check network connectivity to ensure that phone home works properly.
 - Using bash mode, you can use ping or telnet to make sure you can both resolve and reach `api.f5.com`, `asb.f5.com`, and `activate.f5.com` before configuring phone home. See [Bash Mode Commands](#) for instructions on using system tools available in the bash shell.
- Click **Licenses**.
- Click **Activate License**.
- In the Registration Key field, paste the registration key.
- Select the Volume License option.
- Click **Automatic Activation**.
- If not already configured, in the IP Address fields, enter the IP addresses for the three license managers.

- If not already configured, in the Description field, enter a description that identifies this instance of Point™ Load Balancer to the license managers.
 - This is optional, but recommended to permit identification of each instance of Point™ Load Balancer
 - Click **Activate License**.
 - The system accesses the licensing server and automatically installs the license.
 - After you configure licensing, the Unlicensed System message across the top of F5® LineRate Manager no longer displays. If the message remains, confirm the license as described in [Confirming the License](#).
 - (Optional) Configure phone home credentials.
 - If your license expires, which happens only if the renewal payment is not made, the system starts using a Starter Edition license after expiration. However, the system can only activate the Starter Edition license if valid phone home credentials are configured and phone home is working.
 - If you need help with your F5 credentials, go to linerate.f5.com/try and click **LOGIN** to retrieve your username or password. For more information about what phone home does, see [Phone Home Mode Commands](#).
 1. Click **Configure system**.
 2. Click **Phone Home**.
 3. Enter your F5 username and password.
 4. For Amazon EC2, select the "I agree to phone home terms" option.
 5. Click **Save**.
 6. Click **System** and click **Save to Startup**.
1. Install Point™ Load Balancer using the file the file provided by F5® sales personnel.
 - If you do not have the file, see [Downloading Point™ Load Balancer](#).
 2. Make sure the license managers are configured.
 - See [Configuring VLM](#).
 3. Log in to Point™ Load Balancer.
 - See [Accessing LineRate Manager](#).
 4. Make sure the Point™ Load Balancer instance is properly configured to access the Internet.
 - For virtual environments, DHCP is enabled by default, and you should not need to configure anything.
 - For other environments, do one of the following:

To use DHCP	To manually configure a default route and IP address
<ul style="list-style-type: none"> • Use the following commands to enable DHCP on one interface using the CLI: <code>config</code> <code>interface <interface_name></code> 	<ol style="list-style-type: none"> 1. Click Configure system. 2. Click Interfaces. 3. Assign a valid IP address to an interface. 4. Click IP Routes.

To use DHCP	To manually configure a default route and IP address
<pre>ip address dhcp</pre>	<p>5. Click Add static route  and add a default route.</p> <ul style="list-style-type: none"> • Enter one of the following in the Destination field: <ul style="list-style-type: none"> • Default-IPV4 • Default-IPV6 • Be sure that the Destination and Gateway address types match, that is, they are both either IPv4 or IPv6.

Downloading Point™ Load Balancer

If you do not have the F5® LineRate Point™ Load Balancer installation file, you can download one. You can also use this procedure to upgrade to a new version of Point™ Load Balancer.



To download Point™ Load Balancer:

1. Go to downloads.f5.com.
2. Register and log in.
3. Click **Find a Download**.
4. Click [F5 Linerate Application Proxy](#) in the product listing.
5. Click the product name.
6. To accept the terms of the EULA, click **I Accept**.
7. Select an appropriate image type from the list to download.

Upgrading or Renewing a License

If you upgrade your license to different tier or manually renew your license from the linerate.f5.com/try site, you must take some action for the license change to take effect.

If you want to upgrade your license, you need to cancel your existing license on linerate.f5.com/try, then purchase a new one.

To renew your license, go to linerate.f5.com/try and purchase the renewal.



To make the changed license take effect:

- Do one of the of the following from F5[®] LineRate[®]:

If you receive a new registration key	If you do <i>not</i> receive a new registration key
<ol style="list-style-type: none">1. Click Configure system.2. Click Licenses.3. In the Registration Key field, paste the new registration key.<ul style="list-style-type: none">• If licensing is set to auto, the system will retrieve the changed license automatically.4. If licensing activation is <i>not</i> set to auto, continue with step 5 (show dossier) in Manually Installing a Purchased License.	<ul style="list-style-type: none">• Do one of the following:<ul style="list-style-type: none">• If licensing is set to auto, tell the system to update the license information now: config license refresh base• If licensing activation is <i>not</i> set to auto, continue with step 5 (show dossier) in Manually Installing a Purchased License.

Confirming the License

To confirm that the license is enabled, go to **Configure system**, then click **Licenses** to see your license and status or a warning if the license was not enabled. It also shows the expiration date and the applicable limits.

What's Next

You are now ready to [configure management access to the system](#). You may want to review [Working with Bases](#).

Manually Installing a Purchased License

Manually Installing a Purchased License

If your F5[®] LineRate[®] hardware does not have Internet connectivity, your purchased license cannot install automatically and cannot automatically renew at the end of the license period. You can manually install the license initially and manually renew the license before the license period ends.



To manually install a purchased license:

1. Go to linerate.f5.com/try to purchase and download F5[®] LineRate[®].
 - If you already have an F5 account, click **LOGIN**.
 - Be sure to keep the email that contains your registration key. You will need the registration key to enable the Starter Edition license.
2. Copy the registration key.
3. Log in to F5[®] LineRate[®].
4. Click **Configure system**.
5. Click **Licenses**.
6. In the Registration Key field, paste the registration key.
 - You received the registration key in an email after downloading the F5[®] LineRate[®] installation file.
7. Click **Manual Activation**.
8. In the Create License area, copy the dossier text.
9. Go to <https://activate.f5.com>.
10. Click **Activate License**.
11. Paste in your dossier and follow the directions on the page.
12. Either download the license file or copy the license that displays.
13. In F5[®] LineRate Manager, click **Next**.
14. Install the license using one of the following methods:
 - If you downloaded the license file, click **Choose File** and select the license file.
 - If you copied the license file, paste the text in text box.
15. Click **Next**.
16. ????
 - After you configure licensing, the Unlicensed System message across the top of F5[®] LineRate Manager no longer displays. If the message remains, confirm the license as described in [Confirming the License](#).

17. (Optional) Configure phone home credentials.

- If your license expires, which happens only if the renewal payment is not made, the system starts using a Starter Edition license after expiration. However, the system can only activate the Starter Edition license if valid phone home credentials are configured and phone home is working.
- If you need help with your F5 credentials, go to linerate.f5.com/try and click **LOGIN** to retrieve your username or password. For more information about what phone home does, see [Phone Home Mode Commands](#).
 1. Click **Configure system**.
 2. Click **Phone Home**.
 3. Enter your F5 username and password and click **Save**.
 4. Click **System** and click **Save to Startup**.

Working with Bases

1. [Overview](#)
2. [Using Bases \(Templates\)](#)
 1. [Inheritance](#)
 2. [Finding Where A Parameter Is Set](#)
3. [What's Next](#)

Overview

A "base" in F5[®] LineRate[®] is a type of template that allows you to reuse common portions of configuration across multiple objects. Each base can inherit from another base, overriding properties from that base. This lets you create basic configurations that you can reuse and build upon.

Using Bases (Templates)

You can create a base for the following objects:

- Real server
- Virtual IP
- SSL profile



Best Practice: A best practice is to create the most basic base for each object type that includes all settings common to an object. You may also want to create a second tier of bases that inherit settings from the basic base and then add settings for variations you need.

Throughout this guide, when we create an object that can have a base, we go through creating an example base. For an example of how to create a base, see [Creating a Real Server Base](#) (CLI) or [Creating a Real Server Base](#) (F5[®] LineRate Manager).


Inheritance

Bases can inherit properties from another base. The object takes its values from the most specific configuration. For example, an object looks for its settings as follows:

1. Its own settings

2. The base it is configured to inherit from
3. The base that its base is configured to inherit from

The diagram below shows how a real server inherits its admin status in the REST API, the CLI, and F5® LineRate Manager. Here, we use the /config tree in REST.




For information about how to temporarily override the base for a setting, then return to inheriting from the base, see

- CLI— [no command usage for objects with a base](#)
- F5® LineRate Manager—[Understanding the Use Default Option](#)
- REST—[Changing Configurations Locally, Then Back to the Base](#)

Finding Where A Parameter Is Set

You can see where an object is getting its configuration using the CLI show command or the the REST origin node associated with an object, as shown in the diagram below. Here, we use the /status tree in REST.




To see where a real server is getting its configuration:

- In the CLI, type:
show real-server <real server="server" name="name">
example_host# **show real-server rs1**
show real-server rs1
Configuration
Address: 10.1.2.101:80 set locally
Admin Status: online inherited from base1
Max. Connections: 0 default
TCP Options: <none> default
...
- In F5® LineRate Manager:

1. Look at the Base Name column for the real server to see if the real server has a base.
 - In the example below, the real server rs1 has a base called base1.

Name	Description	Is Base	Base Name	IP Address	Admin Status
<input type="checkbox"/> rs1		<input type="checkbox"/>	base1	10.1.2.101:80	online

2. If the object has a base, for the parameter you want, see if the yellow triangle appears. No triangle means the parameter is inherited from the base.
 - In the example above, Admin Status does not have a yellow triangle, so the real server is inheriting its admin status from the base.
3. To confirm the inheritance, click the parameter cell.
 - In the example below, the Use Default option and (base1) show that the real server is inheriting its admin status from the base called base1.



What's Next

Now you are ready to start configuring F5® LineRate®, starting with management interfaces. See [Configuring Management Interfaces](#).

Configuring Management Interfaces

1. [Overview](#)
2. [Configuring the Host Name](#)
3. [Configuring the Management Interface](#)
4. [Configuring the admin User Password \(Amazon EC2 Only\)](#)
5. [Configuring SSH](#)
6. [What's Next](#)

Overview

An initial, basic management configuration includes the following key functions:

- Configuring the host name
- Configuring the management interface
- Configuring SSH

Configuring the Host Name

You should first configure a host name for the system. The default host name is LROS.

For this example, we are naming the host example_host.



To configure the host name:

1. Click  **Configure system**.
2. Below Host in the Name field, enter a new host name.
3. Click **Save to Startup**.

Configuring the Management Interface

You must assign an IP address to an interface to use for management. F5[®] LineRate[®] supports both static and dynamically assigned IPv4 and IPv6 addresses. In virtual environments such as VMWare or Amazon EC2, all interfaces are configured dynamically via DHCP by default. In EC2, we recommend the management interface be left configured to use DHCP.

F5[®] LineRate[®] supports both IPv4 and IPv6. You can specify the IP address and subnet mask in any of the following formats:

- 192.0.2.1/24— example of an IPv4 address with a 24-bit subnet mask using CIDR notation.
- 192.0.2.1 255.255.255.0—equivalent to above using net mask notation.
- 2001:DB8::/64—example of an IPv6 address with a 64 bit subnet mask using CIDR notation.

If you need more information about IP addresses and subnet masks, see these sites for more information:

- http://en.wikipedia.org/wiki/IP_address
- http://en.wikipedia.org/wiki/CIDR_notation

If you want the management interface to obtain its IP address dynamically using DHCP, see [ip address dhcp](#).

In this example, we are configuring one management interface with a static IP address, as shown in the diagram in [Configuring Load Balancing](#):

- Management (on em0)—10.10.10.10/24



To configure the management interface:

1. Click  **Configure system**.
2. Click **Interfaces**.
3. For interface em0, enter 10.10.10.10/24 in the IP Address column.
4. On the left side, click **System**.
5. Click **Save to Startup**.

Configuring the admin User Password (Amazon EC2 Only)

In the Amazon EC2 environment, the admin user does not have a password by default. If you want to log in using a password instead of an SSH key, you must assign a password to the admin user. See [User Name Mode Commands](#).

Configuring SSH

You should also configure secure shell (SSH) to permit access to the system using SSH.

By default, when the F5[®] LineRate[®] is first installed, SSH is automatically configured to allow incoming connections from any SSH client and to allow incoming connections to any IP address configured on the system on TCP port 22. Although this is convenient for first connecting to and configuring the system, best security practices are to limit SSH connections to only those networks and IP addresses where

access is required. So you should remove the automatic settings and replace them with settings that limit access. Typically, you want to allow access only from your management network. You can allow from more than one management network, if needed.



Caution: If you are currently logged into the system via SSH, you may disconnect yourself by changing the allow to or allow from settings if your current connection would no longer be allowed under the new settings. In order to avoid this, you may need to first add additional allow to or allow from lines, possibly make a new connection to the system, using a connection that will be allowed with the new settings, then remove any unwanted allow to or allow from lines.

In this example, we first add the setting for allowing incoming connections from hosts on the 10.10.0.0/24 management subnet and also add the setting to allow incoming connections to only the management IP address. We then remove both of the "any" settings that were added automatically.



To configure SSH:

1. Click  **Configure system**.
2. Click **SSH**.
3. In the **Allow From** tab, click  **Add**.
4. Enter 10.10.0.0/24.
5. Select the "any" and click  **Delete**.
6. Click the **Allow To** tab, click  **Add**.
7. Enter 10.10.10.10 and 22.
8. Select the "any" and click  **Delete**.
9. On the left side, click **System**.
10. Click **Save to Startup**.

What's Next

After configuring management access, you are ready to configure a reverse proxy or a forward proxy. See [Configuring a Reverse Proxy](#) or [Configuring a Forward Proxy](#).



Configuring Proxies

Overview

The sections here describe how to use F5[®] LineRate Manager to configure F5[®] LineRate[®] proxies for common use cases:

- [Configuring a Forward Proxy](#)
- [Configuring a Reverse Proxy](#)

Configuring a Forward Proxy

Overview

The F5[®] LineRate[®] forward proxy capability provides a proxy function from one network to another. A common use case for a forward proxy is for connections from your private network to the Internet.

More about Forward Proxies

A forward proxy lets you insert custom logic created with scripts. Scripts can perform a variety of functions, including gathering usage statistics, redirecting requests to your own cache, blocking of access to specific sites, managing cookies, and much more.


To use a forward proxy effectively, be sure to create and attach a virtual IP that includes the range of Internet IP addresses you want to go through the forward proxy.



Caution: When attaching a virtual IP to a forward proxy, the virtual IP must not include any of the system's own IP addresses. For a virtual IP with a single IP address, do not set the virtual IP's IP address to one of the system's own IP addresses. For a virtual IP with a range of addresses, you must ensure that the IP address range does not contain any of the system's own IP addresses. This may mean you need to break the virtual IP into multiple virtual IPs. See [Configuring a range for a virtual IP with forward proxy](#) for more detail and an example.

Example Configuration

The diagram below shows the configuration for a forward proxy from a private network to the Internet. Requests from clients in the private network go to the default gateway, which is a F5[®] LineRate[®] interface configured with the inside IP address (for example, 10.1.2.110). The virtual IP configured with a range of IP addresses listens for requests on the configured TCP port. This example shows the forward proxy configured to intercept TCP port 80, the typical port for HTTP traffic. The virtual IP passes the request to the forward proxy, where a script may intercept the request and perform specific functions, then passes the request to the outside interface and on to the Internet.



What's Next

After determining that you want to configure a forward proxy, you are ready to configure the data interfaces. See [Configuring Data Interfaces](#).

Configuring Data Interfaces

1. [Overview](#)
2. [Configuring the Data Interfaces](#)
3. [Configuring the Default IP Route \(Gateway\) on F5® LineRate®](#)
4. [Configuring the Default IP Route on Clients](#)
5. [What's Next](#)

Overview

An initial, basic configuration for a forward proxy includes the following key functions:

- Configuring the data interfaces
- Configuring the default IP route (gateway) on F5® LineRate®
- Configuring the default IP route on clients

Configuring the Data Interfaces

You must assign an IP address to at least two interfaces to use for data. The F5® LineRate® software supports both IPv4 and IPv6.

For general information about the IP addresses and formats, see [Configuring Management Interfaces](#).



Note: For virtual environments, all interfaces are configured by default to obtain their IP address via DHCP and to set the system default route with the lowest numbered DHCP-enabled interface. Because the proxy configuration depends on the IP address of the "outside" data interface, you likely want to disable [DHCP](#) for that interface and assign a static IP address.



Note: For Amazon EC2, interfaces by default have IP source/destination check enabled. This will prevent any packets from leaving/arriving at the EC2 interface (eni) that don't match the IP address AWS has associated with that interface. For forward proxy configurations, you must disable source/destination check on the "inside" interface. See http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html#change_source_dest_check.

In this example, we are configuring two interfaces as shown in the diagram in [Configuring a Forward Proxy](#):

- Inside to the virtual IP (on em1)—10.1.2.110/24
- Outside to the Internet (on em2)—192.0.2.10/24



To configure the data interfaces:

1. Click **Configure system**.
2. Click **Interfaces**.
3. For interface em1, enter 10.1.2.110/24 in the IP Address column.
4. For interface em2, enter 192.0.2.10/24 in the IP Address column.
5. On the left side, click **System**.
6. Click **Save to Startup**.

Configuring the Default IP Route (Gateway) on F5[®] LineRate[®]

You should also configure the default IP route (gateway). For this forward proxy example, the inside IP address is the default gateway for all clients to the Internet.

For general information about the IP addresses and formats, see [Configuring Management Interfaces](#).



To configure the default IP route:

1. Click **Configure system**.
2. Click **IP Routes**.
3. Click **Add static route** and use the following settings:

Field	Example setting
Destination	default—This is equivalent to 0.0.0.0/0.
Gateway	192.0.2.1

4. On the left side, click **System**.
5. Click **Save to Startup**.

Configuring the Default IP Route on Clients

You should also configure the default IP route on clients to work with the forward proxy. In this example, each client's default IP route should point to the inside IP address of F5[®] LineRate[®] (10.1.2.110).

What's Next

After configuring the data interfaces, you are ready to configure the virtual IP. See [Configuring the Virtual IP](#).

Configuring the Virtual IP

1. [Overview](#)
2. [Creating a Virtual IP Base](#)
3. [Creating the Virtual IPs](#)
4. [What's Next](#)

Overview

In this forward proxy example, you want the virtual IP to listen for requests from all possible Internet IP addresses on port 80. Therefore, you want to configure the virtual IPs that cover the full range from 0.0.0.0 to 255.255.255.255, but that exclude the IP addresses configured on any F5[®] LineRate[®] interface.

You can attach a virtual IP to only one forward proxy. You cannot use the same virtual IP for both a reverse proxy and forward proxy.



To configure a virtual IP address, complete the following tasks:

1. Create a virtual IP base.
2. Create a virtual IP.

Creating a Virtual IP Base

We recommend creating one or more virtual IP bases. For general information about bases in F5[®] LineRate[®], see [Working with Bases](#). A base lets you configure the most common settings that you want for your virtual IPs. You can also create more than one virtual IP base for settings that you need to be different or more specific for some virtual IPs.

In this example, we are creating a single virtual IP base called vipbase_web1. We recommend giving each virtual IP base a meaningful name that helps identify the base. For example, you might use the application type (such as serving similar web content) or security settings (such as SSL) in the name.



To create an example virtual IP base:

1. In the configuration table, click **Virtual IPs**.

2. At the bottom of the table, click **+** and create a new virtual IP base called `vipbase_web1`.
 - Be sure to select the **Is Base** check box.
 - A new row for the virtual IP base displays in the configuration table.
3. For the new virtual IP base, configure the following:

Column	Example setting	Description
Admin Status	Set Locally, online	Brings the virtual IP online, so it is ready for use.
Service Type	HTTP	Sets the service type to be HTTP for layer 7 load balancing of web traffic.
Keepalive Timeout	5	This sets the keep alive timeout to 5 seconds. This is the time the system waits for a specific client to send a request before closing the connection, reclaiming connection resources. For most use cases, this setting will affect the number of simultaneous connections that the system will have open. A lower setting will usually result in fewer simultaneous open connections. A good rule of thumb is to set this number no higher than 500,000 divided by the number of expected connections per second at peak load. For example, if the load balancer is expected to process up to 100,000 connections per second, 500,000 divided by 100,000 is 5. So the setting should be 5 seconds in this example.

4. Click  **Configure system**.
5. Click **Save to Startup**.

Creating the Virtual IPs



Note: Be sure that you also configure the same IP address on an interface as described in [Configuring Data Interfaces](#).

After creating the virtual IP base, you can create a virtual IP. We recommend giving each virtual IP a meaningful name that helps identify the virtual IP. For example, you might use the application or service type (such as serving similar web content) or security settings (such as SSL) in the name.

For this example, we are configuring a range of IP addresses that cover the entire range of IP addresses on port 80 (HTTP). However, to avoid creating a loop, we must create three virtual IPs to cover the entire range, with IP address ranges that exclude the IP addresses configured on the interfaces (management, inside, and outside). For more information about virtual IP addresses for forward proxies, see [Configuring a range for a virtual IP with forward proxy](#).



To create the virtual IPs:

1. In the configuration table, click **Virtual IPs**.
2. At the bottom of the table, click **+** and create a new virtual IP called vipInternet1.
 - A new row for the virtual IP displays in the configuration table.
3. For the new virtual IP, configure the following:

Column	Example setting	Description
Base Name	Set Locally, vipbase_web1	Sets the virtual IP to use the base.
IP Address Range	Start—0.0.0.0 End—9.255.255.255 Port—802	Configures the IP address range and port.

4. At the bottom of the table, click **+** and create a new virtual IP called vipInternet2.
5. For the new virtual IP, configure the following:

Column	Example setting	Description
Base Name	Set Locally, vipbase_web1	Sets the virtual IP to use the base.
IP Address Range	Start—11.0.0.0 End—192.0.2.9 Port—80	Configures the IP address range and port.

6. At the bottom of the table, click **+** and create a new virtual IP called vipInternet3.
7. For the new virtual IP, configure the following:

Column	Example setting	Description
Base Name	Set Locally, vipbase_web1	Sets the virtual IP to use the base.
IP Address Range	Start—192.0.2.11 End—255.255.255.255 Port—80	Configures the IP address range and port.

8. Click  **Configure system**.
9. Click **Save to Startup**.

What's Next

After configuring the virtual IP, you are ready to configure the forward proxy itself. See [Configuring the Forward Proxy](#).

Configuring the Forward Proxy

Creating a Forward Proxy

For this example, you now need to configure the forward proxy itself. You must give the forward proxy a name, attach the virtual IP to it, and put the forward proxy online. In addition, it is a best practice to always set the response timeout and the keepalive timeout.

To create an example forward proxy:

1. In the configuration table, click **Forward Proxies**.
2. At the bottom of the table, click **+** and create a new forward proxy called fp1.
 - A new row for the forward proxy displays in the configuration table.
3. For the new forward proxy, configure the following:

Column	Example setting	Description
Virtual IPs	vipInternet1 vipInternet2 vipInternet3	Attaches the virtual IPs.
Admin Status	Set Locally, online	Brings the forward proxy online, so it is ready for use.
Service Type	Set Locally, HTTP	Sets the service type to HTTP for layer 7 web traffic.
Response Timeout	5	Sets the response timeout to 5 seconds. This is the time the system waits for the HTTP server to respond to a request. If the server does not respond in this time, the system sends an HTTP 504 error to the client and closes the connection. This ensures that connections are closed and system resources are reclaimed if the server does not respond.

Column	Example setting	Description
		Consider the amount of time you're willing to wait for a target web server to respond to any request. The response-timeout must always be configured to be higher than the amount of time it takes for any of the web servers to respond to a request.
Keepalive Timeout	10	Sets the keepalive timeout to 10 seconds. This is the time the system waits for a new HTTP request from a client to a server. If there are no active HTTP transactions (that is, no active requests or responses) to a server for the specified time (in seconds), the system closes the TCP connection to the server, reclaiming resources. This also helps avoid problems that some HTTP servers have when connections are kept open indefinitely.

4. Click  **Configure system**.
5. Click **Save to Startup**.



Configuring a Reverse Proxy

1. [Overview](#)
2. [More about Reverse Proxies](#)
3. [What's Next](#)


Overview

The F5[®] LineRate[®] system provides a full-proxy front-end to the web and secure web services on back-end servers of the application. As a full proxy, F5[®] LineRate[®] provides a logical front-end to the client's connection to the web and secure web services of the application.

More about Reverse Proxies

To successfully manage load, health, and availability of the application, F5[®] LineRate[®] constantly monitors the traffic flows from the client to the back-end servers to ensure the delivery of the application services. F5[®] LineRate[®] provides several discrete services including health monitoring, transaction statistics, and traffic management to ensure balanced usage of back-end servers and rapid clients service.

The diagram below shows a basic F5[®] LineRate[®] configuration for a load balancing use case. You can have multiple load balancer configurations in the F5[®] LineRate[®] software, as shown by the second configuration in the diagram. Clients access the virtual IP, which goes through a virtual server. Based on the load balancing algorithm set in the virtual server, it passes client requests through the real servers and the inside IP address to the web servers.



What's Next

After determining that you want to configure a reverse proxy, you are ready to configure the data interfaces. See [Configuring Data Interfaces](#).

Configuring Data Interfaces

1. [Overview](#)
2. [Configuring the Data Interfaces](#)
3. [Configuring the Default IP Route \(Gateway\)](#)
4. [What's Next](#)

Overview

An initial, basic configuration includes the following key functions:

- Configuring the data interfaces
- Configuring the default route (gateway)

Configuring the Data Interfaces

The F5[®] LineRate[®] software supports both IPv4 and IPv6.

The recommended configuration for performance and security reasons is to configure IP addresses on at least two interfaces.

For general information about the IP addresses and formats, see [Configuring Management Interfaces](#).



Note: For virtual environments, all interfaces are configured by default to obtain their IP address via DHCP and to set the system default route with the lowest numbered DHCP-enabled interface. Because the proxy configuration depends on the IP address of the "outside" data interface, you likely want to disable [DHCP](#) for that interface and assign a static IP address.

In this example, we are configuring two data interfaces as shown in the diagram in [Configuring Load Balancing](#):

- Outside for the virtual IP (on em1)—192.0.2.1/24
- Inside to web servers (on em2)—10.1.2.1/24



To configure the data interfaces:

1. Click  **Configure system**.

2. Click **Interfaces**.
3. For interface em1, enter 192.0.2.1/24 in the IP Address column.
4. For interface em2, enter 10.1.2.1/24 in the IP Address column.
5. On the left side, click **System**.
6. Click **Save to Startup**.

Configuring the Default IP Route (Gateway)

You should also configure the default IP route (gateway). You can configure additional static IP routes, as needed, to permit access to your networks.

For general information about the IP addresses and formats, see [Configuring Management Interfaces](#).



To configure the default IP route:

1. Click  **Configure system**.
2. Click **IP Routes**.
3. Click  **Add static route** and use the following settings:

Field	Example setting
Destination	default—This is equivalent to 0.0.0.0/0.
Gateway	192.0.2.2

4. On the left side, click **System**.
5. Click **Save to Startup**.

What's Next

After configuring the data interfaces, you are ready to configure load balancing. See [Configuring Load Balancing](#).

Configuring Load Balancing


1. [Overview](#)
2. [Configuring a Virtual IP Address](#)
 1. [Creating a Virtual IP Base](#)
 2. [Creating a Virtual IP](#)
3. [Configuring a Health Monitor](#)
4. [Configuring Real Servers](#)
 1. [Creating a Real Server Base](#)
 2. [Creating Real Servers](#)
 3. [Creating a Real Server Group](#)
5. [Configuring a Virtual Server](#)
6. [What's Next](#)

Overview

The figure below shows a very simple example load balancing architecture, with all of the basic elements you need. It includes the specific names and IP addresses that we use throughout the rest of this section.

This example architecture assumes the following:

- The system has three physical interfaces.
- The architecture uses two real servers.
- You have two web servers.



The remainder of this section walks through how to create this architecture with the F5[®] LineRate[®] software.

Configuring a Virtual IP Address



Note: Be sure that you also configure the same IP address on an interface as described in [Configuring Data Interfaces](#).

The load balancer requires at least one virtual IP address. The virtual IP address is a configuration object that represents the outside interface that clients connect to. In this example, we will configure one virtual IP, but you can configure more, depending on your needs.



To configure a virtual IP address, complete the following tasks:

1. [Create a virtual IP base](#).
 2. [Create a virtual IP](#).
-

Creating a Virtual IP Base

We recommend creating one or more virtual IP bases. For general information about bases in F5[®] LineRate[®], see [Working with Bases](#). A base lets you configure the most common settings that you want for your virtual IPs. You can also create more than one virtual IP base for settings that you need to be different or more specific for some virtual IPs.

In this example, we are creating a single virtual IP base called vipbase_web1. We recommend giving each virtual IP base a meaningful name that helps identify the base. For example, you might use the application type (such as serving similar web content) or security settings (such as SSL) in the name.



To create an example virtual IP base:

1. In the F5[®] LineRate Manager configuration table, click **Virtual IPs**.
2. At the bottom of the table, click **+** and create a new virtual IP base called vipbase_web1.
 - Be sure to select the Is Base check box.
 - A new row for the virtual IP base displays in the configuration table.
3. For the new virtual IP base, configure the following:

Column	Example setting	Description
Admin Status	Set Locally, online	Brings the virtual IP online, so it is ready for use.
Keepalive Timeout	5	This sets the keep alive timeout to 5 seconds. This is the time the system waits for a specific client to send a request before closing the connection, reclaiming connection resources. For most use cases, this setting will affect the number of simultaneous connections that the system will have open. A lower setting will usually result in fewer simultaneous open connections. A good rule of thumb is to set this number no higher than 500,000 divided by the number of expected connections per second at peak load. For example, if the load balancer is expected to process up to 100,000 connections per second, 500,000 divided by 100,000 is 5. So the setting should be 5 seconds in this example.

4. Click  **Configure system**.
5. Click **Save to Startup**.

Creating a Virtual IP

After creating the virtual IP base, you can create a virtual IP. We recommend giving each virtual IP a meaningful name that helps identify the virtual IP. For example, you might use the application or service type (such as serving similar web content) or security settings (such as SSL) in the name.

For this example, we are using the IP address of the outside interface (em1) that we configured already to create the virtual IP, and we are using the virtual IP base we already created. You must also include the TCP port number on which the clients will contact the load balancer.



To create a virtual IP:

1. In the F5[®] LineRate Manager configuration table, click **Virtual IPs**.
2. At the bottom of the table, click **+** and create a new virtual IP called vipweb1.
 - A new row for the virtual IP displays in the configuration table.
3. For the new virtual IP, configure the following:

Column	Example setting	Description
Base Name	Set Locally, vipbase_web1	Sets the virtual IP to use the base.
IP Address	192.0.2.1 and 443	Configures the IP address and port.

4. Click  **Configure system**.
5. Click **Save to Startup**.

Configuring a Health Monitor

A health monitor can monitor multiple real servers. The health monitor for web servers (HTTP) opens a connection to the web server, sends an HTTP request to the web server for something (possibly a specific web page), looks at the response, and determines if the response is correct. The configuration settings on the health monitor object determine what to request from the web server and what the response should be.

In this example, we are creating a single health monitor called hm_web1. We recommend giving each health monitor a meaningful name that helps identify the health monitor. For example, you might use the application or service (such as serving similar web content) or how you are monitoring in the name.



To create a health monitor:

1. In the F5[®] LineRate Manager configuration table, click **Health Monitors**.
2. At the bottom of the table, click **+** and create a new health monitor called hm_web1.
 - A new row for the health monitor displays in the configuration table.
3. For the new health monitor, configure the following:

Column	Example setting	Description
Admin Status	Set Locally, online	Brings the health monitor online, so it is ready for use.
Timeout	1	Sets the timeout to 1 second. The health monitor will determine an individual health probe to fail if it does not respond within this time. One use of this setting is to test the server's response time.
Server Down	Threshold—8 Window—10	Sets the threshold for marking the server DOWN to the failure of 8 out the last 10 health probes. If the health probe fails 8 out of the last 10 times, the system takes the server offline.
Server Up	Threshold—9 Window—10	Sets the threshold for marking the server UP to the success of 9 out the last 10 health probes. If the health probe succeeds 9 out of the last 10 times, the system puts the server back online.
Interval	5	Sets the health monitor to start a health check every 5 seconds.
HTTP Request Method	GET	Sets the type of request the health monitor will send to a GET request

Column	Example setting	Description
HTTP Request Target	health.html	Sets the specific web page that the health monitor will request. If the health monitor is able to retrieve the page, receiving a 200 OK response from the server, the server's health probe is deemed successful.

4. Click  **Configure system**.
5. Click **Save to Startup**.

Configuring Real Servers

Real servers are another required configuration object of the F5[®] LineRate[®] load balancer. Real servers represent and point to actual web servers that the load balancer is distributing client requests to.



To configure real servers, complete the following tasks:

1. [Create a real server base](#).
2. [Create a real server](#).
3. [Create a real sever group](#).

Creating a Real Server Base

We recommend creating one or more real server bases. For general information about bases in F5[®] LineRate[®], see [Working with Bases](#). A base lets you configure common settings that you want for your real servers. You can also create more than one real server base for settings that you need to be different for some real servers.

In this example, we are creating a single real server base called `rsbase_web`. We recommend giving each real server base a meaningful name that helps identify how the base will be used. For example, you might use the application type (such as serving similar web content), hardware capabilities (such as CPU or memory), or security settings (such as SSL) in the name.



To create a real server base:

1. In the F5[®] LineRate Manager configuration table, click **Real Servers**.
2. At the bottom of the table, click **+** create a new real server base called `rsbase_web`.
 - Be sure to select the **Is Base** check box.
 - A new row for the real server base displays in the configuration table.

- For the new real server base, configure the following:

Column	Example setting	Description
Admin Status	Set Locally, online	Brings the real server base online, so it is ready for use.
Health Monitors	hm_web1	Attaches a health monitor called hm_web1.
Max Connect	Set Locally, 1000	Sets the maximum connections to the real server at 1000. F5 [®] LineRate [®] will not open more than 1000 connections to any server.
Keepalive Timeout	Set Locally, 10	Sets the keepalive timeout of the HTTP service to 10 seconds. If there are no requests that get sent to a connection for 10 seconds, the load balancer closes the connection to the server, reclaiming resources. This can help avoid problems that some web servers have when connections are kept open indefinitely.
Response Timeout	Set Locally, 60	This sets the response-timeout of the HTTP service to 60 seconds. The load balancer closes the connection if the HTTP server takes longer than 60 seconds to respond to a request. Consider the amount of time the web server takes to respond to any request. The response-timeout on the load balancer must always be configured to be higher than the amount of time it takes for any of the web servers to respond to a request.
Response Idle Timeout	Set Locally, 60	Sets the response-idle-timeout of the HTTP service to 60 seconds. The load balancer closes the connection if it takes longer than 60 seconds either to receive any part of the response from the HTTP server or to transmit any part of the response to the client. Consider the size of a typical response for your application as well as the user environment to set this value. For example, an application where users download HD videos using mobile devices would need longer timeout than simple web pages due to mobile bandwidth and device processing limitations.

- Click  **Configure system**.
- Click **Save to Startup**.

Creating Real Servers

After creating the real server base, you can create a real server. We recommend giving each real server a meaningful name that helps identify the real server. For example, you might use the application type (such as serving similar web content), hardware capabilities (such as CPU or memory), or security settings (such as SSL) in the name.

In this example, we are creating two real servers (rsweb1 and rsweb2) that inherit properties from our real server base and assigning them the IP addresses of two actual web servers. You must also include the TCP port number on which the load balancer will contact the server.



To create real servers:

1. In the F5[®] LineRate Manager configuration table, click **Real Servers**.
2. At the bottom of the table, click **+** create a new real server called rsweb1.
 - A new row for the real server displays in the configuration table.
3. For the new real server, configure the following:

Column	Example setting	Description
Base Name	Set Locally, rsbase_web	Sets the real server to use the base called rsbase_web.
IP Address	10.1.2.101 and 8080	Sets the IP address and port.

4. At the bottom of the table, click **+** create a new real server called rsweb2.
 - A new row for the real server displays in the configuration table.
5. For the new real server, configure the following:

Column	Example setting	Description
Base Name	Set Locally, rsbase_web	Sets the real server to use the base called rsbase_web.
IP Address	10.1.2.102 and 8080	Sets the IP address and port.

6. Click  **Configure system**.
7. Click **Save to Startup**.

Creating a Real Server Group

You can use real server groups create logical groups of real servers. A real server can be a member of multiple groups. You can use the groups to show information about the real servers or to attach the group to a virtual server for load balancing.

We recommend giving each real server group a meaningful name that helps identify the group use. For example, you might set up a real server group based on the application, floor location (all servers in a specific rack), or data center (all servers in data center).

In this example, we are creating a real server group (rsgroup_websilo1) and are adding the two real servers we already created, using a regular expression. We will then attach this real server group to a virtual server.



To create a real server group:

1. In the F5[®] LineRate Manager configuration table, click **Real Server Groups**.
2. At the bottom of the table, click **+** and create a new real server group called rsgroup_websilo1.
 - A new row for the real server group displays in the configuration table.
3. For the new real server group, configure the Member RegEx using:

```
rsweb.*
```

4. Click  **Configure system**.
5. Click **Save to Startup**.

Configuring a Virtual Server

Each load balancing (reverse proxy) configuration requires at least one virtual server. The virtual server is a configuration object that acts as a reverse proxy and ties together one or more virtual IPs and real servers. You also set the load balancing algorithm on the virtual server.

We recommend giving each virtual server a meaningful name that helps identify the server use. For example, you might name a virtual server based on the application and the resources that the virtual server is load balancing traffic to (real servers).

In this example, we are creating a virtual server and are attaching the virtual IP and real server group to it. We will also set the load balancing algorithm.



To create a virtual server:

1. In the F5[®] LineRate Manager configuration table, click **Virtual Servers**.
2. At the bottom of the table, click **+** create a new virtual server called websilo1.
 - A new row for the virtual server displays in the configuration table.
3. For the new virtual server, configure the following:

Column	Example setting	Description
Load Balancing Algorithm	round robin	Sets the algorithm the system uses for load balancing.
Virtual IPs	vipweb1	Attaches the virtual IP to the virtual server.
Real Server Groups	rsgroup_websilo1	Attaches the real server group to the virtual server

4. Click  **Configure system**.
5. Click **Save to Startup**.

What's Next

After you configure the load balancer, you can test the load balancer ([Monitoring and Troubleshooting Load Balancing](#)) or add security by configuring SSL ([Configuring SSL](#)).

Monitoring and Troubleshooting Load Balancing

1. [Overview](#)
2. [Solutions](#)
 1. [Seeing "404 Not Found" errors when attempting to access a new virtual server, but real-server logs do not show any error_log messages that match](#)
 2. [Seeing "502 Bad Gateway" error messages when attempting to access a virtual server](#)
 3. [Seeing "couldn't connect to host" error messages when attempting to access a virtual server](#)
 4. [Some of my real servers may be more heavily loaded than others. How can I see this?](#)
3. [Related](#)

Overview

This section provides a few possible errors you might encounter while setting up a load balancing or SSL offload use case and how to resolve those issues.

Solutions

Seeing "404 Not Found" errors when attempting to access a new virtual server, but real-server logs do not show any error_log messages that match

The most common issue surrounding 404 errors, especially with a new virtual server, is a missing "default" virtual IP for the virtual server.

When the virtual server is using the service type HTTP, the virtual server parses the HTTP request headers and looks for a hostname match to a given virtual server. If no hostname match is found, the proxy will respond to the client with a 404 error and not pass the request to the real server. This will also happen with requests that request the URL with the IP address of the virtual IP and not using an FQDN.

For example, the following virtual server has a hostname set and no default virtual IP for requests.

```
!  
virtual-server vsl  
  service http  
    tcp-multiplex  
    hostname www.f5.com  
  attach virtual-ip vip1
```

```
attach real-server group reals
!
```

When a request is made to the IP address of the virtual-ip we receive a "404 Not Found" message:

```
curl 201.0.50.1:8080/
<html><head><title>Status 404 Not Found</title></head><body><h1>Status 404 Not
Found</h1></body></html>
```

If we append the Host: header to the request using our defined hostname, the error goes away:

```
curl 201.0.50.1:8080/ -H "Host: www.f5.com"
<html>
<head>
<title>Welcome to the virtual-server!</title>
</head>
<body bgcolor="white" text="black">
<center><h1>Welcome to the virtual-server!</h1></center>
</body>
</html>
```

Similarly, if we use a Host: header name that is not defined for the virtual server, we will get a 404 error:

```
curl 201.0.50.1:8080/ -H "Host: www.lineratesomething.com"
<html><head><title>Status 404 Not Found</title></head><body><h1>Status 404 Not
Found</h1></body></html>
```

Attaching the virtual IP with the "default" setting will allow all of these unmatched requests to be proxied to the back-end real servers.

```
!
virtual-server vs1
service http
tcp-multiplex
hostname www.f5.com
attach virtual-ip vip1 default
attach real-server group reals
!
```

These 404 response codes can be monitored on the system with SNMP or viewed from the CLI. The load-balancer statistics are where we can find the current count of HTTP internal response codes, as none of these will match a defined virtual server:

```
lrs01# show load-balancer statistics detailed
<output omitted>
httpInternalResp404: 6
```

Seeing "502 Bad Gateway" error messages when attempting to access a virtual server

The most common issue surrounding 502 errors, especially with a new virtual server, is when the real servers are either marked down with a failed health monitor or their admin status is set to offline. The default setting for a new real server base is admin status offline.

Check the real servers to make sure they are all in admin status online.

```
lrs01# show real group reals
reals Group Members
Name Address      Port  Svc  Admin  Health  Conns  Rx Mbps  Tx Mbps
-----
rs1   201.0.51.1    8080  http offline up      0       0.0     0.0
rs2   201.0.51.2    8080  http offline up      0       0.0     0.0
rs3   201.0.51.3    8080  http offline up      0       0.0     0.0
rs4   201.0.51.4    8080  http offline up      0       0.0     0.0
rs5   201.0.51.5    8080  http offline up      0       0.0     0.0
```

If the real servers are offline, change their real server base to admin status online or edit each real server individually to set the admin status online.

Check the real servers to make sure they are all passing their health monitor checks and are marked as "up."

```
lrs01# show real group reals
reals Group Members
Name Address      Port  Svc  Admin  Health  Conns  Rx Mbps  Tx Mbps
-----
rs1   201.0.51.1    8080  http  online down   0       0.0     0.0
rs2   201.0.51.2    8080  http  online down   0       0.0     0.0
rs3   201.0.51.3    8080  http  online down   0       0.0     0.0
rs4   201.0.51.4    8080  http  online down   0       0.0     0.0
rs5   201.0.51.5    8080  http  online down   0       0.0     0.0
```

Check the health monitor settings to make sure that the configuration is correct and that the request target is available on each of the real servers.

```
lrs01# show real group reals
reals Group Members
Name Address      Port  Svc  Admin  Health  Conns  Rx Mbps  Tx Mbps
-----
rs1   201.0.51.1    8080  http  online up    0       0.0     0.0
rs2   201.0.51.2    8080  http  online up    0       0.0     0.0
```

Seeing "couldn't connect to host" error messages when attempting to access a virtual server

The most common issue surrounding "couldn't connect to host" error messages is when the virtual IP is admin status offline. The default setting for a virtual IP base is admin status offline.

Check your virtual IP or virtual server to make the virtualIP is online:

```
lrs01# show virtual-server vs1
Configuration
  LB Algorithm:          round-robin
  Service Type:         http
  Real Server Groups:
    Name                 Weight
    reals                1 (default)
  Real Servers:
    <none>
  Virtual IPs:
    Name                 Address      Port Svc  Admin
    vip1                201.0.50.1  8080 http offline
```

Make sure your virtual IP base or virtual IP is set to admin status online:

```
!
virtual-ip vip1
  ip address 201.0.50.1 8080
  admin-status offline
!
```

Some of my real servers may be more heavily loaded than others. How can I see this?

Use a real server group and "show real group <name>" to get all of the traffic and connection distribution counts for each real server.

```
!
real-server rs1 ip 201.0.51.1 8080 base rsb
real-server rs2 ip 201.0.51.2 8080 base rsb
real-server rs3 ip 201.0.51.3 8080 base rsb
real-server rs4 ip 201.0.51.4 8080 base rsb
real-server rs5 ip 201.0.51.5 8080 base rsb
!
real-server group reals
```

```
members by regex "rs.*"
```

```
!
```

```
lrs01# show real group reals
```

```
reals Group Members
```

```
Name Address Port Svc Admin Health Conns Rx Mbps Tx Mbps
```

rs1	201.0.51.1	8080	http	online	up	2	87.0	0.2
rs2	201.0.51.2	8080	http	online	up	2	87.0	0.2
rs3	201.0.51.3	8080	http	online	up	2	87.1	0.2
rs4	201.0.51.4	8080	http	online	up	2	87.0	0.2
rs5	201.0.51.5	8080	http	online	up	2	86.9	0.2

Related

- [Configuring Load Balancing](#)

Configuring SSL

1. [SSL Types Supported in the F5® LineRate® Software](#)
2. [SSL Termination](#)
3. [Configuring SSL Termination](#)
4. [Configuring SSL Initiation](#)
5. [What's Next](#)


Secure Sockets Layer (SSL) and Transport Layer Security (TLS) are closely related technologies that provide communication security over an insecure network, such as the Internet. TLS is a standardized protocol, defined by IETF RFCs, and is the successor to the non-standardized SSL protocol. The F5® LineRate® software supports both TLS and SSL, for both service type TCP and service type HTTP, but the system and documentation refers to both protocols collectively as "SSL," following the most common industry terminology.

SSL Types Supported in the F5® LineRate® Software

The F5® LineRate® software supports two types of SSL connections:

- SSL termination—SSL connection from the client to the F5® LineRate® load balancer.
- SSL initiation—SSL connection from the F5® LineRate® load balancer to the web server.

The diagram below shows the two types of SSL.



By using the SSL termination feature in F5® LineRate®, you can move the computationally intensive SSL processing off your web servers and onto F5® LineRate®, allowing your web servers to concentrate on performing application tasks. Or, if your application requires greater security on your internal network, you can use SSL initiation together with SSL termination to provide end-to-end SSL security, while still allowing the F5® LineRate® to do full layer 7 load balancing.

SSL Termination

Before beginning to set up SSL termination, you will need the following:

- Primary certificate file that identifies the website you wish to set up on the F5[®] LineRate[®] system
- Private key file that corresponds to the primary certificate
- Chain certificate files (also called intermediate certificates) that correspond to the primary certificate are only required if their primary certificate uses them.



To set up SSL termination, you must complete the following tasks:

1. Set up the private key.
2. Set up certificates.
3. Configure SSL on the virtual IP.

For more information, see [Managing SSL](#).

Configuring SSL Termination

You set up a private key object to correspond to each primary certificate you need. The system supports using one private key to generate more than one primary certificate and the use of separate private keys for individual primary certificates.

You need access to your private key file. The F5[®] LineRate[®] software supports keys in PEM format.



Best Practices:

- Configure primary certificates and corresponding keys for each cipher type that clients may use, then attach the configured certificate/key pairs to the SSL profile. For example, for clients that support ECC or RSA ciphers, you may want to configure and attach both RSA and ECC certificates and keys.
- Give each key a meaningful name that helps identify the key. For example, you might use the domain name or security settings in the name.

In this example, we will create a key object in the configuration, give it a name (key_secure.example.com), and import the key text.

To set up certificates, you must have access to your certificate files. The F5[®] LineRate[®] software supports PEM format certificates.




Best Practices:

- Give each certificate a meaningful name that helps identify the certificate. For example, you might use the domain name or security settings in the name.
-

-
- Configure primary certificates and corresponding keys for each cipher type that clients may use, then attach the configured certificate/key pairs to the SSL profile. For example, for clients that support ECC or RSA ciphers, you may want to configure and attach both RSA and ECC certificates and keys.
-

In this example, we will assume that your primary certificate only requires a single chain certificate. So we will create two certificate objects in the configuration, give them names (`cert_prim_secure.example.com` and `cert_chain_secure.example.com1`, as shown below), and paste the certificate text into them.



Certificate chain example

In this example, `cert_prim_secure.example.com` is the certificate that identifies the website "secure.example.com" and `cert_chain_secure.example.com1` is the chain certificate required for that primary certificate. We have to create an SSL profile (`ssl_prof_secure.example.com`) and attach the private key, and the certificates to it. We then attach the SSL profile to the virtual IP. We recommend giving each profile a meaningful name that helps identify it. For example, you might use the domain name or security settings in the name.

Attaching an SSL profile to a virtual IP configures that virtual IP to always use SSL. It will no longer accept connections from clients unless they perform SSL negotiation.



Note: When installing an SSL certificate on the F5[®] LineRate[®], you must remove any passphrases from the certificate file.

When the system first starts the SSL web service, it cannot wait for user input to enter a passphrase before the services will start.



To remove a passphrase using openssl:

1. Make a copy of your SSL key file keeping the original intact.

2. Use openssl to enter the passphrase and output a new key file:
`openssl rsa -in key.pem -out newkey.pem`
3. Use this new file newkey.pem as your SSL private key for upload into F5[®] LineRate[®].



To set up the private key for SSL termination:

1. Copy the private key and each certificate file to the F5[®] LineRate[®] system or any computer accessible on your network.
2. In F5[®] LineRate Manager, click  **Configure system**.
3. On the left side, click **SSL**.
4. In the **Certificates** tab, click  **Add new certificate**, enter the name cert_prim_secure.example.com, select and import the certificate file.
5. Add another certificate called cert_chain_secure.example.com1.
6. Click **Private Keys** tab.
7. Click  **Add new key**, enter the name key_secure.example.com, select and import the key file
8. Close the System Configuration window.
9. In the configuration table, click **SSL Profiles**.
10. At the bottom of the table, click **+** and create a new SSL profile called ssl_prof_secure.example.com.
 - Leave Is Base deselected.
 - A new row for the SSL profile displays in the configuration table.
11. For the new SSL profile, configure the following:

Column	Example setting	Description
Primary Certificate	Set Locally, cert_prim_secure.example.com	Attaches the primary certificate to the SSL profile.
Private Key	Set Locally, key_secure.example.com	Attaches the private key to the SSL profile.
Chain Certificate	Set Locally, cert_chain_secure.example.com1	Attaches the chain certificate to the SSL profile.

12. Click **Virtual IPs**.
13. For vipweb1 that we created already, click the SSL Profile column, select Set Locally, and use ssl_prof_secure.example.com as the Name.
14. Click  **Configure System**.
15. Click **Save to Startup**.

Configuring SSL Initiation

In many cases, the default settings for SSL initiation work well.

In this example, we will set up the SSL initiation profile (ssl_prof_init1), using the defaults, and attach it to the real server base (rsbase_web). We recommend giving each profile a meaningful name that helps identify it. For example, you might use the security settings in the name.

Attaching an SSL profile to a real server configures that real server to always use SSL. If the web server is not configured to accept an SSL connection from the F5[®] LineRate[®] system, the system will not be able to send traffic to that web server.



To configure SSL initiation:

1. In the F5[®] LineRate Manager configuration table, click **SSL Profiles**.
2. At the bottom of the table, click **+** and create a new SSL profile called ssl_prof_init1.com.
 - Leave Is Base deselected.
 - A new row for the SSL profile displays in the configuration table.
3. Click **Real Servers**.
4. For rsbase_web that we created already, click the SSL Profile column, select Set Locally, and use ssl_prof_init1 as the Name.
5. Click  **Configure System**.
6. Click **Save to Startup**.

What's Next

After you have completed all the sections of *Configuring a Reverse Proxy*, you have a basic reverse proxy setup that uses SSL.



Configuring A/B Testing

Configuring Scripts

[Example from David? create, edit, errors, etc.]



LineRate Manager Reference

Overview

This section is a reference for using F5[®] LineRate[®] Manager.

Contents

The guide is broken into the following sections:

- [Configuration Tables](#)
- [Dashboards](#)
- [Forward Proxies](#)
- [Health Monitors](#)
- [Interfaces](#)
- [IP Routes](#)
- [Licenses](#)
- [Line and Area Charts](#)
- [NTP](#)
- [Phone Home](#)
- [Real Servers](#)
- [Real Server Groups](#)
- [REST Server](#)
- [Scripts](#)
- [SSH](#)
- [SSL](#)
- [SSL Profiles](#)
- [System](#)
- [Virtual IPs](#)
- [Virtual Servers](#)
- [Whiteboards](#)

Configuration Tables

1. [Overview](#)
2. [Working with Configuration Tables](#)
3. [Understanding the Use Default Option](#)
4. [Using Filters](#)
5. [Working with Scripts](#)

Overview

Configuration tables are components you can add to a dashboard and let you view and configure some F5[®] LineRate[®] objects.

Working with Configuration Tables


Configuration tables display some configured objects and let you create, edit, and delete those objects.



Note: You cannot configure all objects using F5[®] LineRate Manager. Any objects not listed across the top row of the configuration table must be configured using the CLI or REST API.

You can sort by clicking a column heading. The blue triangle shows the column currently used to sort and the sort order (ascending or descending). You can click a column heading and drag it to another location and resize columns.

Below is an example configuration table showing a few configured virtual IPs.



Name	Is Base	Base Name	IP Address	IP Address Range	Admin Status	SSL Profile	Status
<input type="checkbox"/> vip-01	<input type="checkbox"/>	vipb-01	201.0.45.1:20001		online	sp-01	UP
<input type="checkbox"/> vip-02	<input type="checkbox"/>	vipb-01	210.0.45.1:20002		online	sp-01	UP: no associated system inf
<input type="checkbox"/> vip-03	<input type="checkbox"/>	vipb-01	201.0.45.1:20003		online	sp-01	UP
<input type="checkbox"/> vip-04	<input type="checkbox"/>	vipb-01	201.0.45.1:20004		online	sp-01	UP

Item	Name	Description
1	Object types to configure	Click to see the objects of that type that are currently configured and to create or edit objects of that type. The columns are the parameters you can configure for the selected object.
2	Filter	Use to filter the list of configured objects. See Using Filters .
3	Refresh configuration	Click to refresh the table to reflect the current configuration.
4	Create new	Click to create a new object of the current type.
5	Delete	Click to delete the currently selected object.
6	Copy	Click to copy the currently selected object.
7	Yellow triangle	Denotes that the parameter is not set to the Use Default option, but is directly configured, that is, the parameter value is not inherited from a base. See Understanding the Use Default Option .



After creating an object in the configuration table, you cannot edit the object's name. You can copy the object, give the copy the name you want, then delete the original object.

Understanding the Use Default Option

When you configure some object parameters in the configuration table, you will see a window similar to this:

The Use Default option does one of two things depending on whether an object can have a [base](#):

- Objects that can have bases—The Use Default option removes a direct setting from the object. This option tells the object to inherit that setting from its base. If the object does not have a base or that setting is not configured in its base, the object returns to its default.
 - For example, you may want to create a real server base and set the admin status in the base to online. Then configure the real server to inherit its settings from the base by leaving the Use Default option selected.

When you want to take just one real server, called rs1 offline, set the Admin Status column for rs1 to Set Locally and select offline. This takes the real server offline, overriding the setting

from the base. When you want the real server to inherit its admin status from its base again, set its Admin Status column back to Use Default. This ensures that the admin status is not set locally on the real server, but that it inherits its setting from the base.

- Objects that cannot have bases—The Use Default option sets the object back to its default.

Using Filters

The filter box at the top of the configuration table lets you filter the list of configured objects, so you see a subset of one type of object.

Filters use the following format:

<Object_type> <Parameter> <operator> <value_to_filter>

For example, **vip name =~ web** displays only virtual IPs that contain "web" in their name.

Below is an example of a two-level filter.



Item	Name	Description
1	Filter field	Use to select a preconfigured filter or to create a filter, as shown and as described in the steps below.
2	N/A	Example of a two-level filter that first filters based on virtual IPs whose name include "web," then filters for any that have an admin status of offline. Click the x in each filter to remove that filter.
3	N/A	Shows that just one of the seven existing virtual IPs meets the filter criteria.

The filter uses the following operators:

Operator	Description
=	Equals—Exact match of value.
!=	Does not equal—Exact match of value to exclude.

Operator	Description
>	Greater than—Greater than value.
>=	Greater than or equal to—Greater than or equal to value.
<	Less than—Less than value.
<=	Less than or equal to—Less than or equal to value.
=~	Contains—Regular expression as value filter on.
!~	Does not contain—Regular expression as value to exclude.

The filter uses the following abbreviations for the object types:

Object type	Abbreviation
Virtual IP	vip
Forward proxy	fp
Virtual server	vs
Real server	rs
Real server group	rsg
Health monitor	hm
SSL profile	sp

If you enter multiple filters, the system ANDs the filters, that is, only objects that meet all of the filter criteria display. You can create up to ten filters.

Click the **x** next to a filter to remove it. Click the **x** in the main filter box to remove all filters.

You can enter the complete filter string, or you can let the filter "walk" you through the available options, as shown in the steps below, to create your own filter.



To use a pre-configured filter:


1. Click the magnifier icon .
2. Select a pre-configured filter to use.




To create your own filter:

1. Do one of the following:
 - In the filter box, type one or more letters that start the name of a parameter.

- For example, typing the letter **p** displays all of the parameters with a word that starts with p. Keep typing to narrow the list to include only the word you want (profile, port, private, etc.).



- In the filter box, type the abbreviation for an object type.
 - For example, typing **sp** displays all of the parameters for an SSL profile.



2. Select the parameter you want from the list and press **space bar**.
 - A list of available operators displays.
3. Select the operator you want from the list and press **space bar**.
4. Type or select the parameter value.
 - For parameters with limited options, such as admin status (online or offline), a list of options displays. For parameters with values like a name or IP address, type the value you want to filter on.
5. Press **Enter**.

Working with Scripts

The configuration table also lets you create, delete, copy, and edit inline scripts. Use the create, delete, and copy icons as you do for other objects in the configuration table, as described above.

After you create a script, you can edit it. Select the script and click  **Edit script**. An editor window opens where you can write the script.

The editor validates the JavaScript syntax. Be sure to correct lines that have the warning  icon. F5[®] LineRate[®] will not process scripts that contain these types of errors. The  icon indicates information that you can ignore. The editor cannot determine whether the script contains valid F5[®] LineRate[®] code.

In the lower-left corner, click  to open the F5[®] LineRate[®] Scripting API Reference Guide.

Dashboards

1. [Overview](#)
2. [Dashboard Components](#)

Overview

Dashboards let you configure and view information about your system.

Dashboard Components

Dashboards can include any combination of the following components:

- [Whiteboards](#)—Display a diagram of the selected object, for example, virtual servers.
- [Configuration tables](#)—Display most configured objects and let you create, edit, and delete objects. You cannot configure all objects using F5[®] LineRate Manager. Any objects not listed across the top row of the configuration table must be configured using the CLI or REST API.
- [Line charts](#)—Let you select statistics to monitor and display a line chart of those statistics.
- [Area charts](#)—Let you select statistics to monitor and display an area chart of those statistics.

Below is an example dashboard for a system configured with several virtual IPs, a virtual server, and a real server group that includes several real servers. The dashboard also includes two charts.



Item	Name	Description
1	Add or remove dashboards	Click to see a list of existing dashboards and to add or remove dashboards.

Item	Name	Description
2	Select dashboard	Click to select the dashboard to display.
3	Add a dashboard component	Click to select a component to add to the current dashboard.
4	Get system information	Click to access system user guides and system version information.
5	Configure system	Click to save to the startup config, to view interface and IP route configuration, and to configure SSL and purchased licenses. When the background turns orange, it means you have unsaved configuration changes. See Saving .
6	Line chart	Example line chart showing HTTP requests.
7	Line chart	Example line chart showing HTTP responses.
8	Configuration table	Example configuration table.
9	Whiteboard	Example whiteboard.

You can create as many dashboards as you need. For example, you may want one dashboard with a whiteboard and a configuration table that you use for view your configuration and for changing it. You may want another dashboard with four charts for monitoring requests, responses, client latency, and server latency.



To create a dashboard:

1. Click **Add or remove dashboards** .
2. Click **+** and give the new dashboard a name.
3. Click **Select dashboard**  and select the dashboard to configure.
4. Click **Add a dashboard component**  and select the component you want.
5. Configure and use the components as described in the sections that follow.

DNS

1. [Overview](#)
2. [Understanding DNS](#)
3. [Configuring DNS](#)
 1. [DNS Parameters](#)

Overview

DNS is one of the system configuration objects you can configure using F5[®] LineRate Manager.

Understanding DNS

Use to configure system-wide settings for Domain Name System (DNS), which allows the name-to-IP resolution of domain names.



Note: By default, the system comes with a ready-to-use DNS configuration. You can use the default servers or configure the system to use your own servers.

Configuring DNS

You can configure up to three DNS name servers and up to six domains to search to complete unqualified host names. When you click  **Configure System**, then click **DNS** in F5[®] LineRate Manager, you see any existing DNS configuration.

You can edit the Admin status, Retries, and Timeout fields.

To add a DNS Server or Search Domain, click the corresponding .

DNS Parameters

Each field in the DNS window is a parameter you can configure for DNS. The fields correspond very closely to the CLI commands used to configure DNS.

For the definition of each parameter you can configure using the CLI, see [ip dns](#).

Forward Proxies

1. [Overview](#)
2. [Understanding Forward Proxies](#)
3. [Configuring Forward Proxies](#)
 1. [Forward Proxy Parameters](#)

Overview

Forward proxies are one of the types of objects you can configure using F5[®] LineRate Manager.

Understanding Forward Proxies

The F5[®] LineRate[®] forward proxy capability provides a proxy function from one network to another. A common use case for a forward proxy is for connections from your private network to the Internet.

A forward proxy lets you insert custom logic created with scripts. Scripts can perform a variety of functions, including gathering usage statistics, redirecting requests to your own cache, blocking of access to specific sites, managing cookies, and much more.

To use a forward proxy effectively, be sure to create and attach a virtual IP that includes the range of Internet IP addresses you want to go through the forward proxy.



Caution: When attaching a virtual IP to a forward proxy, the virtual IP must not include any of the system's own IP addresses. For a virtual IP with a single IP address, do not set the virtual IP's IP address to one of the system's own IP addresses. For a virtual IP with a range of addresses, you must ensure that the IP address range does not contain any of the system's own IP addresses. This may mean you need to break the virtual IP into multiple virtual IPs. See [Configuring a range for a virtual IP with forward proxy](#) for more detail and an example.

For additional information and an example configuration, see [Configuring a Forward Proxy](#).

Configuring Forward Proxies

When you click **Forward Proxies** in the F5[®] LineRate Manager configuration table, you see any existing forward proxies. For information about using the configuration table, see [Configuration Tables](#).



Note: After creating an object in the configuration table, you cannot edit the object's name. You can copy the object, give the copy the name you want, then delete the original object.

Forward Proxy Parameters

Each column in the configuration table is a parameter you can configure for a forward proxy. The column names correspond very closely to the CLI commands used to configure forward proxies. For the definition of each parameter, see [Forward Proxy Mode Commands](#).

Health Monitors

1. [Overview](#)
2. [Understanding Health Monitors](#)
3. [Configuring Health Monitors](#)
 1. [Health Monitor Parameters](#)

Overview

Health monitors are one of the types of objects you can configure using F5[®] LineRate Manager.

Understanding Health Monitors

Use to create a health monitor to regularly check that servers are up and able to accept connections or respond to requests. You can attach a health monitor to a real server and to a real server base.

For HTTP health monitors, you can configure the type of requests and responses.

Configuring Health Monitors

When you click **Health Monitors** in the F5[®] LineRate Manager configuration table, you see any existing health monitors. For information about using the configuration table, see [Configuration Tables](#).



Note: After creating an object in the configuration table, you cannot edit the object's name. You can copy the object, give the copy the name you want, then delete the original object.

Health Monitor Parameters

Each column in the configuration table is a parameter you can configure for a health monitor. The column names correspond very closely to the CLI commands used to configure health monitors. For the definition of each parameter, see [Health Monitor Mode Commands](#).

Interfaces

1. [Overview](#)
2. [Understanding Interfaces](#)
3. [Configuring Interfaces](#)
 1. [Interface Parameters](#)

Overview

Interfaces are one of the system configuration objects you can configure using F5[®] LineRate Manager.

Understanding Interfaces

Configure settings for a network interface for management or data use.

Use

Typically, you want to configure one interface for management use and one or more other interfaces for data use.

To see the names of your interfaces, use the following:

- CLI command: **show interfaces**
 - The following information displays:
 - A list of all interfaces on the system displays. It is possible for the system to have an interface that F5[®] LineRate[®] cannot detect.
 - The first line is the interface name and its status. The interface name is based on the driver for the interface type. The remaining lines list the information available about the interface (typically, manufacturer, model, MAC address, speed in kilobits, and more).
 - Below are the names used for some common interfaces:
 - em—Intel 1Gb interface
 - igb—Intel 1Gb interface
 - bce—Broadcom 1Gb interface
 - ix—Intel 10 Gb interface
 - oce—Emulex 10 Gb interface
 - xn—Xen netfront interface (used by Amazon EC2)

- lo—Loopback interface (internal interface)
- po—Port channel interface
- REST API node: /status/system/interface?op=list
 - The names and current settings for every interface display. You can find all of the information that is included in the show interfaces command in the hierarchy below /status/system/interface.

You can create up to 4094 subinterfaces on an interface. Subinterfaces are disabled by default when you create them. Use the [encapsulation](#) command to set up trunked ports for VLANs.

Configuring Interfaces

When you click  **Configure system**, then click **Interfaces** in F5[®] LineRate Manager, you see all existing network interfaces.

Interface Parameters

Each column in the Network Interface table is a parameter you can configure for an interface. The column names correspond very closely to the CLI commands used to configure interfaces.

The interface parameters available in F5[®] LineRate Manager are only a subset of the configurable parameters.

You can edit the following for interfaces in F5[®] LineRate Manager:

- Status
- MTU
- IPv6 link-local
- IP address

For the definition of each parameter and for all parameters you can configure using the CLI, see [Interface Mode Commands](#).

IP Routes

1. [Overview](#)
2. [Understanding IP Routes](#)
3. [Configuring IP Routes](#)
 1. [IP Route Parameters](#)

Overview

IP routes are one of the system configuration objects you can configure using F5[®] LineRate Manager.

Understanding IP Routes

Configure global IP route settings.

Use

Use to configure the IP routes for the system. You can set routes to go through a specific system interface or through another system, likely a router, on your network. Be sure to configure routes to include every subnet the system needs.

A manually configured default IP route will override any default route obtained from DHCP.

Configuring IP Routes

You can configure IPv4 and IPv6 routes. When you click  **Configure System**, then click **IP Routes** in F5[®] LineRate Manager, you see any existing IP routes.

To add a route, click .

To define the default route, enter one of the following in the Destination field:

- Default-IPV4
- Default-IPV6

You can define both defaults, if needed.



Note: Be sure that the Destination and Gateway address types match, that is, they are both either IPv4 or IPv6.

To use the default MTU, leave the MTU field blank.

IP Route Parameters

Each column in the Static Routes table is a parameter you can configure for an IP route. The column names correspond very closely to the CLI commands used to configure IP routes.

For the definition of each parameter you can configure using the CLI, see [ip routes](#) and [ipv6 routes](#).

Licenses

1. [Overview](#)
2. [Understanding Licenses](#)
3. [Configuring Licenses](#)
 1. [License Parameters](#)

Overview

Licenses are one of the system configuration objects you can configure using F5[®] LineRate Manager.

Understanding Licenses

You can use F5[®] LineRate Manager to configure F5[®] LineRate[®], but if you want to run traffic through the system, you must have a license.

Configuring Licenses

When you click  **Configure System**, then click **Licenses** in F5[®] LineRate Manager, you see any existing licenses.

You can activate a license here. For the complete process, see [Configuring Licensing](#).

License Parameters

The license configuration in F5[®] LineRate Manager corresponds very closely to the CLI commands used to configure licenses.

For more information about licensing, see:

- [Licensing Mode Commands \(config\)](#)
- [License Mode Commands \(exec\)](#)
- [About Licensing](#)

Line and Area Charts

1. [Overview](#)
2. [Configuring Line Charts](#)
3. [Configuring Area Charts](#)
4. [Available Statistics](#)

Overview


Line charts and area charts are components you can add to a dashboard and display statistics for monitoring. The charts poll the system for statistic updates once per second.

For more information about monitoring the system, see [Monitoring LineRate](#).

Configuring Line Charts

Line charts let you select statistics to monitor and display a line chart of those statistics. Use line charts to compare multiple values directly.

Below is an example line chart that shows the total number of client requests for two virtual IP addresses (vip-01 and vip-02).



For a description of the available statistics, see [Available Statistics](#).



To configure a line chart:

1. From the dashboard where you want to add the line chart, click **Add a dashboard component**  and select **Line Chart**.
2. In the upper-left corner of the new chart, click **Customize chart** .
3. Select statistics to include on the chart.
 - The chart starts collecting data as soon as you add a statistic.


Drop-down list/ Field	Description
Object Type	Select the type of object whose statistics you want to chart.
Objects	Select a specific object or All. The All option includes aggregate statistics of all configured objects of the selected type.
Statistics	From the drop-down list, select a statistics category. From the list of statistics below the category, select a statistic and click Add statistic . Select up to five statistics with identical units. The system automatically assigns colors to each statistic line.

4. In the upper-left corner of the chart, click **Customize chart**  to return to the chart view.
5. Hold your mouse pointer over the chart to see the chart legend.

Configuring Area Charts

Area charts let you select statistics to monitor and display an area chart of those statistics. Use area charts to see a cumulative total and the relative component values of the total. Read the cumulative total from the Y-axis and individual values from the legend.

The example stacked area chart below is looking at server requests handled by four real servers and load balanced by one virtual server. You can read total requests from the Y-axis (200 total requests at the point of the mouse hover) and compare the number of requests handled by each real server (50 for each). The number of requests are the same (all 50), because the virtual server is using round robin load balancing algorithm.



For a description of the available statistics, see [Available Statistics](#).



To configure an area chart:

1. From the dashboard where you want to add the area chart, click **Add a dashboard component**  and select **Area Chart**.
2. In the upper-left corner of the new chart, click **Customize chart** .
3. Select statistics to include on the chart.

Drop-down list/ Field	Description
Object Type	Select the type of object whose statistics you want to chart.
Objects	Select a specific object or All. The All option includes aggregate statistics of all configured objects of the selected type.
Statistics	From the drop-down list, select a statistics category. From the list of statistics below the category, select a statistic and click Add statistic . Select up to five statistics with identical units. The system automatically assigns colors to each statistic area.

- The chart starts collecting data as soon as you add a statistic.

4. In the upper-left corner of the chart, click **Customize chart**  to return to the chart view.
5. Hold your mouse pointer over the chart to see the chart legend.

Available Statistics

The statistics available vary based on the selected object type. The table below describes the statistics available for both line and area charts.

Statistic category	Statistic	Description	Units	Available for
Client Latency				
	Client Transaction Latency	1 minute average of the average time between receiving the first byte of the HTTP request from a client and sending the last byte of the response for that request to the client.	Seconds	Forward proxy, virtual IP, virtual server
Connect Tunnels				
	Open Tunnels	Number of active HTTP connect tunnels.	Tunnels	Forward proxy, real server, virtual server
	Opened Tunnels Rate	Number of HTTP connect tunnels opened per second.	Tunnels/second	Forward proxy, real server, virtual server
Connection Errors				
	Client Address in Use	Number of client connections refused due to address-in-use error.	Connections	Virtual IP
	Client Connections Filtered	Number of client connections refused because of an IP filter rule.	Connections	Virtual IP
	Client Connections Lost	Number of client connections closed due to death of a worker process.	Connections	Virtual IP
	Client Connections Refused	Number of client connections refused due to an input/output error.	Connections	Virtual IP
	Client Connections Reset	Number of client connections reset by the load balancer.	Connections	Virtual IP
	Client Connections Timed Out	Number of client connections closed due to timeout.	Connections	Virtual IP

Statistic category	Statistic	Description	Units	Available for
	Client Idle Timeout	Client connections closed due to client idle timeout.	Connections	Virtual IP
	Client No Virtual Server	Number of client HTTP requests in the HTTP load balancer mode and the number of client connections in the L4 load-balancer mode errored due to unavailable virtual server.	Requests	Virtual IP
	Client Other Errors	Number of other errors seen on client connections.	Errors	Virtual IP
	Client SSL Negotiation Failed	Number of failed SSL client connection attempts.	Connections	Virtual IP
	Max Client Connections	Number of times a client connection could not be accepted due to reaching maximum connections limit.	Errors	Virtual IP
	Max Embryonic Connections Dropped	Number of embryonic connections dropped because the VIP's maximum embryonic connections has been exceeded. This can occur if the limit is set too low, or if a synflood DOS attack is in progress.	Connections	Virtual IP
	SSL Profile Error	Number of client connections reset because of an unavailable SSL profile.	Connections	Virtual IP
	Other Server Errors	Number of other errors seen on server connections.	Errors	Real server
	Server Address in Use	Number of connections to the servers refused due to address in use.	Connections	Real server
	Server Connections Lost	Number of server connections closed due to death of a worker process.	Connections	Real server
	Server Connections Refused	Number of connections refused by the server.	Connections	Real server
	Server Connections Reset	Number of connections reset by the server.	Connections	Real server
	Server Idle Timeout	Number of server connections closed due to idle timeout.	Connections	Real server
	Server Timed Out	Number of server connections closed due to timeout.	Connections	Real server

Statistic category	Statistic	Description	Units	Available for
Connections				
	Open Connections	Number of client connections opened (accepted).	Connections	Virtual IP
	Open Connections Rate	Number of client connections opened (accepted) per second.	Connections/second	Virtual IP
	Server Open Connections	Number of server connections opened.	Connections	Real server
	Server Opened Rate	Number of server connections opened per second.	Connections/second	Real server
HTTP Requests				
	Client LB Forbidden Requests	Number of client HTTP requests that were forbidden.	Requests	Virtual IP
	Client Request Rate	Number of client HTTP requests per second received.	Requests/second	Virtual IP
	Good Client Requests	Number of client HTTP requests that were good.	Requests	Virtual IP
	Total Client Request Errors	Number of client HTTP requests that had errors.	Requests	Virtual IP
	Total Client Requests	Number of client HTTP requests received.	Requests	Virtual IP
	Server HTTP Request Rate	Number of HTTP requests per second received from clients.	Requests/second	Real server
	Server Queue Size	Total number of HTTP requests in the queue at the real server. This is the number of requests for which an entire response has not yet been received from the server.	Requests	Real server
	Total HTTP Server Requests	Number of HTTP requests received from clients.	Requests	Real server
HTTP requests (Initiation)				
	Proxy HTTP Request Rate	Number of HTTP requests per second sent to the servers.	Requests/second	Forward proxy,

Statistic category	Statistic	Description	Units	Available for
				virtual server
	Total Proxy HTTP Requests	Number of HTTP requests sent to the servers.	Requests	Forward proxy, virtual server
	Total Proxy Queue Size	Total number of HTTP client requests in the HTTP load balancer mode and the total number of client connections in the L4 load balancer mode that are queued. This is the number of HTTP requests for which an entire response has not yet been received from the server and the number HTTP requests and the client connections waiting to be load balanced to a real server.	Requests	Forward proxy, virtual server
HTTP Requests (Termination)				
	Client Request Rate	Number of client HTTP requests per second received.	Requests/second	Forward proxy, virtual server
	Service Unavailable Error Rate	Number of requests per second responded to with 503 due to rate limit exceeded.	Errors/second	Forward proxy, virtual server
	Service Unavailable Errors	Number of requests responded to with 503 due to rate limit exceeded.	Requests	Forward proxy, virtual server
	Total Client Errors	Number of client HTTP requests that had errors.	Requests	Forward proxy, virtual server
	Total Client Requests	Number of client HTTP requests received.	Requests	Forward proxy, virtual server
HTTP Responses				

Statistic category	Statistic	Description	Units	Available for
	Bad Server Responses	Number of server responses where an error occurred.	Responses	Real server
	HTTP Response Rate	Number of HTTP responses per second received from servers.	Responses/second	Real server
	Client Completed Responses	Number of HTTP responses sent to the clients.	Responses	Virtual IP
	Client Idle Timed Out Responses	Number of times an HTTP response could not be sent to the client due to a delay in sending any part of the response from the server to the client for more than the response-idle-timeout period configured for a real server. The delay may be due to the server's inability to send the response or the client's inability to receive the response. The time period starts immediately after receiving the response header from the server.	Unsent responses	Virtual IP
	Client LB 5xx Error Responses	Number of internally generated HTTP 5xx responses.	Responses	Virtual IP
	Client Started Responses	Number of HTTP responses attempted to be sent to the clients.	Responses	Virtual IP
	Client Timed Out Responses	Number of times an HTTP response could not be sent to a client due to not receiving the entire response header from the server within the response-timeout period configured for a real server. The time period begins either when the entire request was sent to the server or when the responses for all prior requests on the same server connection have been received, whichever is more recent.	Unsent responses	Virtual IP
HTTP Responses (Initiation)				
	Bad HTTP Server Responses	Number of server responses that contained errors.	Responses	Forward proxy, virtual server
	Request redirects received via scripting	Number of request redirects received via scripting.	Redirects	Forward proxy, virtual server

Statistic category	Statistic	Description	Units	Available for
	Request redirects sent via scripting	Number of request redirects sent via scripting.	Redirects	Forward proxy, virtual server
	Total Proxy HTTP Responses	Number of HTTP responses received from the servers.	Responses	Forward proxy, virtual server
HTTP Responses (Termination)				
	Client Idle Timeout	Number of times an HTTP response could not be sent to the client due to a delay in sending any part of the response from the server to the client for more than the response-idle-timeout period configured for a real server. The delay may be due to the server's inability to send the response or the client's inability to receive the response. The time period starts immediately after receiving the response header from the server.	Unsent responses	Forward proxy, virtual server
	Client LB 5xx Error	Number of internally generated HTTP 5xx responses.	Errors	Forward proxy, virtual server
	Client LB Forbidden	Number of client HTTP requests that were forbidden.	Requests	Forward proxy, virtual server
	Client Responses Timed Out	Number of times an HTTP response could not be sent to a client due to not receiving the entire response header from the server within the response-timeout period configured for a real server. The time period begins either when the entire request was sent to the server or when the responses for all prior requests on the same server connection have been received, whichever is more recent.	Unsent responses	Forward proxy, virtual server
	Script Gateway Timeout Error	Number of script timeouts causing 504 error responses.	Errors	Forward proxy,

Statistic category	Statistic	Description	Units	Available for
				virtual server
	Script Gateway Timeout Error Rate	Number of script timeouts per second causing 504 error responses.	Errors/second	Forward proxy, virtual server
	Service Unavailable Error Rate	Number of requests per second responded to with 503 due to rate limit exceeded.	Errors/second	Forward proxy, virtual server
	Service Unavailable Error	Number of requests responded to with 503 due to rate limit exceeded.	Requests	Forward proxy, virtual server
	Total Client Responses	Number of HTTP responses sent to the clients.	Responses	Forward proxy, virtual server
Server Latency				
	Server Transaction Latency	Average latency for a single HTTP request/response transaction between the load balancer and the server.	Seconds	Forward proxy, real server, virtual server
Traffic				
	Client L7 SSL Traffic Received	SSL encrypted layer 7 bytes per second received from clients.	Bits/second	Virtual IP
	Client L7 SSL Traffic Sent	SSL encrypted layer 7 bytes per second sent to clients.	Bits/second	Virtual IP
	Server L7 Traffic Received	Layer 7 bytes per second received from servers.	Bits/second	Real server, virtual IP
	Server L7 Traffic Sent	Layer 7 bytes per second sent to servers.	Bits/second	Real server, virtual IP

NTP

1. [Overview](#)
2. [Understanding NTP](#)
3. [Configuring NTP](#)

Overview

Network time protocol (NTP) is one of the system configuration objects you can configure using F5[®] LineRate Manager.

Understanding NTP

Use

Use to set an IP address of a network time protocol server (NTP) to use to control the system time. You can set up more than one NTP server.

After configuring an NTP server for the first time, we recommend using the [ntp_sync](#) command or `/exec/ntp/sync` node to force an NTP update. Forcing the update ensures that there are no drift issues between the system time and NTP server time. so that NTP can make proper time updates.

Configuring NTP

When you click  **Configure system**, then click **NTP** in F5[®] LineRate Manager, you see any existing NTP configuration.

To configure NTP, click  and enter the IP address of an NTP server.

To sync the time immediately, enter the IP address of an NTP server and click **Synchronize Now**.

The NTP functions work the same as the corresponding CLI commands. For more information, see [NTP Mode Commands \(config\)](#) and [NTP Mode Commands \(exec\)](#).



Phone Home

1. [Overview](#)
2. [Understanding Phone Home](#)
3. [Configuring Phone Home](#)

Overview

Phone home is one of the system configuration objects you can configure using F5[®] LineRate Manager.

Understanding Phone Home

The phone home feature allows F5[®] to maintain and improve F5[®] LineRate[®] by gathering usage data including, but not limited to, URL addresses, usage, and configuration information. It is required to activate some license types. <https://linerate.ihealth.f5.com/>

To enable phone home, you will need your F5 username and password, which you created when you downloaded the installation file or obtained a license. Also, the system must be able to connect to the Internet so that it can phone home.

---->DRAFT: The phone home feature allows us to maintain and improve F5[®] LineRate[®] via the reporting of usage data. These data include, but are not limited to, URLs, usage statistics, and configuration information (additional details below). Phone home service is enabled upon activation of certain license types. As successful phone home uploads are required to activate these licenses, you must [configure your F5[®] username and password](#) prior to, or at the time of, license activation.

Phone home sends the following data to F5 once a week:

Data Sent	Data collection frequency	Purpose
Core files	As they occur	To proactively diagnose errors.
Output of <code>show tech-support detailed</code>	Hourly	To provide information about configuration and usage.

Data Sent	Data collection frequency	Purpose
System logs (all files in /var/log/)	Hourly	To proactively diagnose errors.
Script events (create, remove, online, offline, run-time errors, inline script code)	As they occur, written to disk hourly (or when the amount collected exceeds a threshold)	To proactively diagnose errors.
Per-script statistics	Every 30 minutes	To see how much scripting is being used.
On-disk scripts and dependencies (contents of /home/linerate/data/scripting/)	Every 4 hours	To proactively diagnose errors.
username.txt file (contains configured phone home username)	With each upload	To identify the configured phone home user.



Note: The phone home feature makes an effort to not include sensitive data like passwords, SSL private keys, and SNMP community strings from the system configuration in the data sent to F5[®]. However, sensitive data in script configuration, on-disk scripts, system logs, and core files cannot be redacted.

Configuring Phone Home

When you click  **Configure system**, then click **Phone Home** in F5[®] LineRate Manager, you see the phone home status and any configured username and password.

The configuration is the same as when using the CLI. For more information, see [Phone Home Mode Commands](#).

Real Servers

1. [Overview](#)
2. [Understanding Real Servers](#)
3. [Configuring Real Servers](#)
 1. [Real Server Parameters](#)
 2. [Understanding the Is Base Column](#)

Overview

Real servers are one of the types of objects you can configure using F5[®] LineRate Manager.

Understanding Real Servers

Real servers represent a service, for example a web server, that the load balancer (reverse proxy) is distributing the client requests to. Each load balancer requires at least one real server. The IP address and port for the real server must match the IP address and port of the service on the server or proxy server the real server talks to.

You can create multiple real servers, for example, for different application types, hardware capabilities (such as CPU or memory), or security settings (such as SSL).

We also recommend creating one or more real server bases to make configuring real servers more consistent. See [base](#).

Default Setting

By default, no real servers exist.

When you create a real server, the default settings are:

- admin-status—offline
- attach—nothing attached
- base—none
- description—none
- ip address—none
- service—service is set to http

For additional information and an example configuration, see [Configuring a Reverse Proxy](#).

Configuring Real Servers

When you click **Real Servers** in the F5[®] LineRate Manager configuration table, you see any existing real servers. For information about using the configuration table, see [Configuration Tables](#).



Note: After creating an object in the configuration table, you cannot edit the object's name or its Is Base setting. You can copy the object, give the copy the name and Is Base setting you want, then delete the original object.

Real Server Parameters

Each column in the configuration table is a parameter you can configure for a real server. The column names correspond very closely to the CLI commands used to configure real servers. For the definition of each parameter, see [Real Server Mode Commands](#).

Understanding the Is Base Column

Some objects in F5[®] LineRate[®] can have bases. For more information about bases, see [Working with Bases](#).

In F5[®] LineRate Manager, if the Is Base column has a checkmark in it, the object is a base.

Real Server Groups

1. [Overview](#)
2. [Understanding Real Server Groups](#)
3. [Configuring Real Server Groups](#)
 1. [Real Server Group Parameters](#)

Overview

Real server groups are one of the types of objects you can configure using F5[®] LineRate Manager.

Understanding Real Server Groups

Real server groups are an efficient way to configure real servers. You can attach a real server group to a virtual server. Be sure to group real servers based on those served by the same virtual IP and virtual server.

Configuring Real Server Groups

When you click **Real Server Groups** in the F5[®] LineRate Manager configuration table, you see any existing real server groups. For information about using the configuration table, see [Configuration Tables](#).



Note: After creating an object in the configuration table, you cannot edit the object's name. You can copy the object, give the copy the name you want, then delete the original object.

Real Server Group Parameters

Each column in the configuration table is a parameter you can configure for a real server group. The column names correspond very closely to the CLI commands used to configure real server groups. For the definition of each parameter, see [real server groups](#).

REST Server

1. [Overview](#)
2. [Understanding the REST Server](#)
3. [Configuring the REST Server](#)
 1. [REST Server Parameters](#)

Overview

The REST server is one of the system configuration objects you can configure using F5[®] LineRate Manager.

Understanding the REST Server

Use to configure the HTTP server on the F5[®] LineRate[®] system for Representational State Transfer (REST) access. Log in to the REST server using the same login and password that you use for the F5[®] LineRate[®] system. By default, connections to the REST server must use SSL port 8443.

F5[®] LineRate Manager, the GUI for F5[®] LineRate[®], also uses the REST server for access.

The REST server uses the following HTTP verbs: GET, PUT, POST, and DELETE and lets you do the following:

- Configure the system (add, change, or delete configuration)
- Retrieve system configuration
- Retrieve various statistics and counters used to monitor the system

By default, the system configuration permits access to the REST server on any local interface (on port 8443) from any remote host. To connect to the REST server, point your client (custom REST client application or browser) to the REST server's IP address and port to establish a secure HTTP connection.

Configuring the REST Server

When you click  **Configure System**, then click **REST Server** in F5[®] LineRate Manager, you see any existing REST server configuration.



Caution: If you change the SSL Profile, F5[®] LineRate Manager will automatically restart using the new SSL profile, and you will need to log in again. Be sure that the SSL profile you select is valid and active on F5[®] LineRate[®], or you will not be able to connect to F5[®] LineRate Manager or to the REST server.

REST Server Parameters

The configuration corresponds very closely to the CLI commands used to configure the REST server.

For the definition of each parameter you can configure using the CLI, see [REST Server Mode Commands](#).

For information about using REST, see the [REST API Reference Guide](#).

Scripts

1. [Overview](#)
2. [Understanding Scripts](#)
3. [Configuring Scripts](#)
 1. [Script Parameters](#)

Overview

Inline scripts are one of the types of objects you can configure using F5[®] LineRate Manager. You cannot create a separate script file for the script source.

Understanding Scripts

Use to create or change scripts and to configure script settings. For more information about scripts, see the [Scripting Developer's Guide](#) and [Scripting API Reference Guide](#).



Note: If you enable a Point[™] Load Balancer license, you can configure and edit scripts, but the scripts will not run. If you later upgrade to a Precision[™] Load Balancer license, all configured scripts with valid source and admin status set to online will start.

Configuring Scripts

When you click **Scripts** in the F5[®] LineRate Manager configuration table, you see any existing scripts. For information about using the configuration table, see [Configuration Tables](#).



Note: After creating an object in the configuration table, you cannot edit the object's name. You can copy the object, give the copy the name you want, then delete the original object.

For more about the script editor, see [Working with Scripts](#).

Script Parameters

Each column in the configuration table is a parameter you can configure for a script. The column names correspond very closely to the CLI commands used to configure scripts. For the definition of each parameter, see [Script Mode Commands](#).

SSH

1. [Overview](#)
2. [Understanding SSH](#)
3. [Configuring SSH](#)
4. [SSH Parameters](#)

Overview

SSH is one of the system configuration objects you can configure using F5[®] LineRate Manager.

Understanding SSH

Configure options for Secure Shell (SSH) access to the F5[®] LineRate[®] system.

Use

SSH provides secure, remote access to the system.

Configuring SSH

When you click  **Configure system**, then click **SSH** in F5[®] LineRate Manager, you see any existing SSH configuration.

Use the Allow From and Allow To tabs to configure SSH.

To add a setting, click .

SSH Parameters

The configuration corresponds very closely to the CLI commands used to configure SSH.

For the definition of each parameter you can configure using the CLI, see [SSH Mode Commands](#).

SSL

1. [Overview](#)
2. [Understanding SSL](#)
3. [Configuring SSL](#)
4. [SSL Parameters](#)

Overview

Certificates, certificate bundles, and private keys for SSL are some of the system configuration objects you can configure using F5[®] LineRate Manager.

Understanding SSL

Secure Sockets Layer (SSL) and Transport Layer Security (TLS) are closely related technologies that provide communication security over an insecure network, such as the Internet. TLS is a standardized protocol, defined by IETF RFCs, and is the successor to the non-standardized SSL protocol. The F5[®] LineRate[®] software supports both TLS and SSL, for both service type TCP and service type HTTP, but the system and documentation refers to both protocols collectively as "SSL," following the most common industry terminology.

Configuring SSL

When you click  **Configure system**, then click **SSL** in F5[®] LineRate Manager, you see any existing SSL configuration for certificates, certificate bundles, and private keys. You can configure new certificates, certificate bundles, and private keys, and import the associated files. You can also replace an existing certificate, certificate bundle, or private key with another file.

SSL Parameters

The configuration corresponds very closely to the CLI commands used to configure certificates, certificate bundles, and private keys.

For the definition of each parameter you can configure using the CLI, see [Certificate Mode Commands](#) and [Key Mode Commands](#).

For an example SSL configuration, see [Configuring SSL](#).

SSL Profiles

1. [Overview](#)
2. [Understanding SSL Profiles](#)
3. [Configuring SSL Profiles](#)
 1. [SSL Profile Parameters](#)

Overview

SSL profiles are one of the types of objects you can configure using F5[®] LineRate Manager.

Understanding SSL Profiles

Use to define the security settings you want to use for SSL access. You can use an SSL profile for either termination SSL or initiation SSL or both.

Attaching an SSL profile to a virtual IP configures that virtual IP to always use SSL. The virtual IP will no longer accept connections from clients unless they perform SSL negotiation.

For more information, see [Managing SSL](#).

Secure Sockets Layer (SSL) and Transport Layer Security (TLS) are closely related technologies that provide communication security over an insecure network, such as the Internet. TLS is a standardized protocol, defined by IETF RFCs, and is the successor to the non-standardized SSL protocol. The F5[®] LineRate[®] software supports both TLS and SSL, for both service type TCP and service type HTTP, but the system and documentation refers to both protocols collectively as "SSL," following the most common industry terminology.

For an example SSL configuration, see [Configuring SSL](#).

Configuring SSL Profiles

When you click **SSL Profiles** in the F5[®] LineRate Manager configuration table, you see any existing SSL profiles. For information about using the configuration table, see [Configuration Tables](#).



Note: After creating an object in the configuration table, you cannot edit the object's name or its Is Base setting. You can copy the object, give the copy the name and Is Base setting you want, then delete the original object.

SSL Profile Parameters

Each column in the configuration table is a parameter you can configure for an SSL profile. The column names correspond very closely to the CLI commands used to configure SSL profiles. For the definition of each parameter, see [SSL Mode Commands](#).

System

1. [Overview](#)
2. [Saving](#)
3. [Host Name](#)
4. [System Root Certificate Bundle](#)
5. [Shut Down and Restart](#)

Overview

The System window lets you save your configuration, set the system host name, and shut down or restart the system.

When you click  **Configure system** in F5[®] LineRate Manager, the System tab displays.

Saving

The **Save to Startup** button works just like the CLI `write` command. For more information see [Running Config and Startup Config](#) and [Write Command](#).

When you make configuration changes that you have not yet saved, the background of the **Configure system** icon changes from blue  to orange .

Host Name

To change the host name, click the Host Name field and enter the name.

System Root Certificate Bundle

To change the system root certificate bundle, click the Root Certificate Bundle field and enter the name of a certificate bundle. For information about configuring certificate bundles, see [SSL](#). For additional information certificates in F5[®] LineRate[®], see [Certificate Mode Commands](#).

To return to using the default system root certificate bundle, leave the field blank and press **Enter**.

Shut Down and Restart

The shut down and restart functions work just like the corresponding CLI commands. For more information, see [Halt Mode Commands](#) and [Reload Mode Commands](#).

Users

1. [Overview](#)
2. [Understanding Users](#)
3. [Configuring Users](#)
 1. [User Parameters](#)

Overview

You can configure F5[®] LineRate[®] users using F5[®] LineRate Manager.

Understanding Users

Configure local user authentication

Use

Use to configure users who need access to the F5[®] LineRate[®] system. All users have access to all functions. The system tracks user logins and logouts and a few key actions, such as reload.

Configuring Users

You can configure as many F5[®] LineRate[®] users as you need. When you click  **Configure System**, then click **Users** in F5[®] LineRate Manager, you see any existing users.

To add a user, click . If you need to add more than one user, select the Create Another option to keep the fields available after you click **Add**.

User Parameters

Each field in the User window is a parameter you can configure. The fields correspond very closely to the CLI commands used to configure users.

For the definition of each parameter you can configure using the CLI, see [User Name Mode Commands](#).

Virtual IPs

1. [Overview](#)
2. [Understanding Virtual IPs](#)
3. [Configuring Virtual IPs](#)
 1. [Virtual IP Parameters](#)
 2. [Understanding the Is Base Column](#)

Overview

Virtual IPs are one of the types of objects you can configure using F5[®] LineRate Manager.

Understanding Virtual IPs

Create or modify a virtual IP for reverse proxy (load balancing) or forward proxy.

Use

For either a load balancing or forward proxy use case, the system requires at least one virtual IP. The virtual IP is a configuration object that represents the interface that clients connect to. You can create as many virtual IPs as you need. For an overview of how virtual IPs are used in a load balancing use case, see [F5[®] LineRate[®] Overview](#).

Configuring Virtual IPs

When you click **Virtual IPs** in the F5[®] LineRate Manager configuration table, you see any existing virtual IPs. For information about using the configuration table, see [Configuration Tables](#).



Note: After creating an object in the configuration table, you cannot edit the object's name or its Is Base setting. You can copy the object, give the copy the name and Is Base setting you want, then delete the original object. The Status column is a read-only value that F5[®] LineRate[®] reports.

Virtual IP Parameters

Each column in the configuration table is a parameter you can configure for a virtual IP. The column names correspond very closely to the CLI commands used to configure virtual IPs. For the definition of each parameter, see [Virtual IP Mode Commands](#).

Understanding the Is Base Column

Some objects in F5[®] LineRate[®] can have bases. For more information about bases, see [Working with Bases](#).

In F5[®] LineRate Manager, if the Is Base column has a checkmark in it, the object is a base.

Virtual Servers

1. [Overview](#)
2. [Understanding Virtual Servers](#)
3. [Configuring Virtual Servers](#)
 1. [Virtual Server Parameters](#)

Overview

Virtual servers are one of the types of objects you can configure using F5[®] LineRate Manager.

Understanding Virtual Servers

Each load balancing (reverse proxy) configuration requires at least one virtual server. The virtual server is a configuration object that acts as a reverse proxy and ties together one or more virtual IPs and real servers. You also set the load balancing algorithm on the virtual server.

For additional information and an example configuration, see [Configuring a Reverse Proxy](#).

Configuring Virtual Servers

When you click **Virtual Servers** in the F5[®] LineRate Manager configuration table, you see any existing virtual servers. For information about using the configuration table, see [Configuration Tables](#).



Note: After creating an object in the configuration table, you cannot edit the object's name. You can copy the object, give the copy the name you want, then delete the original object.

Virtual Server Parameters

Each column in the configuration table is a parameter you can configure for a virtual server. The column names correspond very closely to the CLI commands used to configure virtual server. For the definition of each parameter, see [Virtual Server Mode Commands](#).

Whiteboards


1. [Overview](#)
2. [Understanding Whiteboards](#)

Overview

Whiteboards are components you can add to a dashboard and show a diagram of the object selected at the top of the whiteboard.

Understanding Whiteboards

Below is an example whiteboard showing the configuration diagram for a virtual server called vs-01.



Item	Name	Description
1	List of object types	Click to see a list of configured objects of the selected type.
2	Object list	Click to see the configuration of the selected object. Scroll down the list to see all objects of the selected type.
3	Diagram	Diagram showing how the selected object is configured. See the table below for the definition of each icon type.
4	Refresh configuration	Click to refresh the diagram to reflect the current configuration.
5	Adjust size	Slide to change the size of the text in the diagram.

The icons in the diagram show the status of each object as follows:

Icon	Color	Description
	Green	Normal state.
	Yellow	Degraded—An attached object is offline or has bad status. Only applies to virtual server, forward proxy, and real server group.
	Gray	Admin offline—The object's admin status is offline.
	Red	Down or bad status—The object is either missing a key configuration, such as an IP address, or the object is down. For virtual servers and forward proxies—Shows that no objects are attached.