TIBCO LogLogic® Log Management Intelligence

TIBCO LogLogic® Enterprise Virtual Appliance

TIBCO LogLogic® Enterprise Virtual Appliance - Container Edition

Administration

Version 6.4.0
February 2022

Contents

Overview

Log data can comprise up to 25 percent of all enterprise data. Log data also contains critical information that can improve security, compliance and availability. Until now most companies have relied on ineffective and inefficient homegrown solutions and manual processes to manage this data.

TIBCO LogLogic® provides the industry's first enterprise class, end-to-end log management solution. Using LogLogic® log management solutions, IT organizations can analyze and archive network log data for the purpose of compliance and legal protection, decision support for network security remediation, increased network performance, and improved availability.

LogLogic log management appliances help you to simplify, automate, and reduce the cost of log data aggregation and retention, eliminating the need for servers, tape libraries, and archival administrators. If the network grows, you can simply rack and stack additional appliances as needed.

Note: LogLogic LMI has been used to represent LogLogic LMI, LogLogic EVA, and LogLogic EVA - Container Edition. All processes and procedures applicable to LogLogic LMI are also applicable to LogLogic EVA and LogLogic EVA - Container Edition, unless explicitly stated otherwise.

Appliance Functions

There are two primary user types on a TIBCO LogLogic® appliance:

- Administrator - Configures and maintains the appliance itself, including managing log sources, user accounts, appliance configurations, running backups, and more.
- User - Monitors appliance operations, runs searches, manages alerts, and creates and runs reports based on collected data.

The appliance GUI provides access to administrator and user functions.

- Administrators can perform all functions on the appliance. *TIBCO LogLogic® Log*

Management Intelligence Administration describes administrator tasks and functions. The default `admin` user can perform certain tasks and functions exclusively. Such tasks are exclusively indicated in the relevant sections.

- Users can perform functions that have been assigned to them by the system administrator, and these might include a subset of the administrator task and functions. The *TIBCO LogLogic® Log Management Intelligence User Guide* describes user tasks and functions.

Reports, search, and alert functions can be opened by clicking their respective icons on the Home page or by clicking their buttons on the top menu on the Home page. Dashboard, management, and administration functions for the appliance can be accessed from the top menu on the Home page. For more information, see the *TIBCO LogLogic® Log Management Intelligence Administration*.

Online Help can be opened by clicking the **Help** icon on the top right corner. Brief video tutorials provide tips and guidance by example for many new LogLogic features. The tutorials are accessible from the Home page and from certain application pages.

LogLogic Appliance Home Page

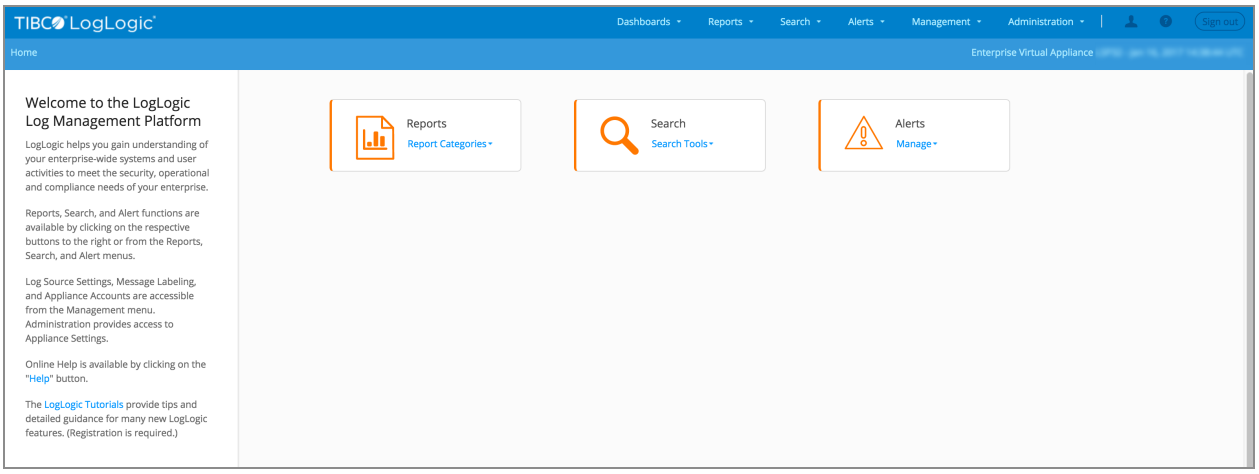

Note:

- The functions in the navigation menu vary depending on the appliance product family. For example, a TIBCO LogLogic® ST Appliance displays fewer options than a TIBCO LogLogic® LX Appliance because certain features are not available on LogLogic® ST Appliances. In addition, reports might show different entries, depending on the TIBCO LogLogic® Log Source Packages installed.
- For all text fields throughout the GUI, null is not a valid entry.

TIBCO LogLogic® Product Families

TIBCO LogLogic® appliances bring visibility of compliance activity metrics to CIOs and CSOs, and control over activities to the compliance team, permitting them to review the compliance timeliness and compliance posture mandated by Sarbanes-Oxley (SOX) and Payment Card Industry Data Security Standard (PCIDSS).

TIBCO LogLogic® appliances provide the highest log collection and analysis performance amongst all log management vendors. Log events are received and indexed in real-time. The TIBCO LogLogic® appliances have clearly stated metrics that cannot be matched.

Hardware Product Families

TIBCO LogLogic® offers the following hardware product families to provide better, faster, and smarter log management, database security, and regulatory compliance solutions to corporations:

- TIBCO LogLogic® LX Appliances are purpose-built appliances for real-time log data collection and analysis. These appliances slash response times to network security and utilization incidents, boost IT productivity, and reduce the corporate cost of security and performance event remediation.
- TIBCO LogLogic® MX Appliances perform real-time log data collection and analysis ideal for mid-size and large companies. These appliances slash response times to network security and utilization incidents, boost IT productivity, and are optimized to provide for log data needs in a non-enterprise environment.
- TIBCO LogLogic® ST Appliances automate the entire log data archival process,

minimizing administration costs while providing more secure log data capture and retention.

To view photographs of the appliance layout, see *TIBCO LogLogic® Log Management Intelligence Hardware Installation Guide*.

Virtual Editions

TIBCO LogLogic® Enterprise Virtual Appliance provides an all-in-one software solution for log management.

TIBCO LogLogic® Enterprise Virtual Appliance and TIBCO LogLogic® Enterprise Virtual Appliance - Container Edition are virtual versions of the TIBCO LogLogic® Log Management Intelligence hardware appliance, and are distributed on the TIBCO eDelivery website.

For a comparison between the platforms and deployment instructions, see *TIBCO LogLogic® Enterprise Virtual Appliance Quick Start.*

LogLogic® LX Appliance Product Family

Featuring a parallel processing architecture, the LogLogic LX Appliances centralize log data collection and retention by simultaneously processing raw log data and metalog data at high volume.

Distributed real-time reporting and targeted queries let administrators take immediate action on network issues from a centralized management console.

These appliances help enterprises harness the power of log data for a safer, more reliable network, while reducing corporate IT costs and providing a rapid return on investment.

Benefits

Appliances in the LogLogic LX Appliance product family offer the following benefits:

- Real-time reports, ad-hoc queries, and fast drill-downs to speed up identification, isolation, and repair of security and network incidents.
- Non-disruptive installation and plug-and-play operation: no changes to network configurations; no dependencies on other systems; no training required; available in

minutes.

- Self-maintaining, embedded database technology eliminates the need for database administration.

To view photographs of the appliance layout, see *TIBCO LogLogic® Log Management Intelligence Hardware Installation Guide*.

LogLogic® MX Appliance Product Family

The LogLogic MX Appliances centralize log data collection and retention by simultaneously processing raw log data and metalog data at any volume.

Designed specifically for mid-size and large companies, LogLogic MX Appliances provide the disk space and processing power required for most non-enterprise environments.

LogLogic MX Appliance features are focused on those needed to harness the power of log data for a safer, more reliable network, while reducing corporate IT costs and providing a rapid return on investment. LogLogic MX Appliances are designed for installations where data must be retained longer than LogLogic LX Appliances provide, but where managing other log appliances is not required.

Benefits

LogLogic MX Appliances offer the following benefits:

- Real-time reports, ad-hoc queries and fast drill-downs to speed up identification, isolation and repair of security and network incidents.
- Features and specifications targeted specifically to mid-size and large companies.
- Self-maintaining, embedded database technology eliminates the need for database administration.

To view photographs of the appliance layout, see *TIBCO LogLogic® Log Management Intelligence Hardware Installation Guide*.

LogLogic® ST Appliance Product Family

The LogLogic ST Appliances archive up to 10 years of log data while eliminating the need for servers, tape libraries, and archive administrators.

They are available in compact, rack-mountable systems with up to 8 terabytes of storage and interfaces to NAS devices. The LogLogic ST SAN (Storage Area Network) appliances offer virtually unlimited archive storage.

When used with LogLogic LX Appliances, LogLogic ST Appliances guarantee the complete and accurate transmission of network equipment logs from anywhere on the enterprise WAN or LAN. LogLogic ST Appliances feature an n-Tier architecture controlled by a management console that centralizes long-term log data archival while allowing for distributed log analysis and broader data accessibility.

Benefits

LogLogic ST Appliances offer the following benefits:

- High volume log data aggregation from centralized and remote log data sources.
- Long-term retention of unaltered, complete, raw log messages at a secure, central location to make archives unimpeachable.
- Distributed architecture of remote collection and central storage make log data collection and retention infinitely scalable.
- Self-maintaining, embedded database technology eliminates the need for database administration.

To view photographs of the appliance layout, see *TIBCO LogLogic® Log Management Intelligence Hardware Installation Guide*.

LogLogic® EVA and LogLogic® EVA - Container Edition

LogLogic EVA is optimized for virtual and cloud platforms such as VMware server, Kernel-based Virtual Machine (KVM), Amazon Web Services (AWS), and Microsoft Azure; and provides an all-in-one software solution for log management.

LogLogic EVA - Container Edition is meant for container platforms such as Docker.

It helps you derive actionable information by capturing, indexing, and compressing log files and flow data.

Benefits

LogLogic EVA and LogLogic® EVA - Container Edition offer the following benefits:

- Provides all the alerting, searching, and reporting you need for both near-term activity and archived data.
- Real-time monitoring lets you alert as needed and monitor system behavior such as VPN session tracking, application distribution, port monitoring, hardware health, and security stance.
- Better IT operations, which leads to reduced time-to-resolution, simplified and improved security, improved IT efficiency, and best-practice implementations.
- Provides easy scalability.

For a comparison between the platforms and deployment instructions, see *TIBCO LogLogic® Enterprise Virtual Appliance Quick Start.*

Scalable Infrastructure

The scalable LogLogic network infrastructure significantly accelerates response time to data center security and availability events, while providing complete log data archives for compliance and legal protection.

LogLogic appliances make log data in enterprise networks truly useful for the first time, improving corporate security, compliance, and network availability, while reducing IT costs and costly network downtime and improving corporations' return on IT investment.

Manage Appliances with Management Station

A Management Station is a single LogLogic appliance configured to perform log management and appliance administration tasks on other remote LogLogic appliances.

It is easier to manage multiple appliances by using the Management Station feature than by configuring and controlling appliances individually.

Any LogLogic appliance can be a Management Station except the LogLogic MX Appliances. Any LogLogic LX Appliance or LogLogic ST Appliance Management Station can access any LogLogic LX Appliance, LogLogic MX Appliance, or LogLogic ST Appliance within your network to perform administration or reporting functions.

Introduction to Management Stations

Note: Only the admin user (the default admin account) has access to the Management Station menu. If the default admin user is disabled, any other user with administrator access rights can access the menu.

Once you configure a LogLogic appliance as a Management Station, you can use that appliance to:

- View and use the GUI on any managed appliance from within the Management Station appliance
- Perform any appliance management function and initiate any report directly on any managed appliance
- Aggregate report results from managed appliances to view a single report containing results from log sources spanning those multiple appliances
- Push the user database of an appliance out to as many appliances as you need

Management Station Clusters

A Management Station cluster consists of a Management Station appliance and all the appliances that you manage from it.

To create a cluster on one appliance, add another appliance under the Management Station (see Creating a Management Station Cluster). This automatically converts the first appliance into a Management Station that now manages the second appliance, plus any other appliances you later add to the cluster.

Once added to a cluster, you can manage an appliance individually from the Management Station by setting the appliance as the Current Appliance. See Designating a Management Station Appliance as Current.

Note: If the Super admin user is disabled, any user who has appropriate admin rights can access or add remote appliances to a cluster.

Warning:

- Starting with Release 5.3, all users, including 'admin', must explicitly be given access to each of the remote appliances for users to have access to the data on the remote appliances.
- Access appliances using the appliances tab of the User Edit page. The Appliances tab is displayed only when an appliance is a Management Station.
- After upgrade, if a report is not showing the same data as before it means the user does not have access to the appliance(s) the report refers to. There is no indication (in the GUI or reports) that data from inaccessible appliances is missing.

Creating a Management Station Cluster

Any TIBCO LogLogic® appliance, except the LogLogic MX Appliances, can be used as a Management Station.

From the LogLogic MX Appliances, only MXVirtual can be used as a Management Station. On any such appliance, after you add another appliance through the Management Station feature, the appliance becomes a Management Station. You can then add more appliances

to the cluster at any time.

All appliances in a Management Station cluster must be running exactly the same LogLogic LMI and LogLogic LSP versions (the hot fix versions need not match)

For information about upgrading the appliances in a Management Station cluster, keeping them on the exact same software release, see the *TIBCO LogLogic® Log Configuration Guides*.

Procedure

1. On the GUI of the Management Station appliance, go to **Management > Management Station**.

   The **Appliances** tab displays a list of appliances, if any.

2. To add a new appliance, click the **Add New Appliance** icon.
3. On the **General** tab:
   a. Enter the IP address or DNS name of the appliance to add to the cluster.
   b. (Optional) Enter a name for the appliance. This unique name distinguishes it from the other appliances in the cluster.
   c. Select the type of appliance that you are adding to the cluster: LX, ST, or MX. If the remote appliance is available, the appliance type from the remote appliance is used.
4. Click **Save**.

   The appliance is added to the cluster.

Result

After you add at least one other appliance to the cluster, the system automatically adds your appliance as IP address 127.0.0.1 to the cluster, and converts it to a Management Station. The menu item **Management Station Dashboard** appears in the **Dashboard** menu.

The **Appliances** tab displays all appliances in the cluster. If you remove all the appliances from the cluster, the Management Station reverts back to being a single LogLogic LX Appliance or LogLogic ST Appliance.

To select the current managed appliance, click the chain icon for that appliance. The GUI of the managed appliance is displayed and you can perform tasks on the remote appliance.

Addition of Appliances to a Management Station Cluster

You can add appliances to a Management Station cluster at any time.

Use the **Configuration** tab procedure in Creating a Management Station Cluster.

> **Warning:** When adding an appliance to a Management Station cluster, the appliances use two-way SSL to verify each other's certificates. If a time gap exists between the appliances, such as different time zones, the earlier appliance could interpret the certificate of the later appliance certificate as being in the future and not accept its certificate as valid.
>
> To avoid this issue, either ensure both appliances have similar time settings or wait until the creation time of both appliance certificates are in the past according to both time zones.

Management of Appliances from a Management Station

From a Management Station, you can:

- Performing Tasks on a Remotely Managed Appliance on a remote managed appliance
- Monitor all managed appliances at once
- Perform Management Station and Regular Appliance Features using managed appliance log data

Performing Tasks on a Remotely Managed Appliance

You can perform tasks on a remote appliance as if you were directly logged into the managed appliance itself.

Procedure

1. Make the managed appliance the current appliance, as described in Designating a Management Station Appliance as Current.
2. The GUI switches to showing the dashboard for the selected managed appliance. The Current Appliance in the top left lists the managed appliance.
3. Perform the task on the managed appliance.

> **Note:** To modify LogLogic® Log Source Packages based log sources on a managed appliance, you must install LogLogic LSP on the Management Station appliance as well as the managed appliance. If you add LogLogic LSP-based log sources to a managed appliance and from the Management Station try to modify one, an error is displayed.

Monitoring all Managed Appliances at Once

You can view the status of your entire Management Station cluster together.

Procedure

1. Open the **Management Station** dashboard.
2. View the individual statistics for each managed appliance, as well as aggregate counts of all new alerts and overall message rate counts across all managed appliances.

Managed Station Status

After you configure an appliance as a Management Station and add appliances to its cluster, **Dashboards > Management Station Status** displays the following information:

- Real-time, condensed status of each managed appliance in the cluster
- Message rate for the default managed appliance
- Aggregate new alerts and message counters across all managed appliances

The colored square by each managed appliance indicates the health of the communication to the managed appliance:

- Green square—managed appliance status received
- Red square—failed SSL tunnel between the nodes due to incorrectly configured certificates
- Clear square—managed appliance status is queried but not yet received

The Management Station dashboard displays the following information:

- Managed Appliances list - Displays message statistics for each managed appliance:
    - Total, Processed, Dropped, Unapproved, Skipped—Displays the number of messages processes on each managed appliance for each message category. Clicking a number in these columns toggles the displayed values between exact numbers and rounded numbers.
    - Message Rate/Sec—Displays the message rate per second for the managed appliance by 1, 5, and 15 minutes. Clicking a message rate value for a managed appliance switches the **Message Rate** graph to 2, 10, and 30-hour timescales, respectively, for that appliance.
    - Time Skew—Displays the time difference, in seconds, between the Management Station appliance and the managed appliance. Clicking a number in this column toggles the displayed values between exact numbers and rounded numbers.
- Message Rate - Graphically displays message traffic over 5, 10, or 15 minute segments for the current appliance. By default, message traffic for the Management Station's appliances is updated every 60 seconds.
    - Pink line—the average number of messages per time segment
    - Blue line—the real time incoming message rate for the appliance
    - Red line—the message rate is at the maximum for the appliance (appears only when the maximum is reached, as a flat line at the maximum level)
- New Alerts - Lists the number of alerts on all managed appliances in the last hour, 6 hours, and 12 hours. Alerts are displayed based on severity (high, medium, low). To view the alerts, click on the displayed number.
- Message Counters - Provides statistics for all managed appliances on each message category, as separately listed previously for each managed appliance. This is useful in calculating Data Retention Settings and maximum syslog message rates.

The message categories are:

Message Category	Description
Total Received	Total number of incoming messages for all categories for all appliances
Processed	Total number of messages received and parsed into the database
Skipped	Total number of messages ignored by the appliance because the associated log source is disabled
Unapproved	Messages received from a log source that is not in the **Manage Devices** table. These messages are discarded. The most recent 100 messages are accessible from the Log Source Status screen. (If auto-identify is on, all messages are auto-identified and no messages are unapproved.)
Dropped	Total number of messages recognized but not processed due to network congestion

Note: It is difficult to troubleshoot why messages are dropped because these messages are not dropped by the application. Though the OS is responsible for tracking and dropping messages, as such there isn't any way to determine why the messages were dropped or where they came from as the OS does not record this information.

Tasks using Managed Appliance Log Data

You can perform various administrative and user tasks affecting one or all managed Appliances.

For more information, see Management Station and Regular Appliance Features.

Designating a Management Station Appliance as Current

To perform tasks on a managed appliance, you must designate it as the current appliance.

To select a managed appliance to be the current appliance on the Management Station, you can do one of the following:

- Use the Appliances tab:
  1. Go to the **Management > Management Station > Appliances** tab.

     A list of all managed appliances is displayed.
  2. Select the appliance.
  3. Click the chain icon to set it as the current appliance.
- Using the Management Status tab:

  Go to the **Dashboards > Management Station Status** tab.

  This option is available only if you have at least one remote appliance. All appliance features perform the corresponding functionality on the selected remote appliance.

Management Station and Regular Appliance Features

You can perform various administrative and user tasks affecting one or all managed appliances. Several regular appliance features are expanded to provide greater functionality in the Management Station configuration.

These features are in addition to Management Stations having access to run all features on managed appliances.

Use of the expanded Management Station components of these features is documented in more detail in the following sections.

- Managed Appliance - Administration Tasks
- Managed Appliance - Configuration Tasks
- Management Appliance - User Tasks

Managed Appliance - Administration Tasks

Several administrative tasks affecting managed appliances can be performed from a Management Station.

Task	Reference topic
Using multiple log sources across multiple Appliances together by creating and managing Global Groups	Device Group Management
For user management:	
• Adding or modifying users on selected managed Appliances • Copying user accounts and their configurations from a Management Station onto managed Appliances • Specifying whether a user has permission to access the Management Station	• Adding or Modifying Users on a Managed Appliance • Replicating Users on a Managed Appliance • User Privileges
Proliferating login page customizations to managed Appliances	Customizing the Login Page

Managed Appliance - Configuration Tasks

The following appliance configuration tasks involve managed appliances in a Management Station cluster:

Task	See this topic
Use the Management Station as an NTP server if the appliances cannot access an Internet-based NTP server	Setting the Time Zone
Use caution when making Management Station cluster changes involving an appliance in a High Availability (HA) pair	Failover

Managed Appliance - User Tasks

From a Management Station, you can perform the following user tasks involving managed appliances:

- For the classic Index search and Regular Expression search:
    - Search log data from a single managed appliance or all managed appliances
    - Select a Global group of log sources on which to perform the search
- Monitor alerts in Alert Viewer for any managed appliance or aggregated for all appliances in a single list
- For report generation:
    - Use the log data from a single managed appliance or all managed appliances
    - Aggregate reports run against all appliances into a single large report or generate them at once in separate reports for each appliance
    - Select a Global group of log sources on which to generate the report
- Run a query on Remote Appliances to which the user has appropriate access permissions.
- Create an aggregation rule that is automatically distributed to Remote Appliances to which the user has appropriate access permissions.
- Run distributed Advanced Search

For more information about these user tasks, see the *TIBCO LogLogic® Log Management Intelligence User Guide*.

Ports

You must ensure that the relevant ports are open for the Management Station and the Remote Appliances to communicate with each other through firewalls.

Task	Port numbers
Distributed Regex search	5514
Distributed Advanced Search	9683, 9685
Remote Appliance communication via Management Station	9443
Advanced Search	9620
Distributed aggregation	9685, 9626

For detailed information about ports, see LogLogic LMI Ports.

Log Source Management

This section includes the following information.

Device Management

Use the **Management > Devices** tab to configure devices (log sources) associated with the appliance.

You can manage log sources to specify:

- Log sources that are allowed to send log messages to the appliance
- Log sources from which appliances retrieve logs through file transfer
- Log sources from which appliances receive SNMP traps
- Log source groupings that enhance and simplify reporting, routing, searching, alerting, and so on

Use the **Devices** tab to view log sources that have been added to the appliance.

If the appliance does not determine the log source type, it is assigned as a General Syslog log source type. It is good practice to manually change the log source type by selecting a log source from the **Devices** tab, which opens the Modify Device tab. Manually changing the device type does not guarantee making the logs for that source parsable. Moreover, by changing the device type of a known source, the deep parsing for reporting purposes might fail, because the rules do not match.

Note: If you have over 4000 log sources, then go to the **Administration > System Settings > General** tab and enable the **Optimize Device Selection List > Show Only Device Groups** option.

- To add a new log source to the appliance, click the **Add New** button.
- To modify an existing log source, click the log source's **Device Name**.
- To modify multiple log sources together, you can:
    - Update the device type for multiple log sources, by clicking their check boxes, selecting from the **Device Type** list, and clicking **Update Type**. See Updating the Device Type.
    - Enable or disable device name resolution for multiple log sources, by clicking their check boxes, selecting from the **DNS Resolving** list, and clicking **Update Name Resolution**. See Updating Device Name Resolution.

      Note: This overrides the **System Settings > DNS Resolve All Device Names** setting. This update occurs asynchronously and is the only way to immediately update a log source name, through a reverse DNS refresh, without waiting for the daily update.
- To remove a log source from the appliance, check the log source's check box and then click **Remove** button.

Collector Domain

Using the Collector Domain feature, users can provide custom identification for different domains throughout their environment, when used in conjunction with TIBCO LogLogic® Universal Collector.

Having this capability allows for unique identification of distinct log sources that use duplicate IP addresses across multiple domains within the same organization. For example, a large organization with multiple satellite locations may re-use the `192.168.0.0/16` IP range throughout its environment. Without the Collector Domain identification introduced by the LogLogic® Universal Collector, the log sources sharing the same IP would all be logged under one combined log source in LogLogic® LMI.

Users can configure a ULDP collector to support a domain identifier from a LogLogic® Universal Collector agent. Log data can be sent from different domains that share the same IP addresses. On the LogLogic LMI these are listed as separate devices. LogLogic LMI supports both devices added on LogLogic Universal Collector and devices added on the LogLogic LMI. If a device is auto-discovered that has a Collector Domain, the domain ID is prefixed to the device name.

Auto-discovery supports the naming convention `domain ID_IP address_device type`, for example: `1_10.5.20.1_pix`

If during auto-discovery a device has the same name as an existing device a random number is appended to the source IP for the newly discovered device.

To add a device through the LogLogic LMI with a Collector Domain ID follow the steps in Adding or Modifying Log Sources.

> **Note:** The device creation described here concerns when a collector domain is specified.

Viewing Devices

You can view devices in different ways.

View: All Devices

The default setting; shows a list of all available devices on an Appliance. (Click drop-down arrow to see other **View** options.). Click the **List** link from the **Groups** column to view the group membership information.

Manage Devices – View All Devices

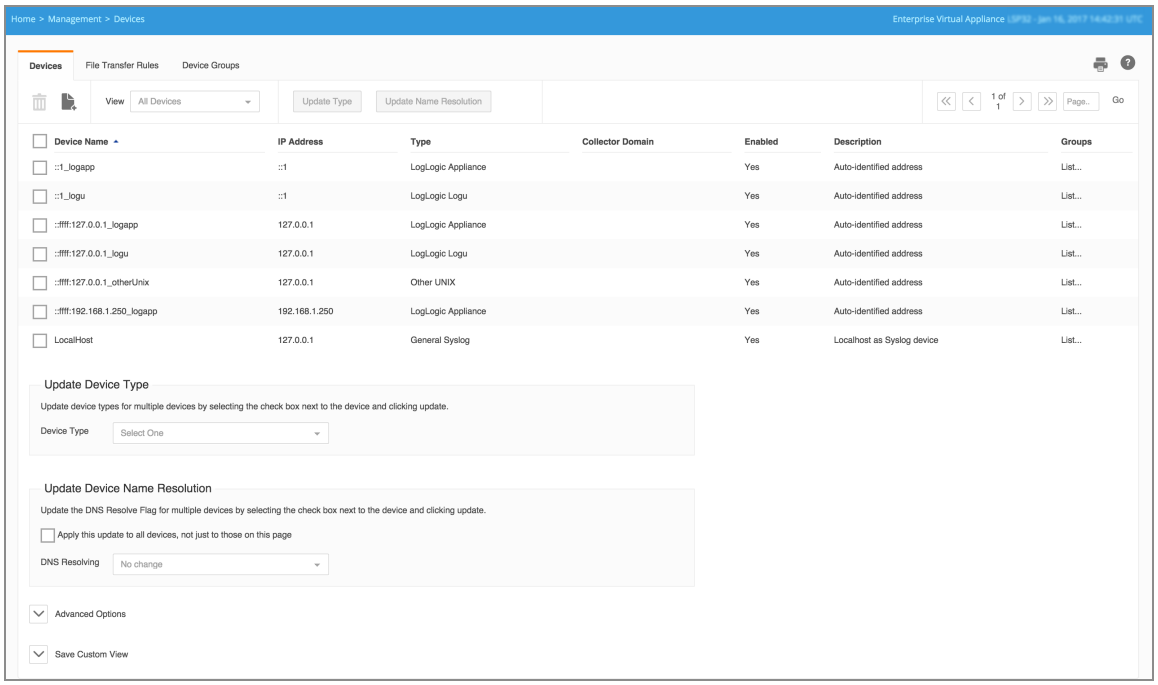

View: Shared IP Devices

Reveals multiple device entries created for a single physical device. The Appliance records different log entries from the same physical device when the log messages are of a different type. This results in same device appearing multiple times. The user may want to examine the details of the duplicated devices, and possibly remove selected ones. Click the **List** link from the **Groups** column to view the group membership information.

View: Stale Devices (0 logs)

Reveals devices that have no incoming log messages. The user can view details of the inactive devices, save the list, and remove selected devices.

Note: “Stale Devices” tend to be localhost or dummy test devices, no longer in use. Or they could be old devices that are not active or have been replaced.

Devices with associated logs do not appear on this list. Click the **List** link from the **Groups** column to view the group membership information.

For appliances with many log sources, you can create customized views of the **Devices** tab that filter the log source information.

Procedure

1. Under **Advanced Options**, select the columns you want to display and any filter expressions for what's included in those columns.

   > **Note:** Days of Inactivity is visible only when Stale Devices (0 Logs) is selected from the **Devices > View** list. 30 days is the default setting; can be set for 1 to 90 days that the selected device has not produced any logs; after which it is considered "stale" by the appliance. The Display box and Days of Inactivity radio button are grayed out because Days of Inactivity is not a displayable field, and hence not sortable.

2. Click **Filter**.

   The **Devices** tab refreshes based on the advanced options selected.

3. Under **Save Custom View**, enter a name and description for this customized view, and indicate whether to share the view with other users.

4. Click **Save View**.

   The **View** list now includes this view.

   Manage Devices – View Advanced Options and Save Custom View

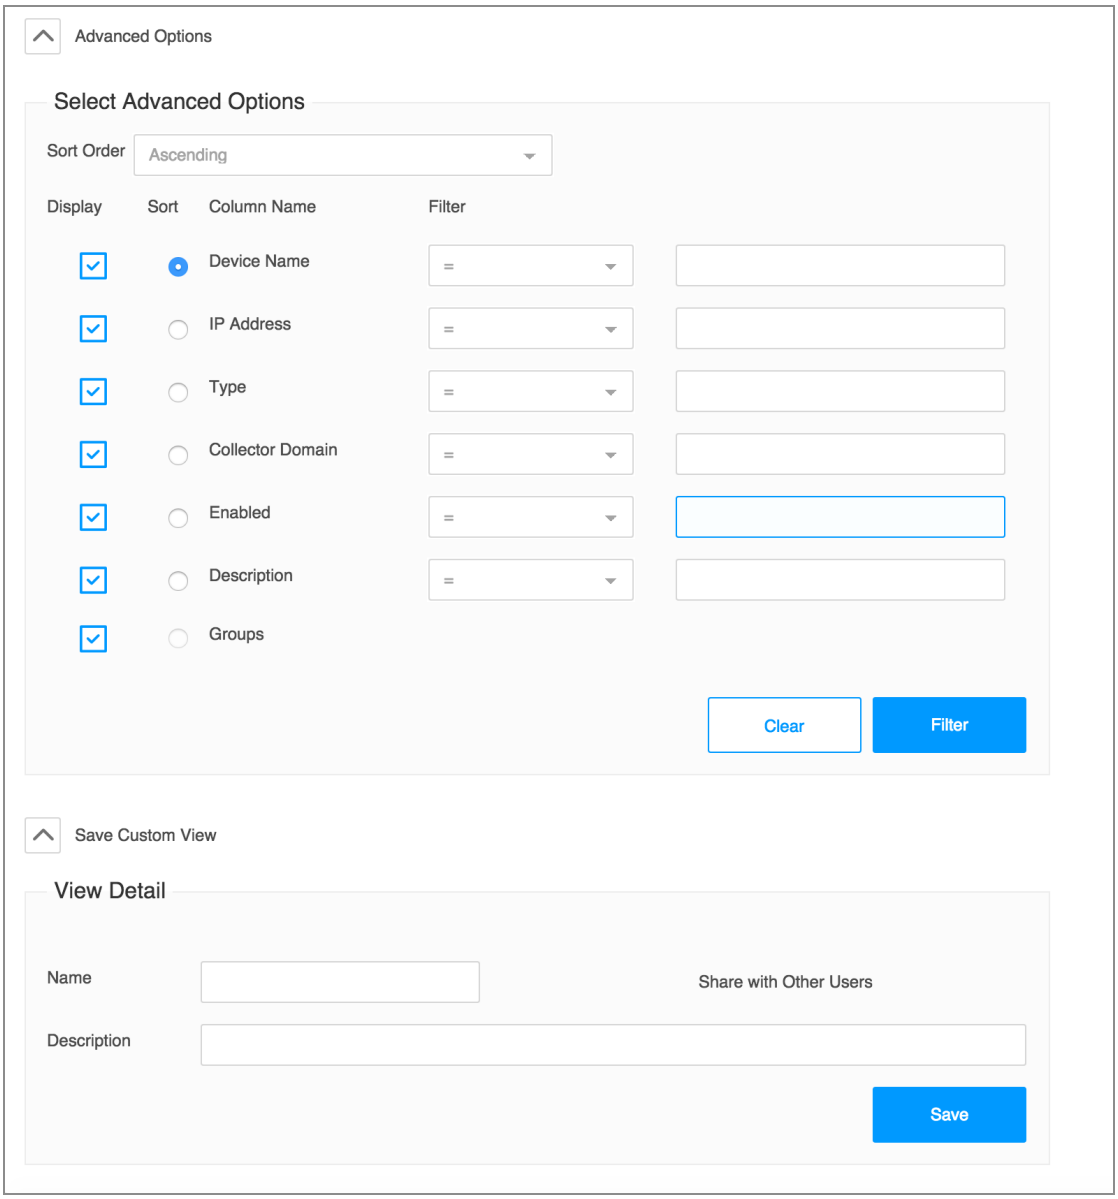

Adding or Modifying Log Sources

You can add or modify log sources from the **Management > Devices > Devices** tab.

- To add a new syslog log source, click the **Add Device** icon .
- To modify an existing syslog log source, click an existing syslog log source name from the list.

Perform the following steps to add or modify a log source profile.

Procedure

1. Type the name of the log source. The length of the log source name should not exceed 63 characters.
2. Type a description of the log source.
3. From the **Device Type** list, select the type of log source you are adding. This cannot be changed after adding the device profile.

   Note:
   - The **File Transfer Rule** tab displays only if you select a device type that supports file transfer rules. Otherwise, the tab does not display.
   - Selecting the **LogLogic Database Security Manager** device type causes the appliance to hang, because the device type is no longer supported. Do not use this device type.

4. In the **Host IP** field, enter the IP address of the log source. When logs arrive from the specified device type, the LogLogic appliance uses the IP address of the added device to map the logs against the device instance.

   Note: For log sources such as Amazon S3 or Tibco® Mashery, you must enter a valid IP address because the field is mandatory. The value in this field is used only for indexing purposes and not for connecting to the log source.

5. In the **Collector Domain** field, enter an identification name to be used for identifying each message sent from this device. This field can be empty. If defined, it must be an unique name with a maximum of 256 characters. Do not include special characters, for example, \ | / " ? * : %. This field is also case sensitive.
6. Under **Enable Data Collection**, select the **Yes** radio button to accept logs from this log source.
7. Select **Refresh Device Name through the DNS Lookups** to have the **Name** field automatically updated with a name obtained through a reverse DNS lookup at the refresh interval configured in the **General tab Refresh Auto-Identified Device Interval** field on the **Administration > System Settings** page. The DNS name overrides any name you assign in the **Name** field.

8. In the **Polling Interval** field, enter the number of minutes between polls to retrieve log data from the Oracle database. The polling interval applies to all Oracle database instances configured for the log source. For example, to poll the Oracle database once every hour, enter 60.
9. (For Blue Coat Proxy SG only)
    - Select the **Use SSL** check box to use SSL to communicate from the appliance to the Blue Coat machine for file transfer.
    - Select the **Use User Authentication** check box to authenticate the user name and password for file transfer from the Blue Coat machine to the appliance. The user name and password should match one of the users listed in the **User** tab.
    - In the **SSL Certificate** field, copy this automatically-generated certificate to the Blue Coat machine. You cannot use SSL without copying the SSL Certificate to your Blue Coat machine. For example, you must copy this certificate on to your Blue Coat machine to enable encryption while transferring files.
10. (For Microsoft SQL Server only) Under the MS SQL Server Collector Configuration section, type in the following information:

Option	Description
Use DBCC TRACEON (optional)	Select this check box to use SQL query “DBCC TRACEON (1903)” before collection of log data.
Use XP Cmd Shell (optional)	Select this check box to use xp_cmdshell.
Authentication	Select SQL Authentication or Windows Authentication.
Domain Name	If you have selected Windows Authentication provide the corresponding domain name of the user.
Database Name	Microsoft SQL Server database instance name
Server Port	Port number for Microsoft SQL Server
UserID	User name for the Microsoft SQL Server sysadmin user or

Option	Description
	Windows Authentication domain user based on the selection of the Authentication type
Password/Confirm Password	Password for the corresponding user authentication type
Polling Interval	Interval in minutes between two instances of data polling
Rows per Collection	Maximum number of rows per collection
No. of Collections	Maximum number of polling intervals per collection run
Trace Files Path	Audit log file name for Microsoft SQL Server. The pathname must be the absolute path to the trace (.trc) file. The LogLogic appliances need to be able to read new trace files that are created after server restart.
Start Collection From Date	Date and time that the LogLogic appliance will begin collecting log data

Note: You can collect data from trace files at multiple locations, to specify different location use the **Add Row** button and enter the trace file path and start time.

11. Click **Add** or **Update** to save your changes.

Copying a Device

You can copy an existing device from the **Devices** tab.

Procedure

1. Navigate to **Management > Devices** tab.
2. Click the name of the device you want to copy.

   The Modify Device page opens.

3. Change any fields of the device, as required.
4. Click **Copy Device** to make a copy of the device.

   The new device is listed in the device list on the Devices tab.

> **Note:**
> - The name of the new device is appended with (copy).
> - As a rule, devices cannot have the same host IP, device type, and collector domain. Therefore, in step 3, you must change at least one of these fields.

Updating the Device Type

Update the device type of an existing device.

> **Note:** Sometimes a file forwarded from LogLogic® Universal Collector may have an empty Device Type field. In this case, use this procedure to enter the Device Type.

Procedure

1. In the devices table, select the device's check box.
2. Under **Update Device Type**, from the **Device Type** list, select the device type to use.
3. Click **Update Type**.

Updating Device Name Resolution

Update the device name resolution for an existing device.

Procedure

1. In the devices table, select the device's check box.
2. Under **Update Device Name Resolution**, from the **DNS Resolving** list, select **Enable**.
3. To update the device name resolution for all devices, select the **Apply this update to**

all devices, not just to those on this page check box.

> **Note:** This overrides any option you have defined in the **Administration > System Settings > General > DNS Resolve All Device Names** radio button. This update occurs asynchronously and is the only way to immediately update a device name, through a reverse DNS refresh, without waiting for the daily update.

4. To execute a full query for device name resolution when the **Update Name Resolution** is clicked, select the **Full DNS query to all selected devices** check box.

> **Note:** Any device names saved in the `/etc/host` file of the LogLogic LMI appliance override the device names on the DNS server.

5. Click **Update Name Resolution**.

Removing Devices

Procedure

1. Select the check box next to the device, and click 🗑.

   The **Remove Devices** tab appears. This tab lets you confirm the removal of the selected log sources from the appliance.

2. Click **Remove**.
3. To cancel the removal of the listed log sources, leaving them on the appliance, click **Cancel**.

File Transfer Rules

Using the File Transfer feature, you can transfer and process log files from a supported log source or application.

The **Management > Devices > Devices** tab and **Management > Devices > File Transfer Rules** together provide the File Transfer feature.

- Configure the remote log source or application that generates/sends logs as a File Transfer device using the **Management > Devices > Devices** tab in the appliance.
- Specify a set of policies or rules for the file transfer using the **Management > Devices > File Transfer Rules** tab.

Note: A command line argument can be used to enable the file collector to ignore files after a certain size (bytes). For more information, contact TIBCO Support.

File Transfer Protocols

File-based logs can be retrieved using various protocols. After the logs are retrieved, they are parsed accordingly.

Protocol	Public Key Copy Required	Supported Search Methods
SFTP	yes	wildcard, CSV (for example, “File1, File2, File3)
SCP	yes	wildcard, CSV (for example, “File1, File2, File3)
HTTPS	no	wildcard, CSV (for example, “File1, File2, File3)
FTP	no	wildcard, CSV (for example, “File1, File2, File3)
FTPS	no	wildcard, CSV (for example, “File1, File2, File3)
CIFS	no	wildcard, CSV (for example, “File1, File2, File3)
HDFS	no	wildcard, CSV (for example, “File1, File2, File3)
Amazon S3	no	wildcard, CSV (for example, “File1, File2, File3)

Compressed Files Collection

To shorten the list of files to transfer, the appliance supports collection of many compressed file formats.

Some of these formats include:

.tar.bz2	.tar.gz	.tar.Z	.tar.z
.tgz	.taz	.tar	.gz
.z	.Z	.zip	.ZIP

No compressed directories or files are allowed in the compressed file (just one flat level). The type of compression is determined by the compressed filename extension.

Compressed `.gz` files created on a 64-bit appliance can be received and extracted on the LogLogic appliance.

Adding a Log Source for File Transfer

You must use the **Add File Transfer** tab to add a remote log source from which you intend to transfer files.

After adding all the remote log sources, you can specify rules using the File Transfer Rules feature.

Procedure

1. On the **Management > Devices > Devices** tab, click  .
2. In the **Name** field, type a name for the log source.
3. Type an optional description for the remote log source.
4. From the **Device Type** list, select the type of log source or application generating the logs to be transferred.
5. In the **Collector Domain** field, enter an identification name that to be used for identifying each message sent from this device. This field can be empty. If defined, it must be a unique name with a maximum of 256 characters. Do not include special characters, for example, \| / " ? * : %. This field is also case sensitive.
6. In the **Host IP** field, type the IP address of the log source from which you want to transfer files.
7. Click the radio button to indicate whether to retrieve log files from this log source.

8. Select **Refresh Device Name through the DNS Lookups** to have the Name field automatically updated with one obtained through a reverse DNS lookup on the configured refresh interval. Configure the refresh interval in the **Refresh Auto-Identified Device Interval** field on the **Administration > System Settings > General tab**. The DNS name overrides any name you type in the Name field.
9. If you selected a log source that uses SSL or user authentication:
   - Select the **Use SSL** check box to use SSL to communicate from your appliance to the log source for file transfer.
   - Select the **Use User Authentication** check box to authenticate the user name and password for file transfer from the log source to your appliance. The user name and password should match a user listed in the **Management > Users > Users** tab.
   - Copy the auto-generated SSL Certificate on to the log source to enable encryption while transferring files. You cannot use SSL without copying the SSL Certificate to the log source.
10. Click **Add** to add the log source.

File Transfer Rules

Use the **Add File Transfer Rule** tab to specify the policy for transferring one or more files from a File Transfer log source.

The appliance computes the checksum for each file and stores it in the database. Both SHA256 and MD5 are supported. The appliance uses the checksum to determine if a log file has been modified since it was last transferred. If the checksum remains the same, the file has not been changed, and therefore is not processed by the appliance.

Some file transfer protocols require the file to be collected by LogLogic LMI before the checksum can be computed.

Note: To include transferred files in the daily summary reports, make sure you are transferring files from the same day that the messages are generated. Also, make sure the time is accurate on the log source from which you are downloading files.

Adding File Transfer Rules

You can add and modify the existing File transfer rules from the **Management > Devices > File Transfer Rules** tab. The options on both tabs are the same.

Before you begin

- For using an Amazon S3 file collector, the user must have:
  - A bucket created in Amazon, to be able to pull files from it
  - Read permissions on the bucket
  - An access key and secret key pair
- If the new file transfer rule uses the SFTP or SCP protocol, perform a `keycopy` to the specified server for the specified user. See system Command.

Procedure

1. On the **Management > Devices > File Transfer Rules** tab, select **Device Type** and **Device**, and then click Add New  .
2. In the **Rule Name** field, specify a name for the rule.
3. From the **Protocol** list, select the protocol type to use to transfer files. The supported protocols include:
   - SFTP
   - SCP
   - HTTP
   - HTTPS
   - FTP
   - FTPS
   - HDFS
   - CIFS
   - Amazon S3

   For details about protocols, see File Transfer Protocols.

Protocol	Notes
SCP	If you specify a wildcard, ensure that the list of files (list of file names only and not the file sizes) does not exceed 32KB. If the list of files exceeds 32KB, consider combining the files to reduce the size of the list to a single file or fewer files. Depending on the number of files you have, you might want to combine the files into an hourly, daily, or weekly file. For example, if you have a large number of `.aud` files in your specified file transfer directory, you can run the command `tar jcf today.tar.bz2 *.aud` to create a single file containing all of the .aud files in the directory. Ensure that you configure your SCP file transfer rule to pick up the combined file; in this example `today.tar.bz2`.
HDFS	LogLogic LMI connects to the default HDFS port 9000. Whether using the default or a custom port, ensure that the HDFS cluster is configured to use the same port that LogLogic LMI uses. If a different port is used, change the port in the `--hdfs-port` parameter in the configuration file `/loglogic/conf/fc_hdfs.conf` and then restart `engine_filecollector`.
Amazon S3	• You must adhere to the Amazon Bucket Restrictions and the rules for naming buckets. • Ensure that the correct retention rules for file transfer are set up in LogLogic LMI. • Enter a valid IP address in the **Host IP** field while creating the log source. This IP address is used only for indexing purposes, and not for connecting to the log source.

4. In the **User ID** field, type the ID to use when accessing the file transfer log source.
5. Depending on the protocol you selected, take the appropriate action:

Protocol	Action
FTP, FTPS, HDFS, HTTP, HTTPS, or CIFS	Type the **Password** and verify the password associated with the User ID.

Protocol	Action
CIFS	a. You can specify the **Domain** or workgroup associated with the directory. b. Specify the **Share Name** for the folder containing the files you want to transfer. To find the share name, view the properties of the shared folder.
Amazon S3	Specify the AWS credentials in the **Access Key** and **Secret Key** fields.

6. In the **Files** field, type the absolute path of the file you want to transfer. For example, `/root/user/LogLogic/IIS/ms_iis_ex030131.txt`. Multiple files can be specified using a comma (,) or semicolon (;) as the delimiter for all protocols. For additional information, see the table in File Transfer Protocols.

> **Note:** The administrator of the remote system must verify that the file path does not include symbolic links. If you cannot specify a file path without symbolic links, then relocate the files to a path that can be specified without symbolic links.

Protocol	Comments
SFTP	File transfer of files with `.xml` extension is not supported.
FTP, FTPS, HTTP, and HTTPS	The file size limit of 20 GB is set in the corresponding configuration file. For example, for FTP files, the `/loglogic/conf/fc_ftp.conf` file specifies the file size limit as: `--max-filesize 21474836480`
Amazon S3	The file path must start with the bucket name. For example: `/<bucket-name>/<dirname>/<Log-file-name>`

Protocol	Comments
	or `/<bucket-name>/<Log-file-name>` where: • `<bucket>` indicates a particular Amazon S3 bucket • `<dirname>` indicates the directory created inside S3 bucket, if any • `<filename>` indicates the name of the file

7. From the **File Format** list, select the format of the files to be transferred.
8. Click the **Test** button to check if the connection parameters can be used to successfully retrieve the file without ingesting the data. The **File Transfer Test Status** window appears. Click **Cancel** to cancel the running test.

> **Note:** Run only one test at a time. Initiating a new test aborts the progressing test, even though it is initiated by different user. It is good practice to test the smaller size file since the entire file is downloaded.

9. In the **Collection Time** section, specify the time interval or schedule on which files should be transferred.

Option	Description
Every	Select the number of minutes (in five minute increments) to wait between intervals.
Every	Select the number of hours to wait between intervals.
Daily at	Select the hour at which files should be transferred every day.
Weekly on	Select the day and time at which files should be transferred every week.

> **Note:** The **View** button that is used to view the history is disabled when you are adding the file transfer rule.
>
> By default, file transfer history is not displayed on LogLogic ST Appliances. To change this, execute a database query. Contact TIBCO Support for more information about the query.

10. For **Use Advanced Data Duplication Detection**, select the appropriate option:
    - Click **Yes** for the appliance to try to detect partial duplication of the newly transferred file.
    - Click **No** for the appliance not to detect partial duplication.

    By default, the appliance automatically detects exact duplicates among transferred files. Advanced detection analyzes potential duplicate data between transferred files, even if the files are not identical. (For example, a partial file at the end of a transfer that is repeated in whole at the start of the next transfer.)

    With or without the advanced data duplication detection, the old file is always replaced with the newly transferred file. However, when duplication (complete or partial) is detected, only the non-duplicated portion of the file is processed.
11. Click **Yes** to **Enable** the transfer rule or **No** to disable this rule.
12. Click **Add** to add the rule.

    After you verify that the rule is displayed on the **Management > Devices > FileTransfer Rules** tab, you can click **Dashboards > Log Source Status** to view the transfer as it occurs. This might take several minutes.
13. If the protocol you select from the **Management > Devices > FileTransfer Rules > Add Rule > Add File Transfer Rule** tab requires a public key copy, register your public key on the server from which you transfer files.

Removing File Transfer Rules

Procedure

1. Select the check box next to the rule, and click Delete  .

   The **Remove File Transfer Rule** page appears. This page lets you confirm the removal of the selected file transfer rules from the appliance.

2. Click **Confirm Remove**.
3. To cancel the removal of the listed log sources, leaving them on the appliance, click **Cancel**.

File Collection Parallelism - Overview

File parallelism allows you to configure concurrent file collections at the rule level in the file collector.

By default, LogLogic LMI processes a single file pull rule at a time by delivering files associated with a single rule. Where you feel appropriate and where the file transfer is relatively slow, LogLogic LMI allows for simultaneous file delivery for several rules simultaneously. This creates a higher throughput of file log processing because LogLogic LMI is capable of processing files at a relatively high speed compared to the time it takes to pull the files into LogLogic LMI. You can experiment by increasing the parallelism until there is a backlog in the `/loglogic/data/filecollector/archiver` directory.

File Parallelism is configured through command-line options that pass to the file collector when it is launched. These options are specified in `/loglogic/conf/node_config.xml`.

The number of concurrent file collections is specified by the command-line option `-p x` (where `x` is the degree of parallelism). The maximum is 100, and the default is 1 (when the command-line option is absent).

Note: File collection parallelism is available in LogLogic LMI 5.6.2 and later.

Changing the Number of Concurrent File Collections

Procedure

1. Using the following command stop all engines:

```
$ mtask stop
```

> **Note:** Using `mtask` causes the GUI, log collecting, searching, reporting, and all other LogLogic LMI functions to stop until `mtask` is restarted.

2. Open the `/loglogic/conf/node_config.xml` file.
3. Add the following line to the file collector, as shown in the following example.

   The following example for an LogLogic ST Appliance sets the degree of parallelism to 50:

```
<node type="ST_SERVICES">
<service
   group="BACKEND"
   start_cmd="/loglogic/bin/engine_filecollector"
   args="-p 50"
   heartbeat_timeout="180"
   escalation="GROUP_RESTART,DISABLE"
   runlevel="8"/>
</node>
```

4. Use the following command to start all engines:

```
$ mtask start
```

   The appliance will be fully functional after a few minutes.

Configuring Parallel File Processing and Parallel File Forwarding

This section describes how to enable and configure the number of parallel file processing and file forwarding threads in the file collector and event forwarding process.

By default, a single process handles a pulled log file on LogLogic LMI. For appliances where multiple files are collected simultaneously, LogLogic LMI allows for parallel file processing, that is, parsing, time extraction, and log conversion. This allows for faster file handling. LogLogic LMI provides the flexibility to adjust overall file handling by selecting an appropriate degree of parallelism. As a rule of thumb, the number of parallel processes should be targeted at the number of cores minus 1. If the number is too high, then processing scalability is not linear. LogLogic LMI selects its processing rate automatically.

It is good practice to perform parallel file processing when a backlog of unprocessed files is accumulated in the `/loglogic/data/filecollector/forwarder` directory. If a backlog is detected in this directory, the user can increase file forwarding throughput by making it parallel as well.

Procedure

1. Create the file `/loglogic/conf/fc.conf` and specify the `NumParallelFwdProcs` and `NumParallelParserProcs` parameters in the file. For example:

```
NumParallelFwdProcs=3
NumParallelParserProcs=10
```

2. Run the command:

```
$mtask -s engine_filecollector restart
```

Keyword	Description
`NumParallelFwdProcs`	Indicates the number of processes that forward files in parallel. The maximum value of this parameter is the number of cores on LogLogic LMI. Example: `NumParallelFwdProcs=3`
`NumParallelParserProcs`	Indicates the number of processes for parallel file processing. The maximum value for this parameter is the number of cores on LogLogic LMI. Example: `NumParallelParserProcs=10`

File Collection Merging

File merging is used for log sources that produce a large number of log files.

The file merging process concatenates multiple files into a single log file, thereby increasing the efficiency of file processing. File merging does not work for files pulled by LogLogic® Universal Collector and sent to LogLogic LMI.

File merging can be used in the following scenarios:

- For any native file-pull rule
- When a rule execution results in many files
- For a pull of a single archive file that could be compressed such as `.tar, .taz, .tar.Z, .tar.gz, .tar.bz2, .tgz`
- When a pulled archive contains more than 5 files smaller than 1 MB in size. Otherwise, merging overhead becomes comparable to the overall file processing time.

You can specify archive types such as `.zip` or `.bz2` for merging. When unpacked, the archive creates multiple small files in LogLogic LMI.

Note:
- See the list of sources for which file merging is not supported.
- Do not configure file merging when you want to leave the pulled files unmodified in any way.

To optimize file processing, the merged file should contain logs sorted on the log time stamp. It is assumed that each small file is already sorted by time. While concatenating files during merging, file merging sorts the files by file name.

The following file merging configuration parameters are specified in the `fc.conf` file. Whitespace is not allowed within the parameters.

Keyword	Description	Example
`Merge=<FileNamePattern>`	Specifies file name prefix for the small files so that files that match the prefix pattern are merged. The value of this parameter must be a literal string.	`Merge=filedata_` Selects files such as filedata_1.txt and filedata_2.txt for merging

Keyword	Description	Example
`Merge=*`	The wildcard * indicates that all files should be merged	
`Rule=<RuleName>`	Optional. Limits the merge operation to a specific file pull rule.	`Rule=ruleA`
	Note: • If `Rule=` is not present, the merge applies to all pull rules on LogLogic LMI. • Only one rule can be specified. If more than one rule should use the same merging, the `Merge=/Rule=/SearchKey= config` block should be duplicated in the configuration file for each file-pull rule. • Whitespace characters are not allowed in the rule name.	
`SearchKey=<SortableFieldSearchPattern>`	Optional. Specifies rules to find a sortable portion of the file name to be used for sorting all text files in the `/loglogic/data/filecollector/archiver` folder, which are unpacked from the `.tar` file. Doing so helps simplify the processing of the merged file.	If the archiver file is `19_192.168.1.10_29309_1461094448_7.txt`, then `SearchKey=_3` skips the underscore character three times to find the timestamp 1461094448 in the filename.

The `SearchKey` parameter specifies rules for finding a sortable portion of the file name. When deciding on the SearchKey value, refer to the structure of the pulled archive file on the log source server, and derive the SearchKey based on original file names on the log source server.

The most obvious one is a time stamp embedded into the file name. If at the beginning of the file name, `SearchKey` is not needed. If in the middle of the file name, `SearchKey` specifies a printable character that precedes the timestamp. If several instances of this character precede the timestamp, specify a number immediately after the character. For example, for the file name `file_A_B_12345678.txt`, the timestamp 12345678 is preceded by an underscore. As three underscores are found before the timestamp, `SearchKey=_3` makes sure the timestamp is extracted after the last underscore.

Usually, the user has no control over names of the small files. Whenever possible, the timestamp should appear at the beginning of the file name.

Sorting is performed on the portion of the file name that begins with the fragment; not the entire file name. Sorting is not mandatory; it helps to optimize log processing. After being sorted, the merged files generate a single file in which the logs are sorted. This makes file processing more efficient and the parsed data better aggregated.

Note:

- File merge works even if logs are not sorted.
- File sorting affects subsequent regex search. During regex search, the results are returned in order they are written into BFQ files.
- Sorting during merging does not affect the result of the Index search because the Index search sorts results by time.
- File parsing is not affected by sorting during file merging.

Configuring File Collection Merging

Procedure

1. Create the file `/loglogic/conf/fc.conf` and add the following line in the file:

```
Merge=* Rule=RuleName SearchKey=_3
```

2. Run the command:

```
$mtask -s engine_filecollector restart
```

Sources that are not supported

File merging is not supported for the following sources:

- Data collected via TIBCO LogLogic® Universal Collector
- JDBC-based file pulls
- All of the file extensions: `.gz, .bz2, .zip, .z, .Z`
- `.tar.z` and `.Z` files: Merging of `.tar.z` and `.Z` is not supported as those files are treated purely as compressed files because of their `.z` and `.Z` extensions. LogLogic LMI checks the last extension and hence these files are not treated as archive files.

Device Group Management

From the **Management > Devices > Device Groups** tab, you can group log sources into a single virtual log source.

A log source must be part of the **Available Devices** list before it can be included in a group.

Device Groups are updated dynamically. For example, if you create a device group for all routers, and later add a new router to the appliance, the new router is automatically added to your router device group.

Note: Device groups without a hyperlink are system-generated groups that contain all log sources of a specific type. For example, “All Cisco ESA.” You cannot modify or delete these groups.

You can perform the following tasks on the device group list:

Task	Description
Add or modify a device group	• To add, click the **Add New** icon  . • To modify, click the device group name and then edit the required information. For the detailed procedure, see Adding or Modifying a Device Group.

Task	Description
Modify a device group	From the **Device Groups** tab, click the name of the device group. The **Modify Device Group** tab appears. Enter your changes and click the **Update Device** button. **Note:** You cannot modify system-generated groups.
View devices in a group	Click the **List** link in the **Devices** column.
Filter the groups using the **Find** field	You can quickly find the desired group by typing the group name in the **Find** field. As you start typing a group name or group description in the **Find** field, the **Device Groups** page is automatically refreshed showing your selection.
Remove a device group	Select the group by clicking its check box and then click the **Remove** icon. **Note:** You cannot delete system-generated groups.

If you are running a Management Station, you can multi-select and group log sources across appliances. These global groups are accessible only from the Management Station on which the global group is created. To view global groups in a search or report, you must select All Appliances in the **Appliance** field.

When a Management Station is reverted to being a regular appliance, its global groups are still visible but can only be deleted. When the appliance becomes a Management Station again, the global groups can be used and modified as before.

Adding or Modifying a Device Group

To add or modify a group, use the **Add Device Group** tab to arrange your log sources into bundles and categories.

You can create a group using log sources of the same type or of different types (for example, Cisco PIX/ASA and Juniper Firewall). The options on both tabs are the same.

Procedure

1. From **Management > Devices > Device Groups**, click **Add New**. The **Add Device Group** tab appears.
2. Type a unique **Group Name** to identify the log sources you are grouping.
3. Select the appropriate **Enable** radio button to indicate whether the Group device is activated for your appliances. The default is **Yes**.
4. Select whether this group is a Local or Global group. Once you set the Group type, you cannot change it.

Option	Description
Local	The group contains log sources on the current appliance only.
Global	The group contains log sources on multiple appliances. (Global groups can be created and accessed on Management Station only.)

> **Note:** Global groups:
>
> - cannot contain another global group as a member.
> - are marked with an asterisk (*) in the **Groups** tab.
> - are not supported with IPv6 addresses.

5. Select **Static** (default) or **Dynamic** if you want the new device group to be updated automatically as new devices are added to the appliance.
6. In the **Description** field, type an optional description for the Group device.
7. Use the **Device Filter** fields to search for log sources connected to your appliance that you want to group together. To perform multi searches, search on more than one field.

> **Note:** If a match is found for your search, the results display in the **Available Device** section.

8. Under **Available Devices**, find the devices available that are available to add to the

group. You can use one or any combination of the following fields:

a. In the **Name Pattern** field, type a name of a log source to search for and add to your group. You can use regex wildcards for this search.

b. In the **IP Pattern** field, type an IP address of a log source to search for and add to your group. You can use wildcards for this search. Regex wildcards are not supported.

c. In the **Collector Domain** field, type a collector domain to search for and to add to your group.

d. From the **Device Type** list, select a log source to add to your group. A group can contain log sources of one type or multiple types.

e. In the **Desc Pattern** field, type a description of a log source to search for and add to your group. You can use regex wildcards for this search. The descriptions that you define in the **Add Syslog Device** or **Add File > Transfer Device** screens are the fields that are searched using the Desc Pattern search.

f. (Management Station and Global Group Types only) From the **Appliance** list, select an appliance on which to search for log sources.

9. Click **Filter** to search for log sources on your appliance with the specified search criteria.

   The **Available Device** table lists all devices matching the criteria. The **Available Device** list contains the following information:

   - Appliance—IP address of the appliance which contains the log source (Management Station only).
   - Name—Log source name.
   - IP Address—IP address for the log source.
   - Type—Log source type.
   - Enabled—Indicates whether the log source is enabled or not.
   - Description—Lists the log source description.

Note:
- All devices that appear in the Available Devices list when the Filter button is clicked are added automatically to the Dynamic Group. It is actually not necessary to click the Filter button for this to occur. New devices auto-discovered or manually added to the system are added automatically to the Dynamic Group if the device matches the pattern.
- Dynamic Groups cannot contain Static Groups as members. However, Static Groups can contain Dynamic Groups as members.

10. (For Static Groups Only) In the **Available Device** list, select the check box next to the log source name and click **Add** to add the log source to the **Current Devices in Group** list.
11. The **Current Devices in Group** table lists the log sources you added from the **Available Device** table. You must add at least one log source to this list before you can save your group.
12. (Optional) From the **Current Devices in Group** list, check the log source name and click **Remove** to move the selected log source to the **Available Device** list.
13. Click **Save** to add the group of log sources to the **Groups** tab.

Note:
- A user must have “all device access” to create or update a Dynamic Group.
- A user can be given explicit permission on the Dynamic Group, but if they do not have “all device access”, they can see and use the Group, but cannot edit it.

Removing Device Groups

Procedure

1. Select the check box next to the group name, and click **Remove**.

   The **Remove** Groups tab appears. The **Remove Groups** tab lets you confirm the

removal of the selected log source group from the appliance.

2. Click **Confirm Remove**.
3. To cancel the removal of the listed groups, leaving them on the appliance, click **Cancel**.

Device Types Management

Use the **Management > Device Types** tab to view most common device types that have been added to the appliance.

Viewing the Device Type

When you create a new Device Type, a unique device type ID is automatically associated with that device type.

Since the appliance uses this unique ID to identify and categorize messages, the device type/ID mapping association is enforced in the Management Station environment. This restriction is enforced when importing a new device type into a system. Therefore, two device types with the same device type ID are not allowed to exist in the system. For example, if an imported device type has an ID that is already being used in the system to map to a different device type, then a device type conflict is reported and the import is aborted. To avoid this scenario, it is good practice that device types be created and distributed from a single node (the Management Station) only, and imported into other nodes that are being managed.

The maximum number of device types that can be created is 1500.

Procedure

1. From **Management** menu, select **Device Types**. A list of All Device Types is displayed.
2. To filter the list of Device Types, type a keyword in the **Name** field and press **Enter**. Device types matching only the keyword are displayed. Clear the entry in the **Name** field and press **Enter** to display the entire list of Device Types again.
3. Click the **Show User Defined** check box. The new user-defined Device Type appears next that shows the Status is Active.
4. To remove the user defined device types, select the row and click **Remove selected**.

> **Note:** You can only remove the user defined device types. You cannot remove the appliance stored device types.

5. To confirm the new user-defined Device Type, click **Dashboards > Log Source Status**.

Adding a New Device Type

Procedure

1. Click **Management > Device Types** from the home page.
2. Click **Create**.

   The Add New Device Type window appears.
3. To see Sample Messages over the last 5 minutes:.
   a. Click the **Reload** button, or click the calendar icon to open the **Date and Time Range Picker**.
   b. Select the desired timeframe.
   c. Click the **Reload** button.
4. Limit the sample messages (10, 100, 1000) by clicking the drop-down arrow on the **Limit To** button.
5. Limit the log source by entering a specific **IP** Address.
6. In the **Device Type Attributes** area, enter a Name and a Description for the new Device Type.
7. In the **Regular Expression** field enter a regular expression for the new Device Type, or click the **Reload** button to see Sample Messages displaying the Device Type token that you can use for the regular expression.

> **Note:** While adding or importing a user-defined device type, `.*` is automatically added at the beginning of the regular expression for the device type. For example, if the regular expression is `abc`, the expression saved in the system is `.*abc`. If the regular expression is `abc.*`, the expression saved in the system is `.*abc.*`

8. The **Transport Type** is defaulted to the **Network** radio button. You can change to **File Transfer**, if appropriate.
9. Click **Enabled** to enable the pattern and then click **Save**.

   The new Device Type name and description appears in the list of all Device Types.

Editing or Removing Device Types

Editing existing Device Types

1. To edit the device type, click the Edit icon next to the Device Type Name.

   The Edit Device Types window appears.
2. Make the necessary changes and click Save.

Removing User Defined Device Types

1. Select one or more tag names and click Remove selected.
2. Click Yes to confirm removal of the selected device type.

Importing Device Types

You can import the device types using an XML file using the Import feature.
The XML file must be a valid XML file that has been exported using the **Device Types > Export** feature (see Exporting Device Types). Once uploaded, the file is validated to check for any conflicts before importing the device types.

The maximum number of device types that can be created is 1500.

Note: While adding or importing a user-defined device type, `.*` is automatically added at the beginning of the regular expression for the device type. For example, if the regular expression is `abc`, the expression saved in the system is `.*abc`. If the regular expression is `abc.*`, the expression saved in the system is `.*abc.*`

Procedure

1. Click **Management > Device Types**, and click the **Import** button.

   The **Import Device Types** wizard appears.

2. Click **Next** to continue. Follow the instructions displayed in the wizard.
3. Click **Browse** to locate and import a Device Types File. This file must be a valid XML file.
4. Click **Next** to continue.

   The Pre-Import Summary window appears. This window displays the summary of imported Device Types.

   - If there are no conflicts detected for the imported Device Types, click **Import** to add the imported Device Types to the Appliance.
   - If there are any Message Patterns, the Validate Message Patterns window appears. Review the message patterns for device types that you have selected and resolve any name conflicts that may have been detected by changing the Name. Double-click the row that you want to modify in Name column, the field becomes editable. Alternatively, you can click the button (located next to the conflicted row) to use the locally defined tag, and the << button to revert to the imported (but you may need to change the name) tag. Make the necessary changes. Make sure to use a unique name to avoid message pattern conflicts.

Exporting Device Types

Procedure

1. Click **Management > Device Types**, and click the **Export** button.

   The Export Device Types window appears.

2. Select a device type from the Available Device Type list and click the **Select** button to move entries from the **Available Device Types** pane to the **Selected Device Types** pane.
3. Remove any selected Device Types by clicking the Device Types name from the Selected Device Types pane, and then click **Remove**.
4. Click **Export** to export the Device Types and save the file on your desktop.

Check Point Log Sources

To collect log data from Check Point devices, you must set up LEA servers to foster log collection under **Management > Check Point Configuration**.

Log Export API (LEA) is used to retrieve and export VPN-1/ FireWall-1 Log data. Check Point Management Interface (CPMI) is used to provide a secure interface to the Check Point management server's databases.

For more information:

For more information about	See this documentation
LogLogic support of Check Point	*TIBCO LogLogic® Log Source Packages Log Configuration Guide for Check Point Management Station*
Check Point LEA servers and CPMI protocols	Check Point documentation

Management of Check Point Log Sources

You define the LEA server and CPMI protocols using the **LEA Servers** tab.

If the firewall or interface for the LEA server is on a different Check Point log source than the LEA server, you must specify it using the Firewalls or Interfaces tabs. The Firewalls or Interfaces tabs are accessible only after you add at least one LEA server to the appliance.

The LEA Servers tab lists the LEA servers defined on the appliance. Using this tab you can:

- Add new LEA servers
- Modify existing LEA servers
- Delete existing LEA servers
- View LEA, CPMI, and LEA server status

- Start or stop LEA servers
- Manually propagate LEA server definitions downstream (all new and updated LEA servers are automatically propagated after their properties are set)

The Firewalls or Interfaces tabs similarly let you add, modify, delete, and view firewalls and interfaces.

When modifying an LEA server, firewall, or interface, you have access to the same parameters and options. Using this tab, you can perform the following:

- To add a new LEA server to the appliance, click **Add New**. The Add LEA Server tab appears. For more information, see Adding an LEA Server.
- To modify an existing LEA server on the appliance, click the server's **Name**. Make the necessary changes using the **Modify LEA Server** screen and click **Update**.
- To remove an LEA server from the appliance, check the server's check box and then click **Remove**.
- To start an LEA server, in its row click  .
- To stop a running LEA server, in its row click  .
- To refresh the LEA Servers tab, click **Refresh**.
- To manually propagate LEA server definitions to downstream syslog receivers, click  .

Propagation of LEA Server Definitions

Definitions are automatically propagated whenever you add or update an LEA server.

For example, you can propagate information from LogLogic ST Appliance to LogLogic LX Appliance, LogLogic LX Appliance to LogLogic LX Appliance, or LogLogic LX Appliance to LogLogic ST Appliances. The  icon appears only if you add at least one LEA server.

Before you can enable this feature, you must perform the following tasks:

- Allow access to TCP port 5514. Use **Administration > Firewall Settings** to configure your ports.
- Verify that at least one appliance in the **Administration > Message Routing** tab

exists on your appliances.

Adding an LEA Server

To collect log data from a Check Point log source, you must define an LEA server on the appliance.
You can define an LEA Server on the appliance from **Management > Check Point > LEA Servers**. This lets you collect log data from that Check Point log source.

If the firewall or interface for this LEA server is on a separate Check Point log source, use the Firewalls or Interfaces tabs instead of the *Add Firewalls & Interfaces* section in step 7.

Procedure

1. Type the **Name** for the LEA server.
2. Select an **Agent Mode** to define how the LEA server starts. The default is **Automatic**, to ensure that the Check Point connection is established during system boot up.
3. Make sure that **Enable Data Collection** is set to **Yes**.
4. (Optional) Type a **Description** for the LEA server.
5. Establish Secure Internal Communication (SIC):
    a. Select the **Establish Secure Internal Communication** check box.
    b. Enter the Check Point server **SIC IP** address.
    c. Enter the **Activation Key** for the OPSEC Application on the Check Point log source.
    d. Enter the **OPSEC Application Name** for the application on the Check Point log source.
    e. Set up the SSL connection to the LEA server:
    f. Select the **SSL Connection to LEA Server** check box to enable it.
    g. Type the **LEA IP** address for the LEA server.
    h. Type the **LEA Port** number for the LEA server.
    i. Type the **LEA Server DN** (domain name).
6. If the firewall and interface are on the same Check Point log source as the LEA server,

configure them.

If they are on separate Check Point log sources, after adding this LEA server, use the Firewalls or Interfaces tabs instead.

a. Select the appropriate **Add Firewalls & Interfaces** radio button:
   - **CPMI Auto Discovery** - Automatically detects any Check Point Management Interface (CPMI) log sources connected to your system.
   - **Manual Input** - Lets you manually input each CPMI log source

b. Type the **CPMI IP** address.

c. Type the **CPMI Port** number.

d. Type the **Check Point User Name**. You must create an Administrator account in your Check Point application before you can use that ID for the **Check Point User Name** field on the LogLogic appliance.

e. Type the **Check Point User Password**. You must create an Administrator account in your Check Point application before you can use that password for the **Check Point User Password** field on the LogLogic appliance.

f. Select **SSL Connection to CPMI Server** to enable the SSL connection to your CPMI server.

g. Type the **CPMI Server DN** (domain name).

7. Click **Add** to add the LEA server. The new server definition is automatically propagated to the downstream syslog receivers.

Adding a Separate LEA Firewall

To collect firewall log data from a Check Point log source, define a firewall to associate with an LEA server defined on the appliance.
The **Add LEA Firewall** tab lets you define a firewall to associate with an LEA server defined on the appliance. This lets you collect firewall log data from that Check Point log source.

If the firewall is on a separate Check Point log source from the LEA server, use the Add LEA Firewall tab. If the firewall is on the same Check Point log source as the LEA server, you would have defined the firewall in the **Add Firewalls & Interfaces** section while adding the LEA server.

- To add a new LEA Firewall to the appliance, click **Add New**. The Add LEA Firewall tab

appears. For more information, see Adding an LEA Server.

- To modify an existing LEA Firewall on the appliance, click the firewall's Name. Make the necessary changes using the Modify LEA Firewall screen and click **Update**.
- To remove Firewalls from the appliance, check the firewall name's check box and then click **Remove**.

Procedure

1. Select an LEA Server from the list to associate with the firewall.
2. Type a **Name** for the firewall.
3. Type a **Description** for the firewall.
4. Select the **Yes** radio button to **Enable Data Collection**.
5. Click **Add** to add the firewall.

Adding a Separate LEA Interface

The interface for an LEA server is the actual log source for the Check Point system. The **Add LEA Interface** tab lets you define an interface for an LEA server defined on the appliance. This interface is the actual log source for the Check Point system, and the interface IP address appears as the origin in LEA messages.

Complete the configuration options listed under. If the interface is on a separate Check Point log source from the LEA server, use this **Add LEA Interface** tab. If the interface is on the same Check Point log source as the LEA server, you would have defined the interface in the **Add Firewalls & Interfaces** section while adding the LEA server.

- To add a new LEA Interface to the appliance, click **Add New**. The **Add LEA Interface** tab appears.
- To modify an existing LEA Interface on the appliance, click the firewall's Name. Make the necessary changes using the Modify LEA Interface screen and click **Update**.
- To remove Interfaces from the appliance, check the interface’s check box and then click **Remove**.

Procedure

1. Select an LEA Server to associate with the interface.

2. Select a firewall to associate with the interface.
3. Type a **Name** for the interface.
4. Type the **Interface IP address**.
5. Type the **Interface IP mask**.
6. For **Enable**, indicate whether to activate the interface. The default is Yes.
7. For **Trusted**, indicate whether to flag the interface as secure. The default is No.
8. For **Log Origin**, indicate whether the interface is the origin of the log message. The default is No. Typically the origin is the interface that is connected to the Check Point Management Station.
9. (Optional) Type a **Description** for the interface.
10. Click **Add** to add the interface.

Column Manager - Overview

Using the **Column Manager** menu you can define which columns to hide from the Searches and Reports when the Data Privacy mode is enabled.

To enable the Data Privacy Mode, navigate to the **Administration > System Settings > General tab > Data Privacy Options** and click the **On** radio button.

For more information, see Data Privacy Settings.

Accessing the Column Manager

You can define which columns to hide from the Searches and Reports, when the Data Privacy Mode is enabled.

Procedure

1. Navigate to **Management > Column Manager**.

   The **Column Manager** window appears. By default, it does not display any information. You can filter the column list by entering the column name in the **Find** field, and press **Enter**. The filtered column list containing the search term is displayed.

2. Click the down arrow next to the **Filters** field, to display advanced filtering options.
3. Select a filter from the list. You can filter based on:
   - Feature
   - Reports Category
   - Report Name
   - Device Type

Result

The following information is displayed:

- Column Name: Name of the column

- Hidden in Data Privacy Mode: Displays ☑ for the hidden columns

Hiding Columns

You can define which columns to hide from the Searches and Reports when the Data Privacy Mode is enabled.

Procedure

1. Access **Management > Column Manager** from the navigation menu.

   The **Column Manager** window appears.

2. Select the **Column Name** checkbox next to the column name that you want to select. You can select multiple columns at a time.

3. Click the **Hide** button to hide the selected columns.

Result

The ☑ icon is displayed in the **Hidden in Data Privacy Mode** column and those columns are hidden from the Search or Report.

Showing Columns

You can define which columns to display in the Searches and Reports when the Data Privacy Mode is enabled.

Procedure

1. Access **Management > Column Manager** from the navigation menu.

   The **Column Manager** window appears.

2. Select the **Column Name** checkbox next to the column name that you want to select. You can select multiple columns at a time.

3. Click the **Show** button to display the selected columns.

Result

The  icon is removed from the **Hidden in Data Privacy Mode** column and those columns are displayed in the Search or Report.

Exporting a Configuration File

Use the Export feature to export the Column Manager settings into an XML file format that can be used for another appliance.
When using Management Station, you can export the Column Manager settings that you configured on one appliance and import them for use on another appliance. This saves you from configuring the settings again on the other appliance.

Procedure

1. Access **Management > Column Manager** from the navigation menu.

   The **Column Manager** window appears.

2. Click the **Export** button to export the configuration file.
3. In the File Download dialog box, specify where in your file structure the downloaded file should be saved. The Exported files are in XML format.

Importing a Configuration File

Use the Import feature to import the Column Manager settings from one appliance to another.

When using Management Station, you can export the Column Manager settings that you configured on one appliance and import them for use on another appliance. This saves you from configuring the settings again on the other appliance.

Procedure

1. Access **Management > Column Manager** from the navigation menu.
2. The **Column Manager** window appears.
3. Click the **Import** button. The **Import Column Manger Configuration** window

appears.

4. Click the **Browse** button.
5. Click the file name to specify a file. The selected file appears in the **File Name** field. Click **Open**.

   Warning: The files must be in valid XML format. Attempting to import other formats results in an error message.

6. Click **Import** to import the Column Manger settings from the exported file.

   The Column Manger window appears showing the exported file settings. You can change the settings at any time as explained in Hiding Columns or Showing Columns.

Generating a Reports Summary

The summary report summarizes the reports that can be generated in the system based on the user selected filters.
Use the Generate Reports Summary feature to generate a summary report. Reports that have different definitions for different device types are shown separately with the applicable device types.

Procedure

1. Access **Management > Column Manager** from the navigation menu.

   The **Column Manager** window appears.

2. Click the **Generate Reports Summary** button to create a summary report.

   The dialog box that appears lets you open or save a report file.

Management of PIX/ASA Message Codes

Use the **PIX/ASA Message Codes** tab to categorize each incoming message based on the PIX/ASA severity and message code combination. The tab is available only on LogLogic LX Appliances and LogLogic MX Appliances.

For example, each incoming message of severity 2 with a 106006 message code is by default stored in the deny table. You can use the PIX/ASA Message Codes tab to change the category for the message.

Messages marked as Off on **PIX/ASA Messages Codes** tab are inserted into RawSyslog and stored. Disabled radio buttons indicate that the category is not applicable for the particular message type.

Note: The PIX/ASA Messages feature is available only on LogLogic LX Appliances and LogLogic MX Appliances. Only those appliances parse each PIX or ASA message.

Each Cisco message has a corresponding Cisco PIX/ASA message code. Your appliance uses this message code to help determine the message type for an incoming message.

The first number in the message code is the PIX/ASA severity level and the second number is the message code. For example, 1-101002, means the severity level is 1 and the message code is 101002.

The categories to define each PIX/ASA message include:

Category	Description
Off	Ignores the message
System	Categorizes the message in the **System** table
Security	Categorizes the message in the **Security** table
Parsed	Categorizes the message in the **Parsed** table

Each supported message code is listed in the **Message Code** column on the **Administration > PIX/ASA Messages Codes** tab.

To view the definition for each message code, click the hyperlinked message code. Use the **Reset** button for resetting the codes to default installation settings.

The appliance automatically sets the message code to the default setting. However, you can click any active radio button to change the setting. .

> ⚠ **Warning:** Changing default message code settings can result in a large amount of disk space consumption that can shorten your log retention cycles.

The following codes are disabled (off) by default because they are redundant, though you can enable any of them if you want:

7-109014	7-109021	7-111009	7-199009	7-304005
7-701001	7-701002	7-702301	7-702303	7-703001
7-703002	7-709001	7-709002	7-710006	6-305001
6-305002	6-305003	6-305004	6-305009	6-305010
6-305011	6-305012	6-611301	6-611302	6-611303
6-611304	6-611305	6-611306	6-611307	6-611308
6-611309	6-611312	6-611314	6-611315	6-611316
6-611317	6-611318	6-611319	6-611320	6-611321
6-611322	6-611323	6-613002	6-614001	6-614002
3-106014	3-305005	3-305006		

Enabling Cisco PIX/ASA Message Codes

You can enable Cisco PIX/ASA Message Codes from the **PIX/ASA Messages Codes** tab.

Procedure

1. Click **Administration > PIX/ASA Messages Codes**.

   The **PIX/ASA Messages Codes** tab is displayed.
2. Select **All** from the **PIX/ASA Severity** list.
3. Scroll down the **Message Code** column until you see the specific message code.
4. Select the radio button for a different category for your specific message code.
5. Click **Update** to save your change(s).

Mapping Cisco Log Source Names to IP Addresses

The LogLogic appliance identifies log sources by their IP addresses.
Some Cisco logs do not contain an explicit IP address but a DNS-type name instead. If you set up a special configuration file, the appliance can recognize these names and replace them with IP addresses. The effects of this can be seen in a variety of places throughout the GUI including, for example, the **Source IP** and **Destination IP** columns in Active FW Connections reports.

The appliance gets its name recognition information from a configuration file that you need to configure and upload to the appliance.

Procedure

1. On the Cisco log source, locate the generated Cisco IP mapping.
2. Every Cisco system can generate such a mapping file. For more information, see your Cisco documentation.
3. In that file, search for a large number of entries of the form:

   ```
   name 10.20.50.51 remote.lan
   name 10.0.25.51 async.wan
   name 10.19.50.10 nemesis-ss1-vs
   name 10.19.83.1 pwddbl0c-9
   ```
4. Copy and paste all the entries into a text file called `pix_name_ip_map.txt`.

5. Copy the file (using SCP) onto the LogLogic appliance, to the directory `/loglogic/conf`.

   After placing this file on the LogLogic appliance, all report results containing log data from Cisco log sources which originally did not have the correct IP addresses (because they could not recognize the names) now have them.

Management of Port Descriptions

Use the **Port Description** tab to view information about the ports on the appliance. The tab is available only on LogLogic LX Appliances and LogLogic MX Appliances.

Different ports are used depending on the application. LogLogic LX Appliances and MX appliances let you add a description of these ports for display in reports. For various Real-Time reports, the appliance lets a description display with Source (SRC) or Destination Port fields. This provides more detail information about the specific Real-Time reports, because you can have custom applications for different ports. The industry standard definitions are included by default.

Note: The Port Description feature is available only on LogLogic LX Appliances and LogLogic MX Appliances. Real-time reports are available only on those appliances.

You can access the **Port Descriptions** tab from **Administration > Port Descriptions** on any LogLogic LX Appliances or LogLogic MX Appliances.

Adding Ports

The **Add Port Description** tab adds a port description to the appliance port registry.

Procedure

1. Select the **Administration > Port Descriptions** from the navigation menu**.**
2. Click **Add New**.
3. In the **Port** field, enter a port number for the source application.
4. In the **Protocol** list, choose the protocol for the source application.
5. In the **Description** field, enter a description for the source port.

Warning: Do not use the < > & “' characters in the **Description** field.

6. Click **Add** to add the new port.

Modifying Ports

The **Modify Port Description** tab modifies a port description in the appliance port registry.

Procedure

1. In the **Administration > Port Descriptions > Port** column, click the hyperlinked port number of the port you want to modify.

   The **Modify Port Description** tab displays.
2. In the **Port** field, enter a port number for the source application.
3. In the **Protocol** list, choose the protocol for the source application.
4. In the **Description** field, enter a description for the source port.

   > ⚠ **Warning:** Do not use the < > & “' characters in the **Description** field.
5. Click **Update** to save your changes.

Removing Ports

The **Remove Port Description** tab deletes a port description from the Appliance port registry.

Procedure

1. In the **Administration > Port Descriptions > Port** column, select the check boxes next to the hyperlinked port numbers of the ports you want to delete.

   The **Remove Port Description** tab is displayed.
2. Click **Confirm Remove** to delete the ports.

Monitoring Console

The Hawk Console is available in LogLogic LMI as the Monitoring Console. It provides a central view of all the distributed components interacting within the Hawk® environment.

It is easier to manage multiple domains from LogLogic LMI than configuring and controlling the domains individually. From the Monitoring Console, you can:

- Configure and manage domains of TIBCO Operational Intelligence Hawk® 6.2.1 HF2 or later
- Monitor and manage distributed applications and operating systems
- Manage alerts generated by Hawk agents
- Take action in response to predefined conditions

By default, the Monitoring Console is disabled in LogLogic LMI. After a user with administrator access enables it, you can access it from **Monitoring > Console**.

Domains

Monitoring Console provides a ready-to-use Hawk domain named `lmi_domain`, to which you can connect external Hawk agents. By default, `lmi_domain` is registered to the Monitoring Console. You can unregister it by clicking the Unregister icon on the domain card. To register it again, follow the procedure described in Configuring a Hawk domain.

You can also register other external domains and external Hawk agents can connect to the domains.

Related Topics

For more information, see the following documentation:

Topic	Reference
TIBCO Hawk Console	TIBCO Hawk Console User's Guide

Topic	Reference
TIBCO Operational Intelligence Hawk®	TIBCO Operational Intelligence Hawk® documentation
Enabling the Monitoring Console	Enabling the Monitoring Console
Configuring a Hawk domain within LogLogic LMI	Configuring a Hawk domain
Securing the communication between the Monitoring Console and other Hawk components	Adding Hawk Certificates to LogLogic LMI

Enabling the Monitoring Console

By default, the Monitoring Console is disabled in LogLogic LMI. As an administrator, you can enable the Monitoring Console so that other users can access it.
You can enable the Monitoring Console from the GUI or the CLI. To enable the Monitoring Console from the CLI, see system Command.

To enable the Monitoring Console from the GUI, perform the following steps:

Procedure

1. Go to **Administration > System Settings**.
2. On the **General** tab, in the **Advanced Features Setting** section, enable the following options in this sequence:
    a. **Advanced Features**
    b. **Monitoring Console**

What to do next

As an administrator, you can perform the following tasks after enabling the Monitoring Console:

- To configure a Hawk domain, see Configuring a Hawk domain.
- To change the default advanced settings such as number of alerts to be retained or alert storage options, see Configuration of the Monitoring Console.

Configuring a Hawk domain

You can register and configure multiple Hawk domains within LogLogic LMI. The Monitoring Console connects to Hawk domains using TCP.

In addition to the default `lmi_domain`, you can configure and register other external domains in LogLogic LMI and external Hawk agents can connect to the domains.

If you have deployed LogLogic EVA on a Docker host, you must use the `proxy` domain type to connect to existing Hawk domains.

Domain type	Description
regular	An external Hawk domain (created outside of LogLogic LMI), to be connected over TCP. Provide a domain name and its TCP transport details.
proxy	An external Hawk domain (created outside of the LogLogic LMI), to be connected using the proxy method. Provide a domain name, URL, and login credentials.
LMI Domain	An internal domain, provided as a built-in, ready-to-use domain. External agents can connect securely to this domain using the `Self/Cluster Manager Host:Port` field.
New Domain	An internal domain, similar to the built-in LMI Domain External agents can connect securely to this domain using the `Self/Cluster Manager Host:Port` field.

Procedure

1. On the Monitoring Console, in the **Domain** section click the plus (+) icon.
2. In the **Configure Domain** dialog box, provide the following information:

   For details about these parameters, see TIBCO Hawk® Console User's Guide.

Field	Description	Default Value
Domain type	Choose one of the following domain types: • proxy • regular • New Domain • LMI Domain **Note:** LMI Domain is the default domain. If you select LMI Domain, all other fields are automatically set to their default values. LMI Domain is available in the list only if it is currently unregistered.	regular
Domain Name	Type a name for the domain	Automatically set to lmi_domain (if you select LMI Domain as the domain type)
Transport	The transport protocol **Note:** Not applicable to a proxy domain.	TCP
• Self Host: Port • Cluster Manager Host: Port	• **Self Host: Port**: Enter the IP address and port of the LogLogic LMI appliance • **Cluster Manager Host: Port**: Enter the IP address and port of the cluster manager	None

Field	Description	Default Value
	Note: These two fields are available when you select the domain type as regular.	
Self/Cluster Manager Host: Port	Enter the IP address and port of the appliance. The port number must be unique and open for bidirectional communication.	• If you select `LMI Domain` as the domain type, the value is set to `<appliance_IP_address>:9688` and the field is not editable. • If you select `New domain` as the domain type, you must enter the appliance IP address and port.
	If you select `LMI Domain` or `New domain` as the domain type, the two fields **Self Host: Port** and **Cluster Manager Host: Port** are combined as the **Self/Cluster Manager Host: Port** field.	

3. To configure a domain that uses SSL-based TCP transport, select the **Additional transport options** check box and provide the values for the following fields:

Field	Description
Key store	Absolute path of the key store that contains the Monitoring Console certificate and key to be loaded while communicating with the Hawk domain. You must provide a custom keystore that has a password-protected key. For example, the default keystore of LogLogic LMI is: `/loglogic/tomcat/conf/keystore`
Key store	Password to access the key store

Field	Description
password	
Key password	Password to access the private key
Trust store	Absolute path to the trust store to be used to validate the Hawk component certificates while communicating with the Hawk domain. For example, the default trust store of LogLogic LMI is: `/loglogic/tomcat/conf/truststore`
Trust store password	Password to access the trust store
SSL protocol (optional)	TLSv1.3 and TLSv1.2 protocols are supported.
SSL Enabled Algorithms (optional)	Comma-separated list of algorithms. Default value: `TLS_RSA_WITH_AES_128_CBC_SHA`
Security Policy	Choose one of the following: • None • Trusted • Trusted with Domains (only for Windows domains) • Custom

4. Click **Configure**.

Result

The domain is configured and displayed in the Domains section.

Note: Whenever the connection of the appliance with the domain is lost, the appliance tries to reconnect. However, if the appliance is not able to reconnect, a reconnect button appears on the Monitoring Console. You can click the button to manually reconnect to the domain.

What to do next

After configuring a Hawk domain, you can connect Hawk agents to it.

Note: When configuring Hawk agents, ensure that you use the domain name, IP address, and port number of the domain in LogLogic LMI.

For information about how to connect Hawk agents to the Hawk domain, contact your Hawk administrator or see TIBCO Operational Intelligence Hawk® documentation.

Adding Hawk Certificates to LogLogic LMI

From **Administration > SSL Certificate > Trusted Certificate**, you can use LogLogic LMI certificates to establish a secure communication between the Hawk components and the Monitoring Console.
Hawk components use mutual TLS authentication for communication with each other. After authenticating the certificates, an encrypted channel is established between the Hawk components.

If the Hawk components are already configured to use the different CA certificates, you must add them to the LogLogic LMI trust store. There is no need to restart the LogLogic LMI engines after enabling SSL configuration for the Hawk domain.

To add the Hawk CA certificates to the default LogLogic LMI trust store, perform the following steps on the LogLogic LMI GUI:

Procedure

1. Go to **Administration > SSL Certificate > Trusted Certificate**.
2. Copy the Hawk component certificate and paste it in the **Import Trusted Certificate** text box.
3. Click **Import**.

Configuration of the Monitoring Console

As an administrator, you can change the configuration of the Monitoring Console.

In the Monitoring Console configuration files, you can modify advanced settings such as defining the number of alerts that can be retained by default, or the storage options. The values in the following sample files are the default values stored in the LogLogic LMI appliance. You can modify the values as required and then perform the steps described in Configuring Advanced Features.

Sample Configuration File for Monitoring Console Parameters

The following sample file includes the Monitoring Console settings:

```
{

  "configurations": [
    {
      "nodePath":
"/unity/system/config/hawkconsolenodes/hawkconsolenode-0000000000",
      "data": {
        "services": {
          "rest": {
            "host": "0.0.0.0",
            "port": 9687
          }
        },
        "alerts": {
          "notificationRetentionCount": 100000,
          "lowRetentionCount": 100000,
          "mediumRetentionCount": 100000,
          "highRetentionCount": 100000
        },
        "domainConfig": {
          "defaultDomainName": "lmi_domain",
          "defaultClusterPort": 9688,
          "domainTransportConfigPath":
"conf/DomainTransportConfigPath.yml",
          "repositoryPath": "conf/repository"
        },
        "storage": {
          "cache": "/loglogic/data/.hawkconsole/hcache",
```

```
          "type": "mysql",
          "host": "127.0.0.1",
          "port": 3306
        },
        "version": 1
      }
    }
  ]
}
```

Parameter Description	Default value
Maximum number of alerts	100,000 alerts `"alerts": { "notificationRetentionCount": 100000, "lowRetentionCount": 100000, "mediumRetentionCount": 100000, "highRetentionCount": 100000 }`
Configuration of the appliance that hosts the Monitoring Console	`"domainConfig": { "defaultDomainName": "lmi_domain", "defaultClusterPort": 9688, "domainTransportConfigPath": "conf/DomainTransportConfigPath.yml", "repositoryPath": "conf/repository" }`
Storage medium of alerts	MySQL database, its IP address, and port `"storage": { "type": "mysql", "host": "127.0.0.1", "port": 3306 }`

User Access Control

As an administrator, you can control the users' access to domains and agents in the Monitoring Console, by specifying the access rights in a configuration file.

The configuration file is stored at the following location on the LogLogic LMI appliance:

```
/loglogic/logu/hawkconsolenode/conf/monitoring_console_user_access_
list.cfg
```

In this file, you must specify the access rights for each user on a separate line. Similarly, if the same user has multiple access rights, then you must specify each access right on a separate line. The access rights must be specified in the following format:

```
<username> <access_rights>
```

When specifying the access rights, consider the following points:

- An exclamation mark (!) at the beginning of the row indicates an access restriction, whereas omitting the exclamation mark (!) indicates an access permission.
- `<username>`: This is the LogLogic LMI user. To indicate all users, you can use asterisk (*) as a wildcard.
- `<access_rights>`: You can specify one or more of the values - agent name, agent DNS, and domain name - as required.
    - If you specify multiple values, then the values must be in this sequence: agent name, agent DNS, and domain name. These values must be separated either by a space or colon (:), and must be enclosed in double quotes (" ").

      For example: `"agent1 DNS1 domain1"` or `"agent1:DNS1:domain1"`
    - To indicate all agents, or all DNS, or all domains, you can use asterisk (*) as a wildcard.

For sample scenarios and the corresponding entries in the configuration file, see the Examples section.

Procedure

1. Edit the `monitoring_console_user_access_list.cfg` file to add the permissions and restrictions.

2. Restart Monitoring Console by running the following command:

```
mtask -s engine_lldaemon restart
```

Examples

Consider that you want to configure the access control scenarios for admin, user-1, and user-2. The corresponding entries for those access rights are listed in the following table.

User	Required access rights	Entry in the `.cfg` file
admin	Allow access to all domains	`admin  *`
user-1	Allow access to `agent-1` in `lmi_domain` with `dns-1`, and to all agents in `lmi_domain-2`	`user-1  "agent-1 dns-1 lmi_domain"` `user-1  "* * lmi_domain-2"`
user-2	Allow access to `agent-2` of any domain, but no access to the `lmi_domain-2` domain	`user-2  agent-2` `!user-2 "* * lmi_domain-2"`

Therefore, make the following entries in the `monitoring_console_user_access_list.cfg` file:

```
admin   *
user-1  "agent-1 dns-1 lmi_domain"
user-1  "* * lmi_domain-2"
user-2  agent-2
!user-2 "* * lmi_domain-2"
```

For detailed information about user-based access control, see TIBCO® Operational Intelligence Hawk® RedTail 7.1.0 Installation, Configuration, and Administration.

Limitations

- The access control list is applicable to Hawk domains that conform to the `Trusted` security policy.
    - For information about domain settings, see Configuring a Hawk domain.
    - For information about security policies, see the "Authorization at the REST API Layer" section in *TIBCO® Operational Intelligence Hawk® RedTail Installation, Configuration, and Administration*
- In the classic Hawk console of TIBCO® Operational Intelligence Hawk® RedTail, you can use this mechanism to control access up to the node microagent and node methods; but not in LogLogic LMI. In the `monitoring_console_user_access_list.cfg` file of LogLogic LMI Monitoring Console, even if you specify access rights at the microagent or method levels, those lines in the file are ignored.

File Transfer History

The **File Transfer History** displays the transfer status for File Transfer Devices in your system.

View the status for individual devices or rules from **Administration > File Transfer History** tab.

The **File Transfer History** button appears in the following locations:

- **Management > Devices** tab, click a *Device-Name* **> Modify Device** tab
- **Management > Devices > File Transfer Rules**: click a *Rule-Name* **> Modify File Transfer Rule** tab

You can use it to view the file transfer history. To view the history, you must add at least one rule.

The appliance computes the checksum for each file and stores it in the database. Both SHA256 and MD5 are supported. The appliance uses the checksum to determine if a log file has been modified since it was last transferred. If the checksum remains the same, the file has not been changed, and therefore, is not processed by the appliance.

Column	Description
Retrieval Date	Date of file transfer
Filename	Name of the file transferred
Log Source	IP address of the device from which the file was transferred
Protocol	Protocol used to transfer the file(s)
Format	File format
Size	File size

Column	Description
MD5	MD5 of the file. This column is visible only when File digest is set to MD5.
SHA256 Digest	SHA256 digest of the file. This column is only visible when File digest is set to SHA256.

File Transfer: Date and Time Formats

LogLogic supports different file types, and date and time formats.

When a data file is transferred, each message or event contains a timestamp that consists of a date and time. The timestamp refers to the file creation date and time for a particular message in the file. When the file is received and the appliance cannot determine a date and time stamp, the date and time of retrieval is used. The files are then parsed and stored.

LogLogic supports the following file types, and date and time formats:

- Blue Coat ProxySG
- Cisco ACS
- Generic W3C
- Microsoft IAS
- Microsoft ISA
- NetApp NetCache
- Other File Devices
- RSA ACE Server
- Squid

Blue Coat ProxySG

The date and time formats for Blue Coat ProxySG are:

- `YYYY-MM-DD hh:mm:ss`
- `DD/MM/YYYY:hh:mm:ss`
- `epoch.optional_millisecs`

Example: Date/Time Format

```
172.20.56.78    anonymous       Mozilla/4.0 (compatible; MSIE 6.0;
Windows NT 5.0)
       2009-03-25      00:00:00        CO-NET-002      -       -       -
      -   - 857     -       -
       GET     http://ff.company_
name.com/selector/image?client=Bet365&placement=Livescore_Soccer_ROW_
120x60_GIF  -       12209
```

Cisco ACS

The date and time format for Cisco ACS is: `DD/MM/YYYY,hh:mm:ss`.

Example: Date/Time Format

```
03/11/2009,12:23:34,,,,NAS Reset,0,,,,,,,,,123.234.23.34,
```

The first part is the date and the second part is the time.

Generic W3C

The date and time format for Generic W3C is: `YYYY-MM-DD hh:mm:ss`.

Example: Date/Time Format

For this date and time format, there is usually a header which indicates the date and time column.

```
#Software: Microsoft(R) Internet Security and Acceleration Server 2000
#Version: 1.0
#Date: 2009-03-25 00:00:00
#Fields: c-ip   cs-username     c-agent date    time s-computername
cs-referred     r-host  r-ip    r-port time-taken      cs-bytes
sc-bytes       cs-protocol s-operation     cs-uri   s-object-source sc-
status 172.20.56.78    anonymous       Mozilla/4.0 (compatible; MSIE
6.0; Windows NT 5.0)      2009-03-25      00:00:00        CO-NET-002 -
     -       -      --       857     -       -       GET
http://ff.connextra.com/selector/image?
client=Bet365&placement=Livescore_Soccer_ROW_120x60_GIF  -       12209
```

In this example, the date is `2009-03-25` and the time is `00:00:00`. The date and time column can be in different places, so parse the header row to find the location.

It can also be x-timestamp, for example:

```
#Software: NetCache NetApp/5.5R6D18
#Version: 1.0
#Start-Date: 2009-03-15 06:24:09
#Remark: http
#Fields: x-timestamp time-taken c-ip x-transaction bytes cs-method cs-
uri x-username x-hiercode rs(Content-Type) x-note 1237075200.157 0.001
192.168.64.61 TCP_HIT/200 0 GET http://www.movievoyager.com/ - -
"text/html" -
```

In this example, x-timestamp is in the epoch time and milliseconds format. The millisecond is optional so it might not appear.

Microsoft IAS

The date and time formats for Microsoft IAS are:

- `DM/DM/YYYY,hh:mm:ss`
- `DM-DM-YYYY,hh:mm:ss`

  `-DM`, the date or month is automatically detected; the default is to use the USA date

format

Example: Date/Time Format

```
"CLIENTCOMP","IAS",03/17/2009,13:04:

33,1,"client",,,,,,,,9,"10.10.10.10","iasclient",,,,,,,1,,
0,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,
```

Microsoft ISA

The date and time formats for Microsoft ISA are:

- `YYYY-MM-DD,hh:mm:ss`

```
172.20.56.78    anonymous       Mozilla/4.0 (compatible; MSIE 6.0;
Windows NT 5.0) 2009-03-25      00:00:00        CO-NET-002      -
    -       -       -   - 857     -       - GET
http://ff.company_
name.com/selector/image?client=Bet365&placement=Livescore_Soccer_
ROW_120x60_GIF  -       12209
```

- `YYYY-MM-DD,hh:mm:ss`

```
172.20.56.78    anonymous       Mozilla/4.0 (compatible; MSIE 6.0;
Windows NT 5.0) 2009/03/25      00:00:00        CO-NET-002      -
    -       -       -   - 857     -       -GET
http://ff.company_
name.com/selector/image?client=Bet365&placement=Livescore_Soccer_
ROW_120x60_GIF  -       12209
```

NetApp NetCache

The date and time formats for NetApp NetCache are:

- `epoch.optional_millisecs`
- `DD/Mmm/YYYY:hh:mm:ss`

Example: Date/Time Format

```
#Software: NetCache NetApp/5.5R6D18
#Version: 1.0
#Start-Date: 2009-03-15 06:24:09
#Remark: http
#Fields: x-timestamp time-taken c-ip x-transaction bytes cs-method cs-
uri x-username x-hiercode rs(Content-Type) x-note 1237075200.157 0.001
192.168.64.61 TCP_HIT/200 0 GET http://www.movievoyager.com/ - -
"text/html" -
```

Other File Devices

The date and time formats for Other File Devices are:

- `YYYY-MM-DD hh:mm:ss[.ms|,ms]` where `ms` is milliseconds)

  ```
  2009-03-31 00:00:00 128.206.230.55 - 216.106.89.10 80 GET
  /fireworks_files/index_navbar/index_navbar_r12_c3.gif 2009-03-25
    00:00:00        CO-NET-002      -       -       -        -   -
  857     -       -
  ```

- `Mmm DD hh:mm:ss`

  ```
  Mar 11 10:36:21 hq-ace.net.com ACESERVER: [ID 302638 user.warning]
  (1) AUTHENTICATION : ACCESS DENIED, syntax error (Login:'binhn';
  User Name:''; Token:'---->'; Group:''; Site:'
  ```

- `YYYY-MM-DDThh:mm:ss[.ms]` (`ms` is milliseconds)

  ```
  2009-03-25T01:02:03  CO-NET-002  -       -       -       -   - 857
  -       -GET      http://ff.company_
  name.com/selector/image?client=Bet365&placement=Livescore_Soccer_
  ROW_120x60_GIF   -       12209
  ```

- `MM/DM/YYYY hh:mm:ss[.ms]` (`ms` is milliseconds)

  ```
  "CLIENTCOMP","IAS",03/17/2009
  13:04:33,2,,"iasclientdc/Users/client",,,,,,,,9,"10.10.10.10",
  ```

- DD-Mmm-YYYY hh:mm:ss|Mmm-DD-YYYY hh:mm:ss

```
209.36.37.147 - - [03-Mar-2009 01:00:02 +0000] "GET
http://dopzatphv.company_name.com/fraser/needam.gif HTTP/1.0" 304
83 209.36.37.147 - - [Mar-03-2009 01:00:02 +0000] "GET
http://dopzatphv.company_name.com/fraser/needam.gif HTTP/1.0" 304
83
```

- DD/Mmm/YYYY:hh:mm:ss

```
209.36.37.147 - - [03/Mar/2009:01:00:02 +0000] "GET
http://dopzatphv.company_name.com/fraser/needam.gif HTTP/1.0" 304
83
```

- YYYYMMDD hh:mm:ss

```
*   20090310 23:00:15 435ED4D70103125C DELIVER VOLUME=12246
MAILBOX=be71bf220765c750e1bbc0f0a987b934@tin.it *   20090310
23:00:15 435ED4D70103125C DELIVER
From=<e4b115b72192e59dd08b50c25f8bae6b@tin.it> Size=1881 *
20090310 23:00:15 435ED4D70103125C DELIVER
Recipient=<be71bf220765c750e1bbc0f0a987b934@tin.it> *   20090310
23:00:15 435ED4D70103125C DELIVER  Message-
ID=<a5118ecdc9f903c9e5707e9808883feb@vsmtp4.tin.it>
```

RSA ACE Server

The date and time format for RSA ACE Server is: `Mmm DD hh:mm:ss`.

Example: Date/Time Format

```
Mar 11 10:36:21 hq-ace.net.com ACESERVER: [ID 302638 user.warning] (1)
AUTHENTICATION : ACCESS DENIED, syntax error (Login:'binhn'; User
Na$me:''; Token:'---->'; Group:''; Site:''; Agent Host:'taz.net.com';
Server:'hq-ace.net.com')
```

Squid

The date and time format for Squid is: epoch.optional_milli_secs.

Example: Date/Time Format

```
1237075200.240 2 172.16.6.39 TCP_NC_MISS/200 3058 GET
http://172.16.6.202:8082/Secure/Local/console/bcsstyle.css -
DIRECT/172.16.6.202 text/html content_filter_not_applied
```

The Squid time and date format is similar to the NetCache NetApp format. In this example, the time is always the first column, epoch seconds plus milliseconds.

Forwarding Logs to Other Appliances (Routing)

Message routing lets you forward a copy of all incoming log messages to one or more destinations by creating an appropriate routing rule.

For example, an LogLogic ST Appliance receiving messages from multiple sources can be configured to forward some or all of the messages to an appliance based on one or more rules. The destination of the routing rule can be a host other than an LogLogic LMI appliance.

Configuring an Appliance for Message Routing

You must configure your appliance for the various forwarding protocols that TIBCO LogLogic® supports. TIBCO LogLogic® supports the following forwarding protocols per log type:

- Real-time logs can be forwarded using all except SNMP protocols.
- File-based logs, including database collection (such as MSSQL or Oracle), can be forwarded using all except SNMP protocols. If you use a protocol other than LogLogic TCP, the source type of the logs are detected as general syslog source on the downstream appliance. If the downstream appliance is an LogLogic LX Appliance model or MX model, the file-based logs sent using a protocol other than LogLogic TCP are not parsed. LogLogic TCP cannot be used for sending to a non LogLogic LMI host.
- SNMP can only be forwarded using SNMP protocol. (In addition, SNMP logs are translated to ASCII format and internally routed to the Syslog port. These translated SNMP logs are just like real-time logs, and can be forwarded using “All Sources” forwarding rules.)

> **Note:** LogLogic LMI does not support sending duplicate logs to the same destination when the same protocol is selected even with different ports. If you want to send the same set of data twice to the same destination, you must use different protocols.

Procedure

1. (Optional) Define Search Filters.
2. (Optional) Define Device Groups to be used as the routing rule source.
3. For the LogLogic TCP protocol, enable TCP port 5514 access on destination appliances for syslog sources and TCP port 4433 for file-based sources. See Network Access Control.

   When forwarding logs through TCP syslog, the log source is only correctly discovered on the downstream LogLogic LMI if syslog priority <N> is present at the beginning of the log. If syslog priority is not present, the log source is considered to be the upstream appliance. When forwarding is done through a secure tunnel, the downstream appliance considers logs to be sent by 127.0.0.1. Such logs are automatically assigned type of LogLogic LMI appliance.

Outbound Routing Rules

You can create a new routing rule to specify the source device (or device group) this rule applies to, the destinations to forward to, and the details of the communication pathway to the destination.

> **Note:** The LogLogic appliances forward logs through UDP and TCP syslog and SNMP protocols to other destinations. The logs forwarded include syslog messages, file-pulled logs, and SNMP traps. For file-pulled logs, the user can set the forwarding speed. The user can turn headers On or Off on a per-routing-rule basis for file-pulled logs and SNMP traps, and can set the forwarding speed for file-pulled logs.

> **Warning:** When using LogLogic TCP, the source and destination appliances must be of the same release and hotfix.

It is good practice to use the following for your Syslog-NG configuration to correctly collect logs:

```
template(“<$PRI>$R_DATE $SOURCEIP $MSG\n”) template_escape(no)
```

Note: If you enable the **Administration > System Settings > Auto-identify Log Sources** option and you have several thousand log sources configured that need to be auto-identified, routing rules and alerts can slow the auto-identify process.

You can create up to 200 routing rules for each appliance. However, you must account for several factors which can affect the number of rules your appliance can manage:

- message rate
- filter (use of regular expressions)
- tunneling
- authentication (authentication is a one-time occurrence)
- compression
- TCP transport

  LogLogic TCP should be used only when required, for example, over unreliable or slow WAN links or when file-based data must be kept in file format.
- number of searches or reports being executed on the appliance
- number of file-base transfer rules (which are not included in the inbound messager rate)
- number of alerts (especially those with regular expressions)
- whether HA is enabled

Note: The log sources specified in each rule have an impact on performance. For example, 10,000 sources in 3 routing rules having their aggregate data set sent to 3 hosts has an additional overhead as compared to 100 log sources having their aggregate data sent to 3 hosts.

Adding Destinations to the All Sources Rule

The All Sources routing rule (All Sources) forwards a copy of all incoming log messages to multiple destinations.

For example, a LogLogic ST Appliance receiving messages from multiple firewalls can be configured to forward all of the messages to one or more destinations as per the All Sources rule.

If you add a new log source, it is automatically added to this rule. You can add more destinations to the All Sources rule by using the **Add Destination** link. You can edit or remove the added destinations.

Procedure

1. Click the **Add Destination** link.

   The **Add Destination** window appears. By default, **LogLogic Forwarding Settings**, and **Other Settings** options are disabled. When you specify the **Destination Type** and/or Protocol, some of these options are enabled.

2. In the **Destination IP** field, type the IP address of the destination to which you want to forward messages.

   This can be another LogLogic appliance, a LogLogic Security Event Management (SEM) appliance, or another machine (with correct port configuration). This is a mandatory field.

3. In the **Destination Port** field, type the port number to which you want to forward messages.

4. From the **Destination Type** list, select where you want to forward messages:

   - LogLogic LMI Appliance
   - LogLogic SEM Appliance
   - Other Destination

5. From the **Protocol** list, select the protocol to use for forwarding messages:

Option	Description
UDP Syslog	Traditional syslog using the UDP protocol

Option	Description
	Note: The UDP Syslog protocol is not supported on LogLogic EVA - Container Edition.
TCP Syslog	Traditional syslog using the TCP protocol. Also known as Syslog-NG New-line (\n) characters are used to break logs in the TCP stream during message forwarding. If a message contains \n, the message breaks up with only the first portion of the message being delivered to the downstream appliance. It is good practice to select a different forwarding protocol if you know your log messages contain characters of this type.
LogLogic TCP	Buffered syslog provided by TIBCO LogLogic®. Uses a proprietary TCP-based protocol and uploads logs in batches every minute **Note:** If you select **LogLogic TCP** protocol, you can specify the **Other Settings** options.

Note: Compared to the UDP protocol, the TCP protocol uses significantly more CPU processing power and hence decreases the maximum message rate the appliance supports.

6. Select the **Enable** check box to activate message forwarding.
7. Using the **Format Settings**:

Note: The **Insert Syslog Header** option is disabled for All Sources rule.

- Set the **Destination Parsing** (**Yes**/**No**) to enable or disable destination parsing. When enabled, the system automatically generates default rules for each protocol for all destinations.

Note:

- The **Destination Parsing** option is enabled when you select LogLogic LMl Appliance as the Destination Type. When you enable this option and click **Add**, three rules are added, one for each protocol type. Based on its log source type, a message is forwarded using one of the three routing rules. All syslog logs are forwarded using TCP protocol. All file-pulled logs are forwarded using LogLogic TCP protocol, and all SNMP trap messages are forwarded using SNMP protocol.

  When three rules are added (after enabling the **Destination Parsing** option), you can go back to **Edit Destination** window to select the configuration rule file for the rules which are using the LogLogic TCP and TCP syslog protocols. The **Format Rule Definition** field is disabled for the rule using SNMP protocol.

  If the **Destination Parsing** option is enabled, the **Format Rule Definition** option to format messages prior to forwarding is disabled, and vice versa.

- If you do not enable the **Destination Parsing** option, only the specified rule for the selected protocol is added. In this case, messages from some of the log source type may not be forwarded if the selected protocol is not compatible with the log source type. For example, syslog source type cannot be forwarded using SNMP protocol.

- (Optional) Specify the **Format Rule Definition** configuration rule file to format messages prior to forwarding. All messages that match the forwarding rule are formatted. For detailed description about defining the configuration rule file and how messages are formatted, see Definition of Configuration Rule Files.

Message Routing – Newly-added destination with three rules

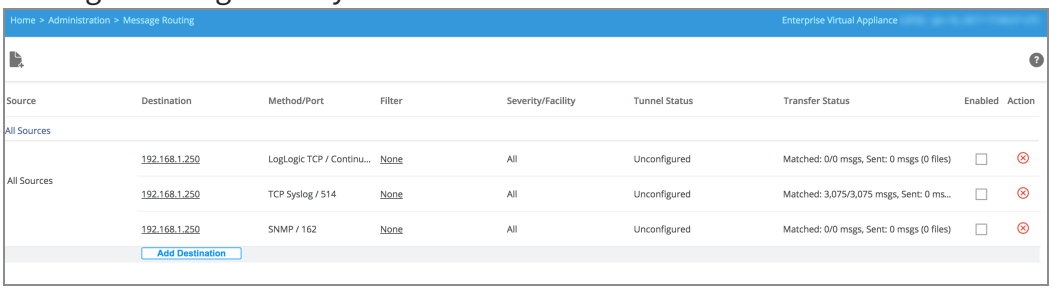

The Tunnel Status column displays the status of the connection between the source and destination when message routing is configured and the **Enable Authentication and Encryption** option is set to Yes. One of the following values is displayed:

Value	Description
Unknown	Initialization in progress. The status is not updated on the page.
Unconfigured	Forwarding is configured not to use encryption or authentication.
Starting	Tunneling is being established by initiating a downstream connection. The connection has not been completed yet.
Connected	The tunnel connection has been established.
An error message	Forwarding either failed to establish a tunnel or forwarding through the tunnel failed.

8. **LogLogic Forwarding Settings**:

   You cannot specify any options. The options are disabled for **All Sources** Rule.

9. **Other Settings**: This section is disabled when using UDP Syslog.
   - Set the **Compression** (**Yes**/**No**) to activate or deactivate compression for message routing. For LogLogic LX Appliances or LogLogic MX Appliances using LogLogic TCP, it is good practice to select **Yes**. The default is **No**.
     - Compression is available only when using LogLogic TCP.
     - You can enable compression or authentication and encryption in the following steps only when the routing destination is another LogLogic LMI appliance.
     - Setting Compression to **Yes** or enabling Authentication and Encryption for any single source/protocol/destination configuration causes all subsequent traffic from the same source sent with the same protocol to the same destination to be either compressed or authenticated and

encrypted. The system does not allow for both encrypted and clear traffic to go to the same IP via the same protocol when sent from the same source. Likewise, all traffic must be either compressed or non-compressed, but not both types.

- Set the **Enable Authentication and Encryption** (**Yes**/**No**) to activate or deactivate authentication and encryption for additional security.

  Using authentication ensures that the data is received by the correct LogLogic LMI appliance.

  - Authentication and encryption cannot be selected separately.
  - The **Enable Authentication and Encryption** option is not available when forwarding messages with the UDP protocol.
  - When you activate the **Enable Authentication and Encryption** option, the authentication and encryption are performed by using the SSH protocol. The `toor` user of the upstream appliance must be authorized to login via SSH to the downstream appliance without entering a password. To configure, type the CLI command `system keycopy` on the upstream appliance and follow the instructions displayed on screen to add the public key of the upstream appliance to the downstream appliance.

    If you select the **Enable Authentication and Encryption** option with **TCP Syslog** as the routing protocol, then for messages that do not contain a syslog priority, the log source is identified as `127.0.0.1_General` instead of the actual IP address of the source device. For messages that contain a syslog priority, the log source is correctly identified with its original source IP. This causes all events without a syslog priority from multiple sources to have their logs associated to the single source `127.0.0.1`.

    If you do not select the **Enable Authentication and Encryption** with **TCP Syslog** as the routing protocol, then for messages that do not contain a syslog priority, the log source is identified as `<upstream LMI IP Address>_General` instead of the actual IP address of the source device. For messages that contain a syslog priority, the log source is correctly identified with its original source IP. This causes all events without a syslog priority from multiple sources to have their logs associated to the single, upstream LogLogic LMI IP address source.

Enable Authentication and Encryption	Routing protocol	Messages contain Syslog priority?	Log source is identified as
Selected	TCP Syslog	No	`127.0.0.1_General`
		Yes	Original source IP address
Not selected	TCP Syslog	No	*<upstream LMI IP Address>*`_General`
		Yes	Original source IP address

10. Click **Add** to add the destination to the All Source rule.

    If you selected the **Enable Authentication and Encryption** option while creating the rule, then after rule creation you must perform the following steps for the changes to take effect:

    a. Run the `system keycopy` command.

    b. Disable and reenable the rule from the Message Routing page.

    The **Message Routing** screen appears showing the newly added destination to the existing All Source rule.

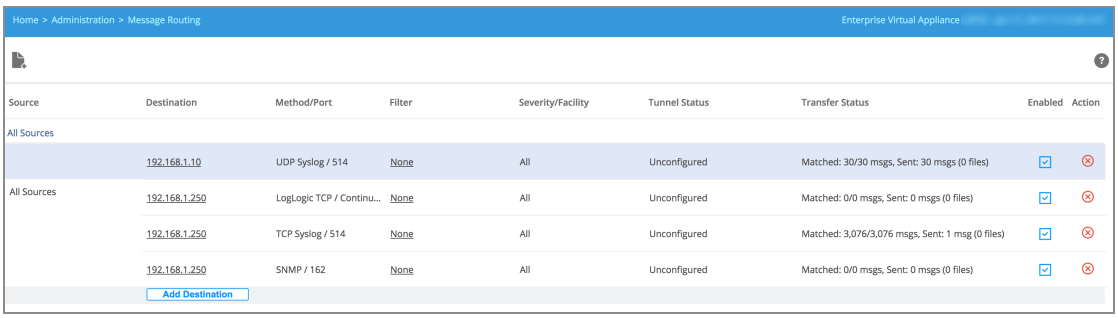

Result

You can enable, disable, or delete the rules; or edit, add, or remove the destinations in the rules.

Related topics

- Editing Routing Rules
- Editing Filters
- Editing Log Sources
- Removing Routing Rules or Destinations

Creating a New Outbound Routing Rule

Procedure

1. Access **Administration > Message Routing** from the navigation menu.
2. Click the **Create New Rule** button  to create a new routing rule.
3. In the **Rule Name** field, enter a name for the routing rule and click **Next**.
4. In the **Add Log Sources** section, click the down arrow next to **Select** and pick a log source filter:
   - Name
   - Collector Domain
   - IP Address
   - Group
   - Type

   a. If you picked Name, enter a Source Name, a specific Source Name or a Name Mask. Wild cards are accepted in this field.

   b. If you picked Collector Domain, enter the name of the Collector Domain.

   c. If you picked IP Address, enter a Source IP Address, a specific IP Address or an IP Address Mask. Wild cards are accepted in this field.

   d. If you picked Group, enter a Group Name, or click the down arrow to the right of the text field and select “All” or one of the other Group names displayed in the drop-down box.

   e. If you picked Type, enter a Source Type (a specific device type), or click the down arrow to the right of the text field and select All or one of the other

Device Types displayed in the drop-down box.

> **Note:**
> - If you select mixed log types from a user-defined Group, the available options such as **Protocols** and **Settings** can be different compared to a rule that contains only a single log type.
> - When adding a large number of devices of the same log type, use the system-defined Group option. Select one or more Groups as long as they are of the same log type, and then click the **<< Add selected log sources** button.

- If required, add a second filter by clicking the **+ sign** and repeating step 4 as often as you like.
- To delete a filter, click the **- sign** to remove the last selection made (repeat if needed).

5. Select a log source by clicking its name. You can select multiple rows.

> **Note:** From the **Administration > System Settings > General** tab, if you have selected **Optimize Device Selection List > Show Only Device Groups** option, **Available Devices** lists only **Device Groups**.

6. Click **<<Add selected log sources** to move the selected log sources to the left.
7. Click **Next**.
8. In the **Destination IP** field, type the IP address of the destination to which you want to forward messages. This can be another LogLogic appliance, or a LogLogic Security Event Management (SEM) appliance, or another machine (with correct port configuration). This is a mandatory field.
9. In the **Destination Port** field, type the port number to which you want to forward messages.
10. From the **Destination Type** list, select where you want to forward messages:
    - LogLogic LMI Appliance
    - LogLogic SEM Appliance

- Other Destination

If you choose a file-based log source such as Blue Coat ProxySG from the **Source Device** list, and Other Destination from **Destinations Type**, then the **Insert Syslog Header Yes/No** radio buttons are displayed. Selecting **Yes** adds <109> at the beginning of the message. The prefix <109> at the beginning of the message is the syslog priority for audit information events. It is included to prevent triggering of intrusion detection systems and firewalls that detect syslog without a proper header.

11. From the **Protocol** drop-down menu, select the protocol to use for forwarding messages:

Option	Description
UDP Syslog	Traditional syslog using the UDP protocol **Note:** The UDP Syslog protocol is not supported on LogLogic EVA - Container Edition.
TCP Syslog	Traditional syslog using the TCP protocol. Also known as Syslog-NG New-line (\n) characters are used to break logs in the TCP stream during message forwarding. If a message contains \n, the message breaks up with only the first portion of the message being delivered to the downstream appliance. It is good practice to select a different forwarding protocol if you know your log messages contain characters of this type.
LogLogic TCP	Buffered syslog provided by TIBCO LogLogic®. Uses a proprietary TCP-based protocol and uploads logs in batches every minute **Note:** If you select **LogLogic TCP** protocol, you can specify the **Other Settings** options.

Note:
- Compared to the UDP protocol, the TCP protocol uses significantly more CPU processing power and hence decreases the maximum message rate the appliance supports.
- Depending on the selected **Destination Type** and **Protocol** values, some of the **Format Settings**, **LogLogic Forwarding Settings** and **Other Settings** options may be available.

12. Select the **Enable** check box to activate message routing.
13. Using the **Format Settings**:

Note: The **Insert Syslog Header** option is only enabled for File-based log messages.

- Select the **Insert Syslog Header** radio button (**Yes**/**No**) to activate or deactivate Syslog headers.
- (Optional) Specify the **Format Rule Definition** configuration rule file to format messages prior to forwarding. All messages that match the forwarding rule are formatted. For detailed description about defining the configuration rule file and how messages are formatted, see Definition of Configuration Rule Files.

14. Using the **LogLogic Forwarding Settings**:

Note: This setting is available only when forwarding real-time logs to a LogLogic appliance using LogLogic TCP protocol.

- From the **Forwarding Type** list, select whether real-time log files are transferred Daily or Continuously.

  Selecting daily minimizes the time frame for performance impact on the network and related systems. However, continuous forwarding allows more immediate use of the log data on the appliance.

  If you select daily forwarding, set the following:

- **Start Time**—Time that daily real-time log file transport starts.
- **End Time**—Time that daily real-time log file transport ends.

Any log files not transported by the end time are the first transferred the next day.

- **Max Bytes/Sec**—The maximum transfer rate allowed for log file transport. 0 means unlimited. The acceptable range is 0 through 125000000.

15. Using the **Other Settings**:
    - Select a **Compression** radio button (**Yes**/**No**) to activate or deactivate compression for message routing. For LX or MX appliances using LogLogic TCP, LogLogic recommends selecting **Yes**. The default is **No**.
        - **Compression** option is available only when you select the LogLogic TCP Protocol.
        - You can enable compression or authentication and encryption in the following steps only when the routing destination is another LogLogic appliance.
        - Setting Compression to **Yes** or enabling Authentication and Encryption for any single source/protocol/destination configuration causes all subsequent traffic from the same source sent with the same protocol to the same destination to be either compressed, or authenticated and encrypted. The system does not allow for both encrypted and clear traffic to go to the same IP via the same protocol when sent from the same source. Likewise, all traffic must be either compressed or non-compressed, but not both types.
    - Select an **Enable Authentication and Encryption** radio button (**Yes**/**No**) to activate or deactivate authentication and encryption for additional security.

      Using authentication ensures that the data is received by the correct LogLogic appliance.
        - Authentication and Encryption cannot be selected separately.
        - The Authentication and Encryption option is not available when forwarding messages with the UDP protocol.
        - When you select the **Enable Authentication and Encryption** option, the authentication is performed using the SSH protocol. The toor user of the upstream appliance must be authorized to login via SSH to the downstream appliance without entering a password. To configure, type

the CLI command system keycopy on the upstream appliance and follow the instructions displayed on screen to add the public key of the upstream appliance to the downstream appliance.

If you select the **Enable Authentication and Encryption** option with **TCP Syslog** as the routing protocol, then for messages that do not contain a syslog priority, the log source is identified as `127.0.0.1_General` instead of the actual IP address of the source device. For messages that contain a syslog priority, the log source is correctly identified with its original source IP. This causes all events without a syslog priority from multiple sources to have their logs associated to the single source `127.0.0.1`.

If you do not select the **Enable Authentication and Encryption** with **TCP Syslog** as the routing protocol, then for messages that do not contain a syslog priority, the log source is identified as `<upstream LMI IP Address>_General` instead of the actual IP address of the source device. For messages that contain a syslog priority, the log source is correctly identified with its original source IP. This causes all events without a syslog priority from multiple sources to have their logs associated to the single, upstream LogLogic LMI IP address source.

Enable Authentication and Encryption	Routing protocol	Messages contain Syslog priority?	Log source is identified as
Selected	TCP Syslog	No	`127.0.0.1_General`
		Yes	Original source IP address
Not selected	TCP Syslog	No	`<upstream LMI IP Address>_General`
		Yes	Original source IP address

16. Click **Next** to define the message Filters including Severity, and Facility or click **Finish** to accept the default message filters. In this case, skip the following steps and go to

step 20.

> **Note:** For file-based log sources, select the desired filter you created earlier from the **Search Filter** list. Boolean searches are not supported for file transfer sources; only three kinds of search filters are supported: “Use Words”, “Use Exact Phrase”, and “Regular Expression”.

17. Select the existing search filter from the **Search Filter** list.

> **Note:** If you want to add a new search filter, use the **Search > All Search Filters** menu. For more information, see "Adding a Search Filter" in the *TIBCO LogLogic® Log Management Intelligence User Guide*.

18. Click the **Forward all except filter matches** check box to forward those messages that do not match the defined search filter.
19. Select the **Message Severity** and **Facility** filters that you wish to select or **Select All** if you want everything forwarded.

    By default, all check boxes are selected for syslog-based log sources. For complete details about your Message Severity and Facility options, see your firewall documentation.

Message Severity - Standard Descriptions

Type	Description
Emergency	System is unusable
Alert	An alert condition exists
Critical	The system is in critical condition
Error	An error condition exists
Warning	A warning condition exists

Type	Description
Notice	A normal but significant condition
Informational	Information message without any serious conditions that exist
Debug	Messages generated to debug the application

Note: To find out how each vendor uses severity values with respect to their messages, see your vendor documentation.

The facility specifies the subsystem that produced the message. For example, all mail programs log with the mail facility (LOG_MAIL) if they log using syslog.

Note: Filtering criteria here applies only to syslog forwarding, not file transfer sources. For details about file transfer, see Adding a Log Source for File Transfer.

20. Click **Finish**.

    If you selected the **Enable Authentication and Encryption** option while creating the rule, then after rule creation you must perform the following steps for the changes to take effect:

    a. Run the `system keycopy` command.

    b. Disable and reenable the rule from the Message Routing page.

    The **Message Routing** screen appears showing the newly added Routing Rule.

Addition of Destinations to the Existing Routing Rule

You can add a destination to the existing routing rule.

Note: You cannot add the same destination IP address and protocol type twice in the same rule.

Once the destination is added, the Message Routing window is displayed, showing the newly added destination under the existing rule.

Editing Routing Rules

You can modify the existing destinations for the routing rule.

> **Note:** If you have enabled the **Destination Parsing** option while adding the destination for **All Sources** rule, some options are disabled from the Edit Destination window.

Procedure

1. Click the destination IP address that you wish to update from the Destination list.
2. Make the necessary changes. For more information about the fields, see Destinations Addition to the Existing Routing Rule.

   > **Warning:** If you have selected SNMP in the **Protocol** field, and the **Destination Type** is LogLogic LMI appliance, then do NOT change the **Destination Port**. Keep the default Destination Port 162. Otherwise, you cannot receive the forwarded messages in the Destination Appliance.

3. Click **Save**.

Result

The **Message Routing** screen displays the updated information.

Editing Filters

You can modify the existing filters for the routing rule.

Procedure

1. Click the Filter name that you wish to update under the **Search Filter** column.
2. The Edit Filter window appears.

3. Make the necessary changes. You can update the filter, Message Severity, and Facility information.
4. Click **Save**.

Result

The **Message Routing** screen displays the updated information.

Editing Log Sources

You can modify the existing log sources for the routing rule.

Procedure

1. Click the **Devices** link to open the Edit Source window.
2. Make the necessary changes.
3. Click **Save**.

Result

The **Message Routing** screen displays the updated information.

Removing Routing Rules or Destinations

You can delete the existing routing rule or destination.
However, you cannot delete the **All Sources** rule.

> **Note:** If you delete the last added destination from the **All Sources** rule, the **All Sources** rule still remains in the system. However, if you delete the last added destination from any Outbound data rule, that Outbound rule is deleted.

Procedure

1. To delete a rule and all of its destinations, click the  button next to the rule.

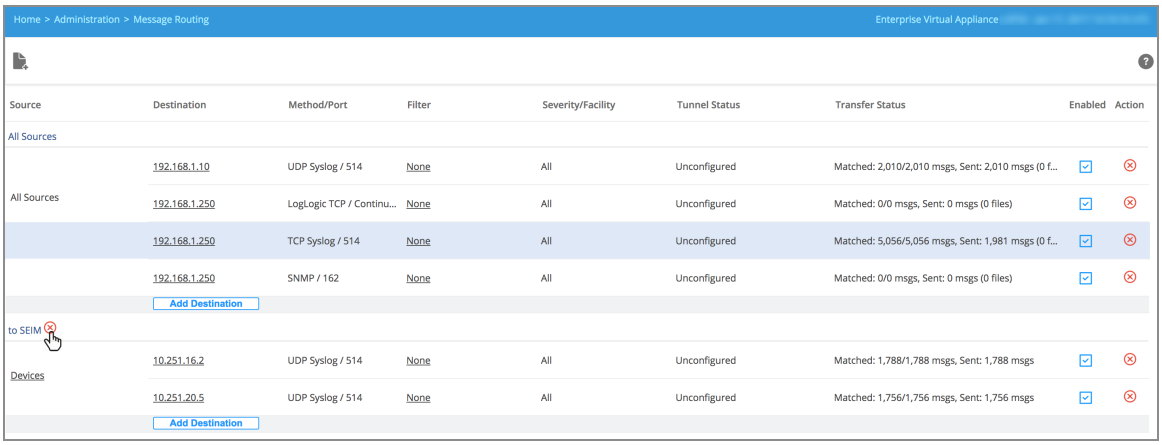

2. To delete a destination from a rule, click the ⊗ button in the **Action** column of that rule.

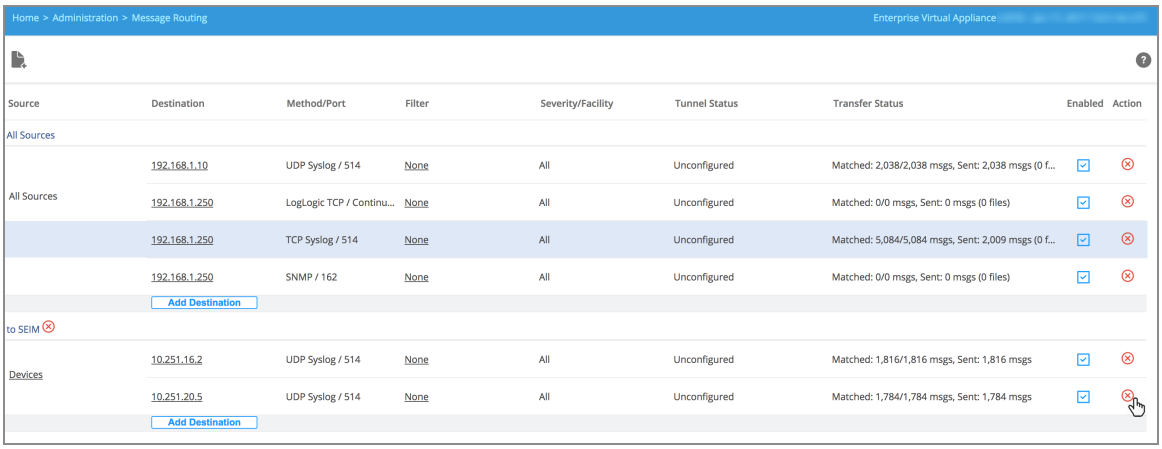

3. Click **Yes** in the confirmation window to delete the rule or the destination.

Replay of Archived Data

Replay lets you re-analyze archived log data by processing it from its archived location on an LogLogic ST Appliance with an LogLogic LX Appliance as its remote appliance.

The LogLogic LX Appliance treats the data as if it were new data, and sends it through the parsing process again.

Because you are replaying archived data, the original time stamps on the log data are kept, so you need to run reports and searches with this in mind. The archived data can then be made available to custom reports and searches.

Replay is particularly useful if you recently added support for new log sources, reports, or Compliance Suites.

Note:
- Replay is not supported on LogLogic MX Appliances.
- Replay only works with IPv4 addresses.

How Replay Works

Replay requires a source LogLogic ST Appliance and a destination LogLogic LX Appliance to be configured in a Management Station relationship.

The LogLogic ST Appliance must be a Management Station that manages the LogLogic LX Appliance. The Management Station relationship ensures that you manage Replay sessions correctly.

Warning: When using Replay, the LogLogic LX Appliance must **not** be set up as a Management Station. If the configuration is not correct, replay does not work.

Note: Archived real-time files on the source LogLogic ST Appliance are always rediscovered during a Replay session whether or not a search filter is used. Rediscovering real-time files lets additional devices be recognized that were not known during the initial capture by the LogLogic LX Appliance or LogLogic ST Appliance. However, file-based logs are not rediscovered at this time.

Pulled files are always replayed as a whole file. However, real-time logs can be subjected to filtering.

The source LogLogic ST Appliance and destination LogLogic LX Appliance manage the progress of each Replay session. Therefore, if at any point a Replay session is interrupted (for example, the network goes down or the appliance service is not available):

1. The source LogLogic ST Appliance keeps trying to replay data infinitely until a connection is re-established.
2. Once the connection is re-established, the data transfer resumes where it left off. After the replay is completed, the Replay Status is updated to 'completed' on the Replay Status tab.

How a Replay session works

1. The scheduled Replay session starts.
2. Replay gathers the appropriate archived data on the source LogLogic ST Appliance based on the Replay rules specified in the Replay session. The source LogLogic ST Appliance notifies the destination LogLogic LX Appliance how many files it is transferring.
3. The source LogLogic ST Appliance transfers the appropriate archived log data to the destination LogLogic LX Appliance. Authentication and encryption are used only if configured for the Replay session.
4. All log data is received by the destination LogLogic LX Appliance, so the LogLogic LX Appliance begins processing the data as new data. Log data is received by LLTCP-HTTP.
5. After all log data is processed by the destination LogLogic LX Appliance, it notifies the source LogLogic ST Appliance that the Replay session is completed.
6. The source LogLogic ST Appliances ends the Replay session and updates the status

to completed.

Note: The maximum replay number is 16. Canceled and completed replays are not included in the total number.

The user must have Search Archived Data privileges on the LogLogic ST Appliance to replay the archived data. For more information on user privileges, see User Privileges.

Configuration of Replay Environment

It is good practice to set up a dedicated destination LogLogic LX Appliance to handle Replay sessions.

Dedicating an appliance lets you focus on only the data you want to re-analyze and does not affect the production environment hard drive space and message handling.

Warning: If you do not use a dedicated LogLogic LX Appliance for Replay sessions, you risk having duplicate data. Reports and searches intended to be done on “production” data can also pick up Replayed data, giving you inaccurate results.

You can configure your Replay environment to support the following source to destination relationship:

- One to one: Single source LogLogic ST Appliance using a single destination LogLogic LX Appliance

Data Retention

If the destinationLogLogic LX Appliance becomes full during a Replay session, standard data retention rules for the appliance apply.

That is, the oldest files are purged to make room for the new messages coming in from the Replay session. It is good practice to configure data retention for the LogLogic LX Appliance you are using to handle Replay sessions.

The retention time is counted from the time the log data was generated by the original log source.

If everything is set to the same expiration policy (for example, 1 year), then the oldest file is purged. However, when an appliance is configured with multiple different retention policies, files with the earliest expiration date are purged first. For example, files that are set to expire in one week are deleted under emergency purging circumstances before files that expire after one year and 10 months.

Configuring the Replay Session to Use Authentication

If a replay session is configured to use authentication, the source LogLogic ST Appliance must present an authentication key to the destination LogLogic LX Appliance.
The LogLogic LX Appliance asks for authentication only if the LogLogic ST Appliance is configured as an upstream device with an authentication key. If the LogLogic LX Appliance is configured without authentication, any upstream device can connect without requiring an authentication key. However, the LogLogic LX Appliance does not need to send an authentication key to the LogLogic ST Appliance.

Procedure

1. Run `system keycopy`. See system Command.
2. When you schedule a replay session, select the **Authentication Required** check box. See Scheduling a Replay Session.

Note: If a key mismatch with an authenticated channel prevents the LogLogic ST Appliance from connecting to the LogLogic LX Appliance, an error message is captured in the sys.log file for both appliances.

What to do next

Auto-Identify Turned On in the LogLogic LX Appliance

If auto-identify is turned on in the destination LogLogic LX Appliance, any forwarding appliance (source LogLogic ST Appliance) can connect without sending an authentication key. However, if the upstream device is configured on the destination LogLogic LX Appliance with an authentication key, the key must match the key from the source LogLogic ST Appliance.

Auto-Identify Turned Off in the LogLogic LX Appliance

If auto-identify is turned off in the destination LogLogic LX Appliance, only configured upstream devices can connect to the LogLogic LX Appliance. If the LogLogic LX Appliance also is configured to use an authentication key, the key must match the key from the source LogLogic ST Appliance.

Configuration of Appliances to Replay Archived Data

To configure appliances to replay archived log data from an LogLogic ST Appliance to an LogLogic LX Appliance, you must configure the LogLogic LX Appliance and then the LogLogic ST Appliance.

1. Configuring the LogLogic LX Appliance
2. Configuring the LogLogic ST Appliance

Configuring the LogLogic LX Appliance

To configure an LogLogic LX Appliance to process archived log data from an LogLogic ST Appliance, you must complete the following steps.

1. Configuring the LogLogic LX Appliance to Analyze Data
2. Clearing All Log Data from the LogLogic LX Appliance

Configuring the LogLogic LX Appliance to Analyze Data

It is good practice to set up the LogLogic LX Appliance that replays archived data as you would a production appliance.
Specifically, to obtain the maximum benefit of replaying archived log data, ensure that you have all of the appropriate components and system settings configured in your Replay Appliance.

Consider configuring at least the following:

Alerts	Configure alerts to send SNMP events or email notification of specific occurrences found in the data in the replay session. **Note:** System Alerts (Message Volume and Ratio-Based) might produce skewed results because the data is being sent all at once rather than over the time period which it was originally sent. It is good practice to use message-based alerts instead.
Reports	Configure reports to analyze the data in the replay session.
Search Filters	Configure search filters to run reports and searches on specific log data.
Devices	Ensure that you have all applicable devices configured.
Full Text Indexing	Consider turning on full-text indexing on all data (parsed and unparsed; unparsed data is log data that is not associated with a supported log source).
PIX/ASA Messages	Enable if the archived data contains PIX/ASA messages (if you enable PIX/ASA Messages and you do not have PIX/ASA messages in the replay session, it does not impact the appliance).
Message Routing	Enable only if you need to forward log data to another device.

Data Retention	Configure how long to retain the data from the replay session on the destination LogLogic LX Appliance (retention time is counted from the time the log data was generated by the original log source).

To speed up the setup process, use the Import/Export tool. For example, you can import components such as search filters and reports from any LogLogic LX Appliance. You must manually set system settings such as global retention settings and full-text indexing. For more information on importing and exporting components from one appliance to another, see Import or Export Entities Between Appliances.

Clearing All Log Data from the LogLogic LX Appliance

Before sending archived log data to an LogLogic LX Appliance configured for replay, consider clearing the appliance of all log data.
A clean appliance lets you run reports and searches on only the archived data you are replaying. If you want to combine log data from multiple replay sessions, do not clear the log data.

> ⚠ **Warning:** The cleanup process removes all log data on the LogLogic appliance. It does not remove configuration data (such as system settings) or reports, search filters, and so on.

When logged in directly to the destination LogLogic LX Appliance or when managing the LogLogic LX Appliance from the LogLogic ST Management Station, Go to **Administration > Clear Log Data**. The Clear Log Data tab is displayed.

For an HA environment

The clear log data operation cannot be performed in an HA environment. To clear log data from a node, first the node must be removed from the HA environment and then perform the operation.

Related topic

Removal of Appliance Log Data

Configuring the LogLogic ST Appliance

To configure an LogLogic ST Appliance to process archived log data to a destination LogLogic LX Appliance, you must complete the following steps.

1. Setting Up a Management Station Relationship
2. Adding a Replay Rule
3. Modifying a Replay Rule

Setting Up a Management Station Relationship

You must set up the source LogLogic ST Appliance as a Management Station with the destination LogLogic LX Appliance as an appliance in the Management Station cluster.

Procedure

1. On the LogLogic ST Appliance, in the navigation menu click **Management > Management Station**.

   The **Configuration** tab appears. See Manage Appliances with Management Station.
2. For the LogLogic LX Appliance to be used the destination for the archived data to be replayed, enter its:
   - Appliance IP or DNS Name
   - Appliance Name
   - Appliance Type
3. Click **Add**.

Result

The LogLogic LX Appliance appears as an appliance in the Management Station cluster.

Adding a Replay Rule

Replay rules let you define specific data to include in a Replay session.
Each Replay rule identifies data from the specific device and time frame, so you can specify to push data associated only with certain devices or from all devices.

For example, you can create a rule that pushes data for your Blue Coat Proxy SG log sources from `03/11/09 at 00:00:00 to 03/12/09 at 23:59:59`. You can also define a rule to push data for a specific Cisco PIX/ASA log source by specifying the device type as Cisco PIX/ASA and the Source Devices as the specific log sources.

Procedure

1. On the destination LogLogic ST Appliance, navigate to **Administration > Replay**.
2. Click the **Replay Rules** tab.

   The **Replay Rules** tab appears listing all existing Replay rules in the appliance.
3. Click the **Add Rule** button.

   The **Add Replay Rule** tab appears.
4. Enter the following information:

Option	Description
Rule Name	Name of the rule
Device Type	Select the device or application generating the logs to be transferred
Source Device	IP address of the device from which you want to transfer files
Search Filter	Select the Pre-Defined search filter to use to filter the archived log data
Time Interval	Time interval for the archived data you want to process

5. Click **Save** to save the Replay rule.

What to do next

After you add your Replay rule you can schedule a Replay session that uses your Replay rules.

Modifying a Replay Rule

Procedure

1. From the navigation menu of the destination LogLogic ST Appliance, click **Administration > Replay**.
2. Click the **Replay Rules** tab.

   The **Replay Rules** tab appears listing all existing Replay rules in the appliance.
3. Hover over the name of an existing **Replay rule** and left-click. The **Modify Replay Rule** tab appears.
4. Enter the following information:
   - **Rule Name**—Name of the rule.
   - **Device Type**—Select the device or application generating the logs to be transferred.
   - **Source Device**—IP address of the device from which you want to transfer files.
   - **Search Filter**—Pre-defined search filter to use to filter the archived log data.
   - **Time Interval**—Time interval for the archived data you want to process.
5. Click **Save** to modify the Replay rule or **Cancel** to discard modifications.

What to do next

After you modify your Replay rule you can schedule a Replay session that uses your Replay rules.

Replay of Archived Data

After you configure the LogLogic LX Appliances and LogLogic ST Appliances and set up Replay rules, you can schedule a replay session.

Note: If you run a report in the destination appliance on newly replayed data, you might see only a portion of the data since the appliance needs time for aggregation. Specifically, if you run a report, the count (number of entries) might not match the actual detailed data that you see when you drill down on the count. Try modifying the search interval or run the report later.

- Scheduling a Replay Session
- Viewing Replay Progress

Scheduling a Replay Session

You can schedule a Replay session to run immediately or at a scheduled time in the future. You can schedule multiple Replay sessions to run from the same source LogLogic ST Appliance, but the destination LogLogic LX Appliance must be different. Replay sessions are serialized and start in sequence.

Warning: When scheduling a replay, if you select **Authentication** and **Encryption** options, type the CLI command `system keycopy` on the LogLogic ST Appliance and follow the instructions displayed on the screen to add the public key to the LogLogic LX Appliance.

Note: The real-time logs can be replayed multiple times. Duplicate logs are not rejected by the LogLogic LX Appliance. However, file-based logs are accepted only once and duplicate logs are rejected by the LogLogic LX Appliance. Pulled files are always replayed as a whole file. However, real-time logs can be subjected to filtering.

Procedure

1. From the navigation menu of the LogLogic ST Appliance, click **Administration > Replay**.

   The **Replay Status** tab appears listing all existing Replay sessions in your system.

   If no Replay session has been scheduled yet, the **Replay Status** tab displays No match found in database even though you may have added a Replay rule (or modified an existing one) and configured the Time Interval—you still must schedule the Replay rule to run in a Replay session.

2. To schedule a Replay session, click the **Schedule Replay** button.

   The **Schedule Replay** tab appears.

3. In the **Schedule Replay** tab, enter the following information:

   - **Destination**—IP address of the destination LogLogic LX Appliance used for the Replay session.
   - **Replay Rules**—Select the appropriate Replay rule (to add additional rules, click **Add Rule**).

- **Authentication Required**—Select the checkbox to enable authentication between the source LogLogic ST Appliance and the destination LogLogic LX Appliance.
- **Encryption Required**—Select the checkbox to require encryption of the data sent from the source LogLogic ST Appliance to the destination LogLogic LX Appliance.
- **Schedule replay to run immediately**—Select the checkbox to schedule the Replay session to run immediately upon clicking the Save button.
- **Start Time**—Start date and time to run the Replay session.

4. Click **Save**.

Result

The **Replay Status** tab appears with the new scheduled replay session. If you scheduled the Replay session to occur in the future, the State appears as pending.

All completed Replay sessions remain in the Replay Status page showing their state. You can remove a Replay session.

Viewing Replay Progress

When using replay, you can view the progress of the log data as it is gathered and sent from the source LogLogic ST Appliance, as well as the progress of incoming log data to the destination LogLogic LX Appliance.

Viewing replay progress in the source LogLogic ST Appliance

The LogLogic ST Appliance lets you view the progress of a running Replay session as well as the status of all schedule Replay sessions.

Procedure

1. On the LogLogic ST Appliance, go to **Administration > Replay**. The Replay Status tab displays all existing replay sessions in the system.
2. Go to the appropriate Replay session in the Replay Status table and view the State

and Status tabs.

The State column lists the current state of the Replay session:

State	Description
canceled	Replay session was canceled by a user
completed	Replay session is complete
in progress	Replay session is currently running
pending	Replay session is scheduled to run

The Status column lists the status of the Replay session:

State	Description
Messages	The total messages to process
retrieved	The total messages to be sent to the destination LogLogic LX Appliance
sent	The total messages sent to the destination LogLogic LX Appliance

Canceling a Replay Session

You can cancel any Replay session that is in progress or that is scheduled to run.

Depending on the state of the replay session, you might need to do further clean-up of the appliance. Specifically, you might want to clear the log data in the destination LogLogic LX Appliance if the log data in the replay session was being parsed.

Note: If you cancel a Replay session that is in progress, the Replay session finishes the file it is currently processing before stopping.

Procedure

1. From the navigation menu of the LogLogic ST Appliance, click **Administration >**

Replay.

The **Replay Status** tab appears listing all existing Replay sessions in your system.

If no Replay session has been scheduled yet, the **Replay Status** tab displays No match found in database even though you may have added a Replay rule (or modified an existing one) and configured the Time Interval—you still must schedule the Replay rule to run in a Replay session.

2. Select the appropriate Replay sessions, and then click **Cancel**.
3. Confirm that you want to cancel the Replay session.

Result

The **Replay Status** tab appears with your Replay session showing a state of cancelled.

Backup and Restore

Backup and restore is intended for use in securing important data if an appliance failure occurs or old data is needed for other reasons.

You can restore the backup only to an appliance that has identical LogLogic LSP settings.

Backup and restore is not a solution for migrating data from one system to another; use the LogLogic data migration solution instead. For more information, see Data Migration Between Appliances.

It is good practice to:

- Maintain regular backups in the event of appliance failure or other needs to recover old data
- Back up the appliance before upgrading it to a newer LogLogic software release

Backup and Restore Architecture

The various backup methods let you choose between greater security and performance.

You can back up LogLogic LMI data using the following methods:

- SAN - only on appliances that support Storage Area Network (SAN)

  For a list of SAN appliances, see *TIBCO LogLogic® Log Management Intelligence Hardware Installation Guide*.

- SCP, NFS, or Amazon S3 - on any appliance

You can schedule daily backups or run an immediate backup at any time. The complete data set is copied from the appliance to the backup device, without stopping data collection during the backup operation.

You can restore an appliance from backups made using any method when needed.

Backup Methods

You can back up the appliance data by using different backup methods; some offer better speed, and some offers better security.

Using the backup methods, you can do the following:

- Schedule regular backups to run weekly on specified days or daily.
- Run a backup immediately after configuring the backup settings.
- Designate backups as full or incremental (optimized) backups.

> **Note:** You cannot perform incremental or optimized backups to Amazon S3.

Method	Backup system requirements	You provide...	Fast?	Secure?
SCP	Any server that supports SCP with a user name, typically any UNIX or Linux system.	A server name, user name, and the directory path location where you want to save the backup files	No	Yes
NFS	An NFS volume mounted by the appliance. The versions v1-v4 are supported.	A server name and directory path location where you want to save the backup files	Yes	No
SAN	A SAN device available to the appliance.	An HBA and UUID device number	Yes	No
Amazon S3	An Amazon S3 bucket	Identity, AWS access key ID, secret key, and bucket name	Yes	Yes

Limitations

Backup is not supported in the following scenarios:

- Backup to a Microsoft Windows system
- NFS backup using the UDP protocol

Backup Interfaces

Backup and restore are both controlled entirely through the LogLogic appliance GUI **Administration > Backup/Restore Configuration** page.

The only exception is that before running an SCP backup for the first time you must configure SCP through the appliance CLI.

The **Backup Configuration** tab is accessible through the GUI on LogLogic LX Appliance, LogLogic MX Appliance, and LogLogic ST Appliance.

Use the **Backup Configuration** tab to specify settings for saving a backup copy of your appliance log and configuration data. The default settings for backup method are **None** and **Optimize**.

Note: Your backup location must support hard links.

Select **SCP** to back up any server which supports SCP with a user name, typically any UNIX or Linux system.

Select **SAN** to backup any volume on the appliance to a Storage Area Network (SAN) device available to the appliance. **SAN** backup is available only on LogLogic ST Appliances.

If the **Optimize** check box (the default setting) is selected, it indicates that the appliance performs incremental backups instead of a full backup each time. An incremental backup copies only data that has changed since the previous backup, as opposed to a full backup which copies everything each time. These optimized backups save a lot of time for each backup, but still provide full restore capability if needed.

Note: The initial backup for the appliance is always a full backup.

Bypass disk space checking — Because the LogLogic application is very conservative, and because the total space required for backing up logs is difficult to accurately predict, the system might state that not enough storage space is available for backing up your log files when in fact you know that sufficient space is available. The default condition for this check box is not selected, meaning the appliance calculates (and possibly overestimates) the amount of backup space required. LogLogic estimates the space required based on the

number of backups specified to be retained. For example, if 3 backups are specified, LogLogic estimates the disk space required for 4 backups. This is because the first (oldest) backup is deleted only after the fourth one is complete.

If the **Bypass disk space checking** check box is selected on the **Backup Configuration** pane, it prevents the appliance from possibly overestimating the amount of storage space necessary for a full backup of your log data.

If the **Config only** check box is selected, the appliance backs up only the configuration data on the appliance.

You can schedule regular backups to run at a specific time regularly on any days of the week, or to run daily at that time.

What Is Backed Up?

The backup process copies the configuration and real time databases, and all raw log data from the appliance.

Backed up Databases and Files

Backed Up Data	Description
Configuration Database	A full SQL configuration database dump, created every time at the scheduled time.
Raw Log Files	The general raw syslog files, representing all logs collected from the appliance’s log sources, and stored in the appliance's local file system. If an appliance is configured with a supported archiving method, for example NFS remote server, the appliance pushes its older data such as raw syslog files, file-based log data, and indexed log data (if enabled) to the NFS remote server. In this case, the backup does not copy files from the NFS remote server. You must copy files from the NFS remote server manually by evaluating your business needs and creating a plan.
LEA Certificate files	Log Export API certificates specific to the Check Point Management Interface for the appliance.
Real Time	The Real Time data assembled by the appliance based on its collected logs.

Backed Up Data	Description
Database	The Real Time database is much larger for LogLogic LX Appliances and LogLogic MX Appliances than for LogLogic ST Appliances. The database on LogLogic ST Appliances consists of only a few tables.
System Configuration files	• Text files related to the appliance • LogLogic application configuration on it • Configuration files related to Advanced Features (Bloks, Data Model, Advanced Dashboard, and Advanced Search)

LogLogic backs up files on the appliance up to the midnight before when the backup is run. That is, if you start a backup at 10 a.m. the backup captures data through the midnight ten hours prior.

The only exception is that the backup collects only through 6 hours prior. That is, if you start the backup between midnight and 6 a.m. the backup copies data through midnight the previous day, not the most recent midnight.

How Backup/Restore Works

LogLogic backup copies a complete snapshot of appliance data to a different system, and updates that snapshot incrementally to minimize time and resource requirements for recurring backup processing.

Backup/Restore Processing

Appliance backup uses a system utility such as rsync to copy appliance data as a complete snapshot (the image of your data at the time of backup). The utility determines the difference between the data on the appliance and the previous backup data on the remote backup device before the data transfer.

- First backup - all data is copied
- Subsequent backups - only the difference since the previous backup is copied. This saves appliance resources and network bandwidth. The snapshot on the backup

system is updated to include the differences copied in the latest incremental backup.

Data collection continues during backup processing.

The backup system creates a directory named for the IP address of the backed-up appliance and stores its backup data there. This allows multiple appliances to be backed up to a backup system large enough to hold their cumulative data, as long as each appliance has a unique IP address.

After each backup completes, both data sets (the original appliance and the backup server data set) are identical. When restoring an appliance, the backed up information is copied back to the appliance. Processing does not continue during a restore process. For the duration of restore operation, all engines that access BFQ or MySQL (except `engine_backup` and `engine_tomcat`) are stopped. All log processing is resumed after the restore is completed.

Failure Situations & Workarounds

If the appliance crashes during a backup, the data in previous successful backups is protected. Each backup, including incremental backups, is written to a new location so no overwriting of existing backup data occurs.

Differences between Backup on Appliance Product Families

The backup feature functions similarly on all appliance product families.

The only significant difference is in the amount of disk space taken by different types of backed up data:

Backup Appliance Comparison

Appliance Product Family	Configuration Database	Real Time Database	Raw Log Files
LogLogic LX Appliance, LogLogic MX Appliance	Similar	Larger	Similar
LogLogic ST Appliance	Similar	Very small	Similar

Performance Metric

The following equation calculates the minimum time required for a backup or restore:

```
data-size / network-speed = minimum-time
```

- *data-size* is the total size (in MB) to backup or restore
- *network-speed* is the network speed (in MB/second)
- *minimum-time* is the least possible time the process takes (in seconds)

For example, the minimum time needed to back up or restore 10,000 MB (10 GB) data through a 100 Mb (12.5 MB/sec) network is:

10,000 / 12.5 = 800 sec = 13 min 20 sec

Performance beyond this minimum time depends on many factors, including network bandwidth, and processor speed and availability.

It is good practice to use a 1 GB interface to provide a higher network bandwidth.

Backup Storage

The backup data sets of an appliance are typically sent to and stored on a remote server using the SCP or NFS method.

Backup solution outline using external storage

LogLogic LMI appliances support multiple backup methods including NFS or SCP to transfer backup data sets to a remote storage system. Before successfully transferring the data to a remote server, the appropriate steps are required, for example:

- NFS method: create and export a volume on the NFS server
- SCP method: set up the SSH public-key-based authentication

Backup and Restore Scenarios

LogLogic LMI supports three primary backup/restore scenarios.

- Single system
- High Availability
- Disaster Recovery

These descriptions are cumulative. That is, the High Availability description details how to expand upon the single system backup/restore setup to set up High Availability backup/restore.

For more information, see the following sections:

Backup	Restore
Single System Backup	Single System Restore
High Availability Backup	High Availability Restore
Disaster Recovery Backup	Disaster Recovery Restore

Single System Backup

It is good practice to schedule daily backups through the GUI, at a time when network bandwidth and CPU usage are typically at a minimum. This ensures regular, complete backups that can most easily be restored from at any time if necessary.

You can backup multiple appliances to the same backup system, as long as the backup system has sufficient space and each appliance has a unique IP address.

Single System Restore

You can restore a backup to the original backed-up appliance, or to a replacement appliance as long as the replacement appliance has the same hardware, IP address, and software release (including maintenance releases and hotfixes) as the backed-up appliance.

Note: You can only restore a backup to an appliance for which a backup configuration is already defined, and for which the backup configuration is identical to that on the backed-up appliance. For example, if the backup location was `/home/john/lmi-backup` and options such as **Optimize** and **Bypass disk space checking** were used on the backed-up appliance, the same location and options must be used while configuring backup on the replacement appliance.

High Availability Backup

Backup and restore function similarly in a High Availability pair as they do for a single appliance.

The differences primarily involve the dynamic of having the second appliance involved.

When a High Availability pair is initially configured, the standby appliance automatically has backup disabled. The active appliance is the one that gets backed up as specified in the GUI.

High Availability Restore

Restoring data to the active appliance is similar to a single appliance restore.

Before restoring the active appliance, you must disable or shutdown the standby appliance. After the restore is complete, re-configure the standby appliance to join the HA. It should not be necessary to perform a restore to the standby; the restored data should mirror from the active to the standby as part of normal High Availability mirroring.

If the standby becomes the active appliance and you need to restore data to this new active appliance, it should perform as a normal restore would as long as the normal restore requirements (same IP address, appliance model, and software release) are met.

In case of an HA pair, the IP address checked is the virtual IP. Therefore, as long as the standby appliance being used is the one from the same HA pair, it can be used to restore the backup captured originally by its peer node.

Restoring in an HA pair

To restore a backup to both appliances in a high availability (HA) pair, you must use a different restore procedure

Before you begin

If currently enabled in your HA environment, disable Advanced Features on the active appliance by running the following command from the CLI:

```
> system logu disable
```

Disabling Advanced Features automatically disables the Monthly Index, Advanced Aggregation, and Monitoring Console features.

> **Note:** In HA environment, you can disable the Advanced Features only from the active appliance and only from the CLI. Disabling on the active appliance also disables all these features on the standby appliance.

Procedure

1. From the CLI, run the following command on the standby appliance to remove it from the HA pair:

   ```
   > set failover disable
   ```

2. Using the virtual IP of the HA pair, restore from the backup as described in Restoring an Appliance.

   After clicking **Restore**, if the following message is displayed, then wait for a few minutes:

   ```
   This site can't be reached
   ```

3. Run the following command on standby appliance to configure the standby appliance to rejoin the HA pair:

   ```
   > set failover configure
   ```

4. When prompted `This appliance shall be the destination of the initial`

`data migration`, enter Y.

Result

After the failover is configured and the migration completes, the HA pair is available with the restored data on both appliances.

What to do next

If you had disabled them before starting the procedure, reenable Advanced Features by running the following command on the active appliance:

```
> system logu enable
```

Note: In HA environment, you can enable the Advanced Features only from the CLI on the active appliance. Enabling on the active appliance also enables the feature on the standby appliance.

Enabling the Advanced Features does not automatically enable the Monthly Index, Advanced Aggregation, and Monitoring Console features. You must separately enable these features as required. See the system command.

Disaster Recovery Backup

The backup and restore features function similarly for a disaster recovery scenario as they do for either a single appliance or High Availability appliance pair.

The differences primarily involve the placement of the backup system, the method used, and the use of tape copies.

For disaster recovery, you typically set up the backup system at a remote location separate from the appliance being backed up. It is good practice to use the SCP backup method for disaster recovery situations because it is the more secure backup method. You must ensure that your network's security allows access to the remote location for successful connection between the appliance and backup system.

For optimal disaster recovery protection, copy the backup snapshot daily to magnetic tape, and secure the tapes in a separate location from the backup system.

Disaster Recovery Restore

Restore in a disaster recovery scenario functions similarly to restore for a single appliance or High Availability scenario. For more information, see the preceding sections on restoring in the relevant scenario.

Backup Recommendations

Read the following additional recommendations to get the best protection for appliance data:

- Run daily backups, scheduled for a time when network bandwidth and appliance CPU usage are both at their lowest.
- When configuring backup for a disaster recovery installation, backing up to a remote location, use SCP due to its better security.

  ⚠ **Warning:** The rsync version on the backup server must be higher or equal to the version running on the appliance. Otherwise the backup process fails. To check the rsync version, run the following command:

  ```
  rsync --version
  ```

- Move a copy of the backup snapshot daily to another, more reliable location such as a magnetic tape. For optimal disaster recovery protection, send the copy offsite to a secure data center.
- The following storage platforms are recommended for use with LogLogic backup:

Backup method	Storage platform
NFS	CentOS version 6.6
SCP/SSH	OpenSSH_7.4p1, OpenSSL 1.0.2k-fips

Note: Stop all data migration activities before you take a backup or perform restore operations.

Configuring SCP Backup

A one-time configuration is required before you run or schedule SCP backups. For other backup methods, this configuration is not required and you can directly schedule or run a backup.
Before you can use SCP backup for the first time, you must set up and test the SSH key on the appliance using the `system keycopy` CLI command. In a failover setup, you must perform the test on active and standby nodes.

Before you begin

Before setting up and testing the SSH key, ensure that you meet the following requirements:

- You must have the credentials and appropriate permissions to access a server that supports SCP, typically any UNIX or Linux system.
- `rsync` must be installed on the backup server.

Warning: The `rsync` version on the backup server must be higher or equal to the version running on the LogLogic LMI appliance. Otherwise, the backup process fails. To check the `rsync` version run the command:

```
rsync --version
```

Procedure

1. In the appliance CLI, copy the public SSH key of the appliance to the SCP server:
   a. Run the system `keycopy` command:

      ```
      > system keycopy
      ```

   b. When prompted whether to test or copy the key, enter C to copy the key.

The key is copied to the SCP server and its path is displayed.

2. Note down the displayed SCP server path where the key is copied. Later you need to append this file to `~/.ssh/authorized_keys` on the SCP server for the SCP account of the user (this must be identical to the user in step 4).

> **Note:** The actual directory that '~' maps to is different for each user, because the shell maps it to the user's home directory based on the username that is logged in.

3. (For LogLogic LMI 6.2.0 or earlier): Set the permissions of the `~/.ssh/authorized_keys` file to 600 by running the following command:

```
$ chmod 600 ~/.ssh/authorized_keys
```

Unless the file has permission 600, the files cannot be backed up to the server.

4. When prompted, enter the following information:
   - SCP server IP address (provided by your administrator)
   - SCP user name (provided by your administrator)
     confirm the displayed host IP address and RSA key fingerprint
   - password
5. Log in to the SCP server and append the appliance key to `-/.ssh/authorized_keys` on the server. For example:

```
SCP Server: IP-address login as: scpdata
============================================================
Machine Name:  sqalinux Owner: SQA Administrator Groups:
RE/SQA/Documentation Last Update: Mar 25, 2019
============================================================ SCP_
server:~> ls -l /tmp/LOGLOGICPUBKEY -rw-r--r--    1 scpdata  users
       611 2019-03-08 18:07 LOGLOGICPUBKEY SCP_server:~> cat
/tmp/LOGLOGICPUBKEY >> ~/.ssh/authorized_keys
```

If you want to configure multiple keys, you must perform separate copy operations using separate SSH sessions and enter the password during each session.

The SCP setup is complete.

6. Verify the SCP setup:

a. Run the system `keycopy` command:

```
> system keycopy
```

b. When prompted whether to test or copy the key, enter T to test the key.

c. When prompted, enter the SCP server IP address (provided by your administrator).

d. When prompted, enter the SCP user name (provided by your administrator).

The appliance copies a test file (`scptestfile`) to the SCP server and then copies it back to the appliance.

After the test copies complete successfully, a message is displayed.

What to do next

Schedule a backup or run a backup immediately. See Running Scheduled or Immediate Backups.

Running Scheduled or Immediate Backups

You can schedule a backup to run on specified days or run a backup immediately.

This section explains how to schedule or run a backup using the following methods:

- NFS
- SCP
- SAN (This option is available only on LogLogic ST SAN appliances.)
- Amazon S3

Before you begin

Before scheduling or running a backup, ensure that you meet the following prerequisites as per the backup method you are using:

Prerequisites for SCP backup

- You must set up the SSH key using the CLI command `system keycopy`. See Configuring SCP Backup.

- Ensure that the `rsync`, `df`, and `awk` utilities are installed on the SCP server. Without these, the backup works but free space on the remote server is not reported.
- The `rsync` version on the backup server must be higher or equal to the version running on the appliance. Otherwise the backup process fails. To check the `rsync` version, run the following command:

```
rsync --version
```

Prerequisites for NFS backup

- Create and export a volume on an NFS remote server.
- By default, the NetApp Snapshot feature is enabled on NetApp systems to copy data from an NFS storage location to the snapshot directory so that you can create restore points for files. You must disable NetApp Snapshot on your NetApp device because it affects some LogLogic LMI components via the mounted snapshot and is not supported in the LogLogic LMI appliance. For instructions about how to disable NetApp Snapshot, see article #000036898 on the TIBCO Support website.

Prerequisites for SAN backup

- Access to a LogLogic ST SAN appliance
- A Host Bus Adapter (HBA) installed on the LogLogic ST SAN appliance

Prerequisites for Amazon S3 backup

- Access to an Amazon S3 console with credentials, key ID, and secret key
- An Amazon S3 bucket to store the backup data
- The appliance system time must be precisely set to match the local timezone.

Procedure

1. In the appliance GUI, go to **Administration > Backup/Restore Configuration**.
2. Select the backup method:
   - NFS
   - SCP
   - SAN (This option is available only on LogLogic ST SAN appliances.)

- Amazon S3

A section with fields specific to the selected backup method is displayed.

3. (Optional) Select to **Optimize** the backups so only incremental backups are run instead of a full backup each time.

   An incremental backup copies only data that changed since the previous backup, while providing full restore capability if needed. This saves a lot of processing time for each backup after the initial backup for the appliance.

   Note: This option is not available for Amazon S3.

4. (Optional) When the **Bypass disk space checking** check box is not selected (Default), the system checks for available backup space prior to backing up your log files.

   Because the LogLogic application is very conservative, and because the total space required for backing up logs is difficult to accurately predict, the system might state that not enough storage space is available for backing up your log files when in fact you know that sufficient space is available. The default condition for this check box is cleared, meaning the appliance calculates (and possibly overestimates) the amount of backup space required.

   Selecting the **Bypass disk space checking** check box on the **Backup Configuration** pane prevents the appliance from possibly overestimating the amount of storage space necessary for a full backup of your log data.

   Note: This option is not available for Amazon S3.

5. (Optional) Select the **Config only** check box to back up only the configuration data of the appliance.

6. Depending upon the backup method you selected, enter the following information:

Backup method	Fields/tasks	Description
SCP	• User • Server	Enter the username and server IP address.

Backup method	Fields/tasks	Description
NFS	NFS Server	Enter the server IP address
SAN	In the Device pane, view the Host Bus Adapters (HBAs) and World Wide Port Number (WWPNs): • Hover over the HBA to reveal the local port number assigned to it. • Click + to reveal the disks attached to the HBA. • Hover over disk numbers to see attached labels. • Click a disk to display (in the Information pane) its operational status, size, and the Linux device mapped to it. (The UUID and Label fields are not used.) • Click + again to reveal the disk partitions and sizes. Hover over partitions to see attached labels. • Click a disk partition to display (in the Information pane) its operational status, Universally Unique Identifier (UUID), label, size, and Linux device mapped to it. The partition is what the user must choose to back up to, not the actual disks themselves.	The **Current Configuration** window shows the target disk and partition selected for backing up to. If the window is blank, no disk and partition have been selected yet. To select or change the target disk and partition, click a partition in the Device pane and roll your mouse over it to reveal the presence of a label. If no label appears, check the partition size in the Information window. If the partition size is adequate for backing up your log data, click the **Label** button above the Device pane, and enter a label for your backup. Click **OK** to close the window, or **Cancel** to close without saving.

Backup method	Fields/tasks	Description
Amazon S3	Identity	Select Identity and Access Management (IAM) User or IAM Role as per the permissions assigned to you by the administrator.
	• AWS Access Key ID • AWS Secret Access Key	Enter the AWS access key ID and secret access key of the AWS console.
	Role ARN	Enter the Amazon Resource Name (ARN). This field is displayed only if you select identity as **IAM Role**.
	Amazon S3 Bucket Name	Enter the bucket name where you want to back up the appliance data.

7. Set the backup to automatically run as per a schedule or to run immediately:
   - Select the recurring **Backup Schedule** days and time.

     Select the days of the week on which you want automatic backups to occur, or select **Everyday**.

     Set **Backup Time** for the automatic backups to begin on the selected days. Appliance data is backed up on the selected days and at the beginning of the scheduled hour.
   - If you want to run a backup immediately, select **Start Backup Immediately**.
8. Select the number of **Backups To Retain** that can be available later for restoring. The maximum is 30.
9. Click **Update** to save the backup configuration.

Result

After the backup is completed, the following message is displayed:

```
Backup configuration update completed
```

For SAN backups, if any of the following messages is displayed, select a different partition and try performing the backup again:

- SAN Config: The selected device can't be used for SAN.
- SAN Config: The selected device conflicts with an existing backup/restore configuration

To avoid accidental overwriting by another appliance, the backup process creates an additional folder `ll_bkup_<ip-address>` using the IP address of the appliance. Depending on the requirements of the backup server file system, the <*ip-address*> might be either the IP address using decimal points, or a format where colon characters are replaced by other characters. All backup sessions for an appliance are contained in its corresponding backup folder. Examples of folders created on NFS, SCP, or SAN servers:

```
$ ls -l
total 24
drwxr-xr-x    7 root root 4096 Mar 30 11:10 .
drwxr-xr-x    3 root root 4096 Mar 27 17:11 ..
drwxr-xr-x   14 root root 4096 Mar 27 19:12 ll_bkup_session_1214586678_
1238181143
drwxr-xr-x   14 root root 4096 Mar 27 21:27 ll_bkup_session_1238188349_
1238189277
drwxr-xr-x   14 root root 4096 Mar 28 11:10 ll_bkup_session_1236038400_
1238238608
drwxr-xr-x   14 root root 4096 Mar 29 11:09 ll_bkup_session_1238324400_
1238324960
drwxr-xr-x   14 root root 4096 Mar 30 11:10 ll_bkup_session_1238410800_
1238411403
```

What to do next

(Optional) To verify if the backup directory has been created on the backup server or Amazon S3 bucket:

1. Log in to the backup server or Amazon S3 bucket.
2. Open the backup folder `ll_bkup_<ip-address>` for the selected appliance and view the files.

Monitoring Backup Status

Monitor the progress of a current backup or review a recent backup from the **Administration > Backup/Restore Configuration > Backup Status** tab.

The **Backup Status** tab displays the latest backup processes and their details. The status of a backup in progress is updated every 20 seconds.

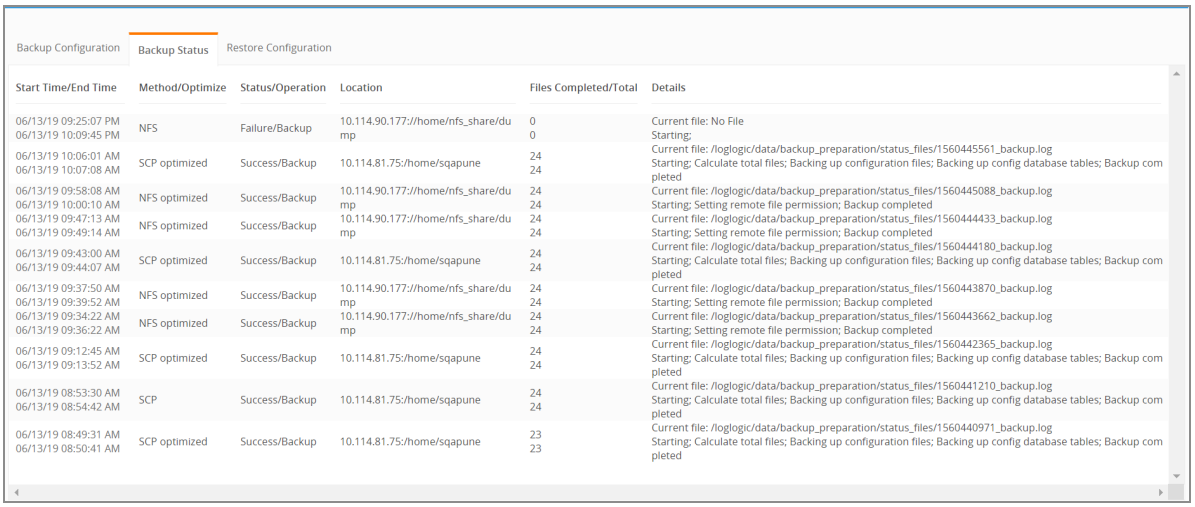

Backup Configuration | Backup Status | Restore Configuration

Start Time/End Time	Method/Optimize	Status/Operation	Location	Files Completed/Total	Details
06/13/19 09:25:07 PM 06/13/19 10:09:45 PM	NFS	Failure/Backup	10.114.90.177://home/nfs_share/dump	0 0	Current file: No File Starting;
06/13/19 10:06:01 AM 06/13/19 10:07:08 AM	SCP optimized	Success/Backup	10.114.81.75:/home/sqapune	24 24	Current file: /loglogic/data/backup_preparation/status_files/1560445561_backup.log Starting; Calculate total files; Backing up configuration files; Backing up config database tables; Backup completed
06/13/19 09:58:08 AM 06/13/19 10:00:10 AM	NFS optimized	Success/Backup	10.114.90.177://home/nfs_share/dump	24 24	Current file: /loglogic/data/backup_preparation/status_files/1560445088_backup.log Starting; Setting remote file permission; Backup completed
06/13/19 09:47:13 AM 06/13/19 09:49:14 AM	NFS optimized	Success/Backup	10.114.90.177://home/nfs_share/dump	24 24	Current file: /loglogic/data/backup_preparation/status_files/1560444433_backup.log Starting; Setting remote file permission; Backup completed
06/13/19 09:43:00 AM 06/13/19 09:44:07 AM	SCP optimized	Success/Backup	10.114.81.75:/home/sqapune	24 24	Current file: /loglogic/data/backup_preparation/status_files/1560444180_backup.log Starting; Calculate total files; Backing up configuration files; Backing up config database tables; Backup completed
06/13/19 09:37:50 AM 06/13/19 09:39:52 AM	NFS optimized	Success/Backup	10.114.90.177://home/nfs_share/dump	24 24	Current file: /loglogic/data/backup_preparation/status_files/1560443870_backup.log Starting; Setting remote file permission; Backup completed
06/13/19 09:34:22 AM 06/13/19 09:36:22 AM	NFS optimized	Success/Backup	10.114.90.177://home/nfs_share/dump	24 24	Current file: /loglogic/data/backup_preparation/status_files/1560443662_backup.log Starting; Setting remote file permission; Backup completed
06/13/19 09:12:45 AM 06/13/19 09:13:52 AM	SCP optimized	Success/Backup	10.114.81.75:/home/sqapune	24 24	Current file: /loglogic/data/backup_preparation/status_files/1560442365_backup.log Starting; Calculate total files; Backing up configuration files; Backing up config database tables; Backup completed
06/13/19 08:53:30 AM 06/13/19 08:54:42 AM	SCP	Success/Backup	10.114.81.75:/home/sqapune	24 24	Current file: /loglogic/data/backup_preparation/status_files/1560441210_backup.log Starting; Calculate total files; Backing up configuration files; Backing up config database tables; Backup completed
06/13/19 08:49:31 AM 06/13/19 08:50:41 AM	SCP optimized	Success/Backup	10.114.81.75:/home/sqapune	23 23	Current file: /loglogic/data/backup_preparation/status_files/1560440971_backup.log Starting; Calculate total files; Backing up configuration files; Backing up config database tables; Backup completed

The **Status/Operation** column indicates the status of the operation and the type of the operation - backup or restore.

The **Location** column displays the backup or restore location.

Note: In case of SAN backup, the location of the previous backup might be displayed in the **Location** column.

The **Files Completed/Total** column displays the number of files backed up and the total files. However, after restoring a full backup, the value changes to 0/0 even though the restore is complete.

The **Details** column indicates the phases of the backup process as they progress. These are high-level process phases, similar to the comment fields that appear in the backup log.

Note: Sometimes you might see a message, CAN NOT FIND EXACT MATCH MIDNIGHT DIRECTORY, in the **Details** column. This message is harmless and can be ignored.

To view the step-by-step detailed log of a listed backup process, including one currently processing, click its **Start Time End Time** entry. The complete log of the selected process appears.

Backup Status Detail Example

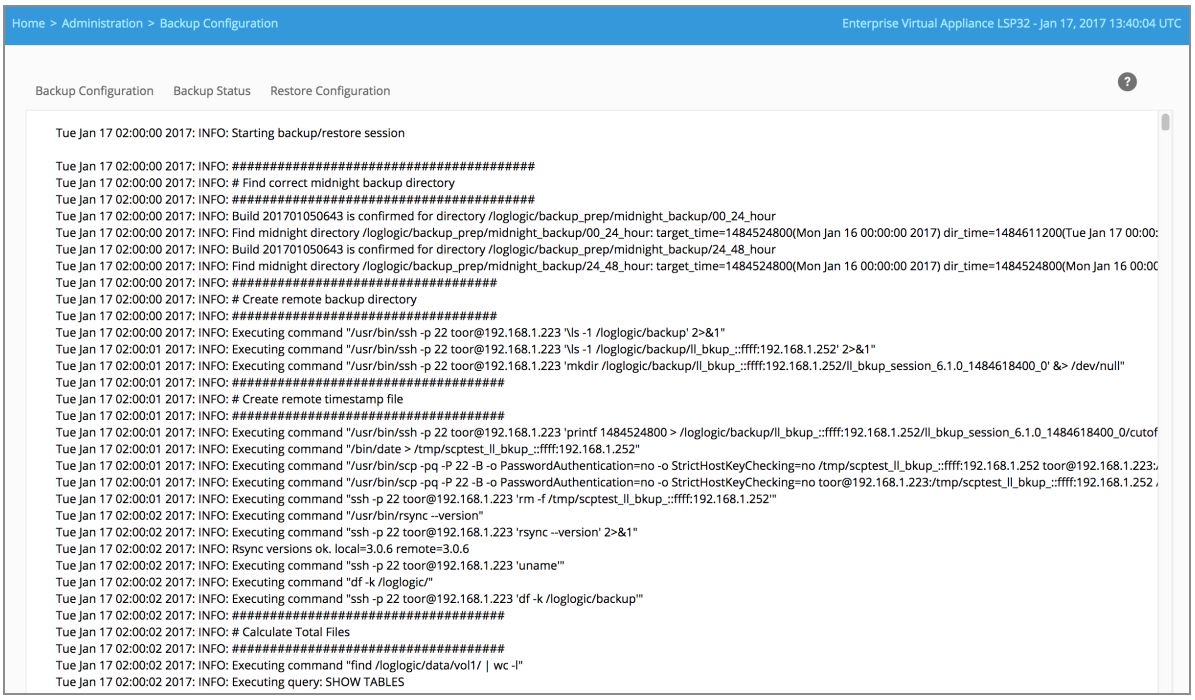

When the backup is completed, the end of the status details includes entries similar to:

```
Tue Jan 17 02:02:11 2017: INFO: ####################################

Tue Jan 17 02:02:11 2017: INFO: # Trim excess backup on remote host

Tue Jan 17 02:02:11 2017: INFO: ####################################

Tue Jan 17 02:02:11 2017: INFO: Executing command "/usr/bin/ssh -p 22
toor@192.168.1.223 '\ls -1 /loglogic/backup/ll_bkup_::ffff:192.168.1.252'
2>&1"

Tue Jan 17 02:02:11 2017: INFO: Executing command "/usr/bin/ssh -p 22
toor@192.168.1.223 'mv /loglogic/backup/ll_bkup_::ffff:192.168.1.252/ll_bkup_
session_6.1.0_1484618400_0 /loglogic/backup/ll_bkup_::ffff:192.168.1.252/ll_
bkup_session_6.1.0_1484618400_1484618531' &> /dev/null"

Tue Jan 17 02:02:12 2017: INFO: ####################################

Tue Jan 17 02:02:12 2017: INFO: # Backup Completed

Tue Jan 17 02:02:12 2017: INFO: ####################################

Tue Jan 17 02:02:12 2017: INFO: Executing command "rm -rf /loglogic/backup_
prep/staging &> /dev/null"
```

Backup Errors

System utility errors that occur during backup are reported on the **Backup Status** tab.

Because the backup is live, data is modified while the backup is running. This might cause some system utility errors which are reported in the backup status.

Whenever the backup process encounters an error, the backup process stops; it does not retry the action that caused the error. If the appliance repeatedly reports errors during backup, run the backup at a different time when the CPU and network bandwidth are less busy.

During an Amazon S3 backup, if the Amazon Web Service (AWS) CLI used for backup faces congestion issues with the Amazon authentication server, error messages are displayed in the **Details** column of the **Backup Status** tab, and the backup process terminates. Try to perform the backup after some time.

Note: The system utility error “Partial transfer due to vanished sources files” is expected because source files are archived or purged due to retention, often as part of appliance maintenance. If the backup runs during the regular appliance maintenance cycle, it encounters this error.

Restoring an Appliance

You can restore an available backup to a LogLogic LMI appliance from the **Restore Configuration** tab.

On the **Restore Configuration** tab, the **Source** list includes all available backups from which a restore can be performed. If no backups are available, only `NULL` is displayed in the list.

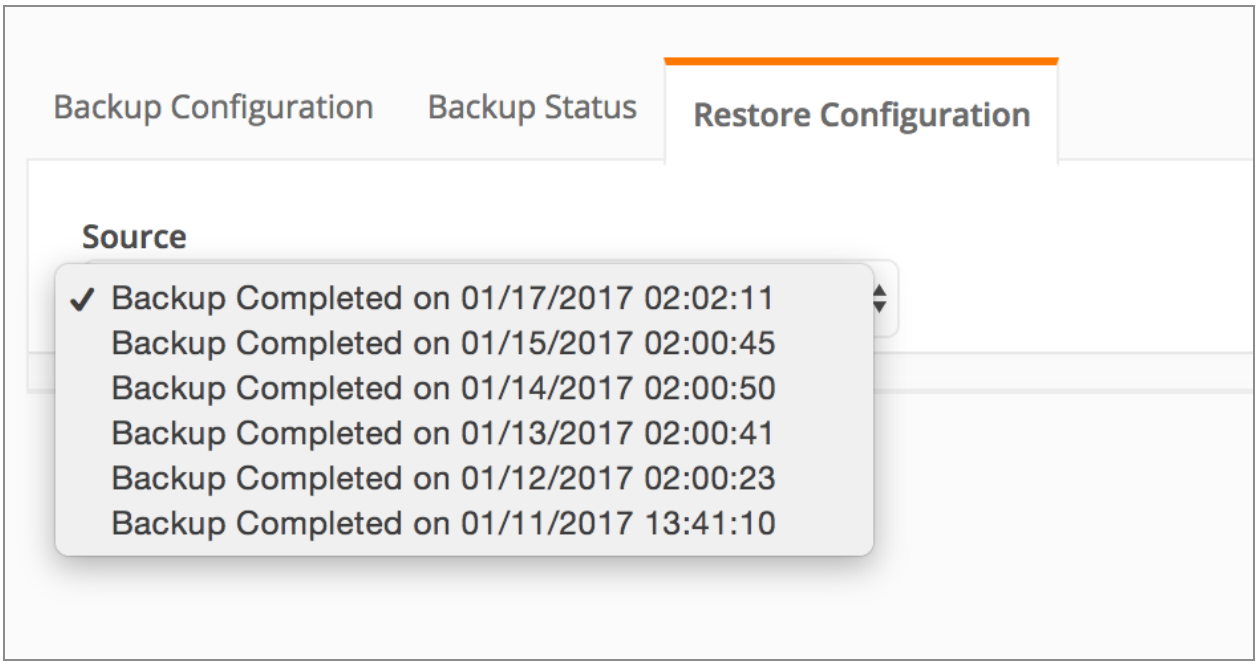

If you select the **Config only** check box on the **Backup Configuration** tab, the Source list displays a list of the "config only" backups of that backup method. If the **Config only** check box is not selected, the list displays all full backups of that backup method.

Important Considerations

- You can only restore a backup to an appliance for which a backup configuration is already defined and for which the backup configuration is identical to that on the backed-up appliance. For example, if the backup location was `/home/john/lmi-backup` and options such as **Optimize** or **Bypass disk space checking** were used on the backed-up appliance, the same location and options must be used while configuring backup on the replacement appliance.
- If the **Backup Method** field on the **Backup Configuration** tab is set to None, you cannot restore any data even if recently backed-up data is available.
- You cannot restore from a backup that was taken with an older version of LogLogic LMI than the current version.
- Performing a restore overwrites the current log and configuration data on the appliance with the data stored in the backup. Data is automatically restored using the method (NFS or SCP) in which it was backed up.
- Data can be restored to any appliance, including the appliance from which it was

backed up. However, if your intention is to move data between appliances, it is good practice to use the data migration feature and not the backup and restore feature. For more information about data migration, see Data Migration Between Appliances.

Before you begin

Before you start restoring data, ensure that you meet the following prerequisites:

- The `rsync` version on the backup server must be higher or equal to the version running on the appliance. Otherwise the backup process fails. To check the rsync version, run the command: `rsync --version`.
- When restoring to a new appliance, it must have the same IP address of the appliance from which you originally ran the backup.

Procedure

1. Go to the **Administration > Backup/Restore Configuration > Restore Configuration** tab.
2. From the **Source** list, select the backup from which you want to restore appliance data.
3. Click **Restore**.
4. Click **Confirm** on the confirmation dialog box.

Result

After restore is successful, the following message is displayed on the page:

```
Restore successful
Click here to redirect back to application.
```

Click the link and log in to the appliance.

However, if multiple interfaces are configured on the appliance, of which, the eth0 interface is set to the default IP address (10.0.0.11), the following error is displayed:

```
The connection has timed out
```

This occurs because the appliance redirects to the default IP address (that is assigned to the eth0 interface) instead of the IP address that you had used for restoring. In such case, log in to the appliance from the IP address that you had used for restoring.

Successful restore operations are not listed on the **Backup Status** tab. If an error occurs during the restore operation, details of the error are displayed on this tab.

After restoring data on a Remote Appliance that is accessed via a Management Station, you can continue to use the Management Station while the restore is in progress. However, the Remote Appliance restarts and is not accessible until the restore is complete.

Data Encryption

The data vault feature manages encryption of all data volumes including archives. By default, the data vault is disabled and the data volumes are in unlocked state.

Enabling the data vault feature begins the encryption of data volumes; but after encryption is complete, the data volumes are unlocked and are accessible to users. An administrator logged in via the CLI as `root` has the rights to enable or unlock the data vault, to check the status of the data volumes, to change the password of the data vault, or to enable or disable the auto-unlock option of the data vault by running the `system data_vault` command.

> ⚠ **Caution:** Exercise caution before enabling the data vault feature. Once enabled, it you cannot disable it. Also, you cannot migrate data when this feature is enabled.

If the system is restarted for any purpose by any user, the data volumes are locked. If N was entered at the prompt Save the password to automatically decrypt the data on boot time? (y/N) while enabling the data vault feature, an administrator user must run the `data_vault` command to unlock the data volumes to resume using the data volumes. Whereas, if y was entered at the prompt while enabling the data vault feature, the data volumes are automatically unlocked after the system restarts.

Once the data vault is enabled, the local volume and the remote archive storage are encrypted, including the existing archive mount point and the archive mount point that is added after turning on the data vault feature.

After turning the data vault feature on, only the new data on the remote archive is encrypted. To encrypt existing data on the remote archive before turning the data vault on, you must run the script `/loglogic/scripts/dv_convert.py`. However, no additional action is required in case of local storage.

From LogLogic LMI 6.3.0 onwards, encryption of remote archive volumes is optional while encrypting data using the data vault feature.

Important Considerations

- Once enabled, you cannot disable the data vault feature.
- Data migration is not supported when the data vault feature is enabled.
- Enabling the data vault takes some time depending on the appliance model configuration. To avoid data loss, do not switch off the appliance while the enabling is in progress. Data lost in such a scenario cannot be recovered.

Purge Stale Devices

LogLogic LMI automatically purges stale devices based on the last time data was received.

By default, the feature is off. To enable the feature, click the Purge Stale Devices slider to **On** and in the **Purge Stale Devices Period** field, specify the number of days after which devices should be purged.

This setting applies to Syslog devices, and does not apply to file log devices. A device is purged after its state changes from active to stale. If a device exists but has never received logs, the device is not purged and remains in the system.

This setting triggers purging of the device from the system, and not the log data of the device. After the device is purged, its log data remains in the system but is not searchable if the query includes the specific device. To search the log data of a purged device, the query must search all data in the system or include the device type. If you add the same device again to the system, you can search its log data.

Viewing Log Files

You can view all log files on the **Administration > Data Files** page.
Because data files are compressed, you must save them to your local computer and decompress them for viewing. You can download and view a data file in the raw format or the ASCII format.

Procedure

1. To limit the list of data files by time, select the year, month, day, and time to view.
2. Click the raw or ASCII link next to the data file you want to view.

   The compressed file is downloaded to your computer.
3. Decompress the downloaded file and view it in a text editor.

Verifying the SHA Digest on Data Files

You can ensure the integrity of your data files by verifying that the SHA Digest has not changed since the LogLogic LMI appliance captured the data.
Use the **Data Files** page to verify the SHA Digest. (See the CLI section for information on setting the appliance Digest.)

> **Note:** To further ensure the integrity of data on an LogLogic ST Appliance, consider using a WORM (write once read many) storage server such as Network Appliance SnapLock.

Procedure

1. In the navigation menu, click **Administration > Data Files**.

   The **Data Files** page appears.
2. Click the check box to the left of each data file to verify.
3. Click **Verify** to start the verification process.

When verification completes, refresh the page. A flag and verification date appear in the **Digest Verified** column:

- A green flag indicates successful verification of the Digest of the data files. The timestamp next to the green flag identifies the date and time the verification succeeded.
- A red flag indicates failed verification. Mouse over the failure message for more information on the reason. A failure can mean any of the following:
    - The file was modified. Mouse over the failure to view the new Digest.
    - The file is no longer accessible. The file might be inaccessible for various reasons such as the location of the file has changed or the network connection is down and your file is on a storage server such as NFS volume.

Note: The verification process has a low priority in the appliance. If the system is busy processing log data, the verification process might take longer than expected.

Archiving of Log Data (LogLogic ST Appliance and LogLogic EVA Only)

To extend the storage capacity of a LogLogic ST Appliance or LogLogic EVA, you can enable archiving to external storage. This lets you archive log data to a remote storage device to free up the local storage space on the appliance to use for incoming log data.

From the **Administration > Archive Configuration** page, you can attach a storage volume from a SAN or NFS remote server to a LogLogic ST Appliance; or one or more volumes from an NFS remote server to a LogLogic EVA instance.

When configured with a LogLogic ST Appliance or LogLogic EVA instance, an attached volume from an NFS remote server acts as a data storage extension where log data can be accessed as if it were still on the appliance. Therefore, the same access to reporting, indexing, searching, and alerting exists on the storage server as on the appliance itself.

Starting release 6.3.0, in addition to EXT2 and EXT3 file systems, you can use XFS and EXT4 file systems on SAN devices for archiving data.

How Archive Storage Works

The archive process copies files to the attached data volume of the SAN or NFS remote server or Amazon S3 buckets, and updates the record in the database.

Archival to NFS remote server

If the Direct Attached Storage (DAS) capacity reaches the watermark set in the Archive Threshold field on the Archive Configuration page, or if the files reach their retention time limit, the files are deleted.

By default, the NetApp Snapshot feature is enabled on NetApp systems to copy data from an NFS storage location to the snapshot directory so that you can create restore points for files. You must disable NetApp Snapshot on your NetApp device because it affects some LogLogic LMI components via the mounted snapshot and is not supported in the LogLogic

LMI appliance. For instructions about how to disable NetApp Snapshot, see article #000036898 on the TIBCO Support website.

Note:

- Set the archive threshold to a non-zero positive number. For best results, use a number greater than 10.
- In case of the attached NFS data volume, LogLogic LMI owns the contents of the archival top directory, if one archive volume is:`1.2.3.4:/ExportedShare` (on the command line it appears as `/loglogic/data/archive/mount_#/vol1`), then every file under `.../archive/mount_#/vol1` comes from the archival process.
- If the connection with the NFS servers is lost or the NFS server is shut down after configuring the NFS archive server, the LogLogic LMI appliance cannot be shut down or restarted for a timeout period of about 30 minutes. After the appliance restarts, restart the NFS server and run the mount command. For a workaround or a resolution, see this article on the Red Hat website.

In HA mode, the node ID might have to be set up to avoid mount point conflicts due to `multipathd` being enabled. See Enabling a Node ID in the Archive Path.

Archival to Amazon S3 buckets

You can archive data to Amazon S3 buckets in AWS cloud by setting up a Storage Gateway of the type "File". Archiving to an Amazon S3 bucket is identical to archiving to an NFS remote storage. To configure the file gateway:

1. On AWS cloud, follow the procedure Creating an NFS File Share in the *AWS Storage Gateway User Guide*. In the **Create file share** section, ensure that you set **Squash level** to No Root Squash.
2. On the LogLogic LMI appliance, configure the archival from the **Administration > Archive Configuration** page. Follow the procedure in Configuring NFS Remote Storage by selecting the **Enable NFS Remote Storage** option.

Storage Volume Watermarks

Pertaining to the attached volumes from the NFS remote server, SAN storage, and LogLogic LMI local disk storage, the watermark thresholds in the `/loglogic/conf/archive_config` file determine how log storage is treated on these devices.

These watermark thresholds are described in the following sections:

- localVolHighWatermark
- localVolHighArchiveWatermark
- nasVolLowWatermark

localVolHighWatermark

The `localVolHighWatermark` watermark pertains to the appliance's local volume. If the percentage of used local disk space (on the `/loglogic` partition) exceeds the specified watermark threshold, the appliance starts incrementally deleting the files with the earliest expiration date from the DAS to accommodate new data.

If everything is set to the same expiration policy (for example, 1 year), then the oldest file is purged. However, when an appliance is configured with multiple different retention policies, files with the earliest expiration date are purged first. For example, files that are set to expire in one week are deleted under emergency purging circumstances before files that expire after one year and 10 months.

Typically, the `localVolHighWatermark` threshold is not reached when volumes from the NFS remote server or SAN storage are attached. This is because another watermark threshold, `localVolHighArchiveWatermark`, would be reached first and would trigger the `engine_archive` to start archiving files off the DAS and over to the attached volumes on the NFS remote server or SAN storage, thus keeping the local DAS disk space relatively free.

However, if the NFS remote storage or SAN storage is not configured properly or if the network is not available, it is possible that the local logs are not archived (moved) and the local `/loglogic` partition can become full. Therefore this `localVolHighWatermark` watermark prevents an undesirable full partition by beginning the file deletion process once the watermark is reached.

This watermark value can be adjusted from 0% to 99%.

localVolHighArchiveWatermark

The *localVolHighArchiveWatermark* watermark is the other watermark threshold pertaining to the appliance local volume DAS. If the percentage of used local disk space (on the `/loglogic` partition) exceeds this specified threshold, the appliance starts incrementally archiving the oldest files to the attached volumes on the NFS remote server or SAN storage. Although files are copied temporarily, they ultimately only exist on archive storage after all operations are performed for a given file, thus freeing up space on the DAS. This is considered an effective move operation because a copy operation is used to transfer data to the archive storage, followed by a delete operation to remove the local copies of the file. This ensures a safer data transfer than using a formal move command.

This watermark value can be adjusted from 2% to 94% with a default level of 75%.

Note: Immediate Archiving: Setting this watermark to a very low value (`localVolHighArchiveWatermark=2`) triggers near-immediate archiving when the volumes from the NFS remote server or SAN storage are successfully mounted. Although the lowest allowed value is 2%, the recommendation is minimum 5%, which can still provide near-immediate archiving capability after files are initially created on the local disk.

Setting the archiving threshold to a value below 5% might increase the chance of interference between the indexer and the archiver processes, and `FileNotFoundException` messages in `indexer.log` might be displayed. This is because, as soon as files are collected in the LogLogic LMI appliance, the archiving process starts moving files while the indexer is trying to access them. The lower the archiving threshold, the greater is the chance of interference.

nasVolLowWatermark

The *nasVolLowWatermark* threshold watermark (in gigabytes) pertains to the available disk space on the attached volumes on the NFS remote server or SAN storage. If the available disk space on the attached volumes goes under this specified threshold, the appliance starts incrementally deleting the files with the earliest expiration date on the attached volumes to accommodate new data.

This value of the *nasVolLowWatermark* watermark threshold can be adjusted from 10 GB to 1024 GB with a default level of 20 GB x the number of mounting points set (up to 32). The value is specified in gigabytes instead of a percentage number. This is because the calculated free space of the attached volumes from the NFS remote server and SAN storage

would vary greatly. For example, if the user has 2 mounting points for NAS, the *nasVolLowWatermark* value is changed to 2x20 GB=40 GB instead of only 20 GB.

Warning: Up to 256TB of aggregate archive space is supported across up to 32 mount points. If you have already connected the appliance to a volume on an NFS remote server or SAN server and then want to change the threshold or the mount directory, you will no longer have access to your stored data.

By reverting the configuration to the old archive storage information, you can access the data again.

Warning: On the NFS remote server, you must create a directory as a file share for the appliance to mount, with `rw`, `no_root_squash` access permission.

NAS SnapLock Protection

For additional protection, LogLogic lets you configure the appliance to use the NetApp SnapLock® software.

If you enable SnapLock for your NetApp NAS Server, all archived raw data is protected and cannot be modified for the time period defined by your data retention settings. Index files are not locked. If SnapLock is enabled, you cannot purge protected files from the NAS server, so make sure you have enough available space to handle the archiving from the appliance.

Note: If you reduce or extend the data retention settings for your archived data, the new retention time applies only to the new files that are SnapLock enabled. All archived data files on the NAS prior to the change in retention time retain their existing retention time settings.

External Storage in an HA Pair

External storage for LogLogic ST Appliances and LogLogic EVA is not replicated as part of the failover configuration.

When an appliance fails over, its external storage is automatically switched over to the standby appliance as it becomes active.

Enabling a Node ID in the Archive Path

> ⚠ **Warning:** This setting is not persisted on upgrades if enabled through LogLogic LMI 5.6.2HF3 or later, or 5.7.0.
>
> Starting LogLogic LMI 6.1.0, archiving with and without the node ID is supported. However, the setting is disabled for new installations and after an upgrade. Hence, if you have been using this feature in prior versions, you must create the `archive_config` file before upgrading to LogLogic LMI 6.1.0 or later.
>
> If you were not using this feature in prior versions and want to use it after the upgrade, you must re-enable the setting after the upgrade is complete.

For more information and instructions on configuring external storage for an LogLogic ST Appliance or LogLogic EVA in an HA pair, see Failover and External Storage.

Procedure

1. Create the archive configuration file `archive_config` under the `/loglogic/conf` directory.
2. Add the following line to the file:

   ```
   archivedPathWithId=1
   ```
3. Restart the `engine_archive`.

Configuring NFS Remote Storage

You can use an NFS server in conjunction with a LogLogic ST Appliance or LogLogic EVA as a dedicated storage system for raw data files. See NAS SnapLock Protection.

Before you begin

By default, the NetApp Snapshot feature is enabled on NetApp systems to copy data from an NFS storage location to the snapshot directory so that you can create restore points for files. You must disable NetApp Snapshot on your NetApp device because it affects some

LogLogic LMI components via the mounted snapshot and is not supported in the LogLogic LMI appliance. For instructions about how to disable NetApp Snapshot, see article #000036898 on the TIBCO Support website.

Procedure

1. Go to **Administration > Archive Configuration**.
2. On the Archive Configuration page, select the **Enable NFS Remote Storage** option to enable the configuration settings for the NFS remote server.

   > **Note: Archive Threshold** relates to local storage capacity of the appliance. The user can change the default setting (75% of max) at which point the oldest files stored on the appliance begin moving to the external archive storage device (until the Archive Threshold value is met again).
   >
   > Setting the archiving threshold to a value below 5% might increase the chance of interference between the indexer and the archiver processes, and `FileNotFoundException` messages in `indexer.log` might be displayed. This is because, as soon as files are collected in the LogLogic LMI appliance, the archiving process starts moving files while the indexer is trying to access them. The lower the archiving threshold, the greater is the chance of interference.

3. Create and export a directory to the NFS remote server.
4. In the **NFS Configuration** section, click **Add** and enter the IP address of the remote server that has the NFS export directory.

   The appliance can be connected to up to 32 mount points.

   > **Note:** The IP address field is case sensitive and format specific. If the same file name and path exists, the data is overwritten. If you designate a partition and change it later, you might not be able to access the files you saved on the old partition. In addition, if this is the first time you are specifying an archive server, make sure the mount directory is empty.

   For example, the format for the IP address and mount directory is:
   `xxx.yyy.zzz.ddd:/ExportedShare`

   where:

   - `xxx.yyy.zzz.ddd` is the IP address of the NFS remote server

- `ExportedShare` is the mount directory

5. Select the **Enable SnapLock** check box to let the appliance archive data to a NetApp Server with SnapLock.
6. Click **Update** to save your settings.

Configuring SAN Archive Storage

A Storage Area Network (SAN) device can be attached to the appliance as a dedicated and secure storage system for your log data files, providing a suitable Host Bus Adapter is fitted to your appliance.
When using the archive feature in a HA configuration ensure that SCSI reservations have been enabled on your SAN appliance.

Procedure

1. In the navigation menu on the appliance, select **Administration > Archive Configuration**.
2. Select the **Enable SAN** radio button.
3. In the **Device** pane, view the Host Bus Adapters (HBAs) and World Wide Port Number (WWPNs).
    a. Rollover the HBA to reveal the local port number assigned to it.
    b. Click **+** to reveal the disks attached to the HBA. Rollover disk numbers to see attached labels.

    > **Note:** If the SAN administrator who published the LUN for you did not partition it, the + is not displayed. To partition it, see the Knowledge Base article #000029103

    c. Click a disk to display (in the **Information** pane) its operational Status, Size, and Linux Device mapped to it. (The UUID and Label fields are not used.)
    d. Click **+** again to reveal the disk partitions and sizes. Rollover partitions to see attached labels.
    e. Click a disk partition to display (in the **Information** pane) its operational Status, Universally Unique Identifier (UUID), Label, Size, and Linux Device

mapped to it. The partition is what the user must choose to back up to, not the actual disks themselves.

f. The Current Configuration window shows the target Disk and Partition selected for backing up to. If the window is blank, no Disk and Partition have been selected yet.

g. To select or change the target Disk and Partition, click a partition in the Device window and roll your mouse over it to reveal the presence of a label. If no label appears, check the partition size in the Information window.

h. If the partition size is adequate for backing up your log data, click the **Label** icon over the Device window, and enter a label for your backup. Click **OK** to close the window, or **Cancel**.

i. (Optional) You can remove one or more partitions (mount points) by clicking on them in the Current Configuration pane and then clicking Remove (red **X**). A warning message appears, advising that once removed, any data stored on the selected partitions is lost and no longer retrievable.

j. Click the **Update** button at the bottom of the page. The message: `Update Completed` appears at the top of the Archive Configuration page, just under the tab.

k. However, if you see any of the following messages, select another partition and try again:

```
SAN Config: The selected device can't be used for SAN
SAN Config: The selected device conflicts with an existing
archive/restore configuration
```

To disable archive on this appliance, under **Archive Configuration** select **None**.

> ⚠ **Warning:** If the user's SAN archive system goes offline for more than 60 seconds, it is necessary to reboot the appliance after the SAN archive system is restored.

Supported Cable Distances

Some cable distances and types that can be used with the TIBCO LogLogic® ST2025-SANR1 Appliance and TIBCO LogLogic® ST2025-SANR2 Appliance appliances.

Supported Cable Distances and Type

Rate	Cable Type and Distance (meters)		
	OM1	OM2	OM3
2 Gbps	150	300	500
4 Gbps	70	150	380
8 Gbps	121	50	150

Data Retention Rules

By configuring data retention rules, an administrator can manage the time duration for which data can be retained on the appliance.

Multiple Data Retention rules can be defined for managing data.

The Raw data retention time is the duration for which the data is retained on the appliance. The Indexed data retention time is the duration for which the raw data is indexed for searching.

Use the **Administration > Data Retention Rules** tab to define Data Retention rules. For each rule, you can specify the retention time period for raw and indexed data. The maximum value of indexed data retention can be 10 years. Log sources should be assigned to a specific Retention rule.

> ⚠ **Warning:** You must have the **System Configuration** privileges and the **Access all devices in the appliance** check box enabled to manage Data Retention rules for any appliance.

During installation, some pre-defined Retention rules (one Default and multiple Custom rules) are created. The number of pre-defined Custom rules may vary depending on each appliance model. You can create new Custom Retention rules. The Custom rules are prioritized in the order as they appear (from highest on the top) in the Custom rules list. You can change the priority by moving them up or down in the Custom Rules list. See Prioritizing Custom Rules.

GUI option	Description	For more information, see...
View All Rules	Lists all log sources and their effective rules. You can view the effective rule for a particular log source.	Viewing Retention Rule Details
Custom	Specifies the raw and indexed data retention time for log	Assigning Log

GUI option	Description	For more information, see...
Retention Rules	sources assigned to the custom rule.	Sources to a Data Retention Rule
Default Retention Rule	Specifies the raw and indexed data retention time for log sources that have not been assigned to any custom rule. If any log source is not assigned to Custom Retention rule, it is automatically assigned to the Default Retention rule. You can modify the time period; however, you cannot delete this rule.	Viewing Retention Rule Details

During installation, the following default period is set for raw data retention:

Appliance	Default period of raw data retention
LogLogic LX Appliance	90 days
LogLogic MX Appliance	365 days
LogLogic ST Appliance	365 days

Viewing Retention Rule Details

The **Custom rules** section specifies the raw and indexed data retention time for log sources assigned to the custom rule. You can create, modify, or delete custom rules.

The **Default Data Retention Rule** section applies to all log sources that are not assigned to any custom rules. You can modify the time period; however, you cannot delete this rule.

Procedure

1. Navigate to **Administration > Data Retention Rules**. The Data Retention window appears.

The left pane of the screen displays all rules including the default rule, and pre-defined custom rules. Once you select a rule, the right pane displays the rule details.

2. To view all rules, click **View All Rules** in the left pane. You can view every log source and its effective rule in the right pane.
3. From the left pane, select a rule. The right pane displays the rule description, and a list of all log sources assigned to the selected rule.

   You can filter the device list by **Groups and devices** or **Only groups**. You can also filter the list by entering text in the **Find** field, and press **Enter**. The filtered list containing the search term is displayed.
4. Clear the search term in the **Find** field and press **Enter** to see all log sources.

   You can resize and move the columns to the positions you prefer by clicking on them and dragging. To sort the column in descending or ascending order, click the column header.

Creating a New Retention Rule

You can create new retention rules to specify the retention time period for specified log sources.

Procedure

1. Access **Administration > Data Retention Rules** from the navigation menu.
2. Click the **Create New** button from the left pane, to create a new retention rule. The New Retention Rule window appears.
3. Enter the **Name** and **Description** of the rule.
4. In the **Raw Data Retention time** section, select the time period for the raw data from the **months**, and **days** lists.

   Note: The maximum raw data retention period for all appliance models is 10 years.
5. To enable indexing, select the **Index the data (for searching)** check box.

   The **Indexed Data Retention time** section is displayed.

6. Select the time period for retaining the indexed data.

   Warning: You must select the Data Retention time period from the list. The **Indexed Data Retention time** period cannot be more than **Raw Data Retention time** period.

7. Click **Save** to save the rule settings. The new rule is now listed in the **Custom Rules** list in the left pane.
8. Click the **Commit Changes** button to commit the latest changes. Click the **Revert Changes** button to revert to the last committed changes.

   Note: After you make **all** changes, click the **Commit Changes** button to commit to all changes. After you click the **Commit Changes** button, the new data is stored as per the new rule.

What to do next

After you create a new retention rule, you can assign log sources to this rule, see Assigning Log Sources to a Data Retention Rule.

Modifying Rule Settings

You can modify the retention rule settings from **Administration > Data Retention Rules** menu.

Procedure

1. Click the Edit icon to edit the retention rule.

   The Edit Retention Rule window appears.

2. Make the appropriate changes.

   Note: The maximum raw data retention period for all appliance models is 10 years.

3. Click the **Save** button to save the new settings.

4. Click the **Commit Changes** button to commit the latest changes. Click the **Revert Changes** button to revert to the last committed changes.

> **Note:** After you make all changes, click the **Commit Changes** button to commit to all changes. After you click the **Commit Changes** button, the new data is stored as per the new rule.

Assigning Log Sources to a Data Retention Rule

You can assign log sources to a Data Retention rule at any time.

After you assign a log source to a particular Retention rule, the log source acquires the Retention time settings of the rule.

Procedure

1. Select a rule from the rule list or click the **View All Rules** button. Make sure that the desired log source is listed on the right pane.
2. Select the log source from the list and drag it to a different rule in the left pane. Press the **Shift** or Ctrl key to select multiple log sources from the list. The confirmation window appears.
3. Click **Yes** in the confirmation window to change the data retention rule for the specified log source.
4. Click the **Commit Changes** button to commit the latest changes. The new data retention time will be applied to the selected log source. Click the **Revert** Changes button to revert to the last committed changes.

> **Note:** After you make **all** changes, click the **Commit Changes** button to commit to all changes. After you click the **Commit Changes** button, the new data will be stored as per the new rule. The retention policy of data already stored on the appliance for the log source does not change.

Prioritizing Custom Rules

The Custom Rules are prioritized in the order as they are displayed (from highest on the top) in the Custom Rules list.
The priority of rule determines the Retention rule that is applied to which log source. You can change the priority of custom rules by moving them up or down in the Custom Rules list.

To view the Retention rule that is applied to a device or group, click the **View All Rules** button. The screen displays the Effective Rule for every device and device group. The Effective Rule Name column displays the Rule that is in effect after considering rule prioritization.

Example

1. We created two Custom Retention Rules named RR1Week and RR3Months that have Raw Data Retention period of 1 Week and 3 Months respectively.
2. Then the device group named Windows Machines is assigned to the RR1Week retention rule.
3. Similarly, a log source named FrontDesk1 is assigned to the RR3Months retention rule.

   Note: The FrontDesk1 is also part of the device group Windows Machines.

4. Since FrontDesk1 is assigned to the RR3Months retention rule with a retention time of 3 Months and is also a part of the group Windows Machines with a retention time of 1 Week, to decide the Effective Retention Rule (that is, the retention time applicable to the data received from FrontDesk1) the system uses the data retention rule priorities. If the rule RR1Week is prioritized higher, then the Effective rule for FrontDesk1 is RR1Week.
5. However, if the rule RR3Months is prioritized higher, then the Effective rule for FrontDesk1 is RR3Months.

Procedure

1. From the left pane, select the **Custom rule** from the custom rule list. Using the drag-

drop method, re-arrange the custom Rules in the priority order. The **Change the Priority** confirmation window appears.

2. Click **Yes** in the confirmation window to change the priority of the selected rule. The Rules are prioritized in the order they are displayed (highest on the top).
3. Click the **Commit Changes** button to commit the latest changes. Click the **Revert Changes** button to revert to the last committed changes.

> **Note:** After you make all changes, click the **Commit Changes** button to commit to all changes. After you click the **Commit Changes** button, the new data is stored as per the new rule.

Deleting a Custom Rule

You can delete a custom retention rule.

Procedure

1. Click the  icon to delete the retention rule.

   The Delete Rule confirmation window appears.
2. Click **Yes** in the confirmation window to delete the selected rule. All log sources under this rule are automatically moved to the **Default** Retention Rule.
3. Click the **Commit Changes** button to commit the latest changes. Click the **Revert Changes** button to revert to the last committed changes.

> **Note:** After you make all changes, click the **Commit Changes** button to commit to all changes. After you click the **Commit Changes** button, the new data is stored as per the new rule.

Working with Suites

A Compliance Suite edition consists of alerts, custom reports, search filters, or any combination of these components. From **Management > Suites**, you can group the alert rules, custom reports, and search filters of any Compliance Suite.

By grouping, you can:

- Import and export all the components together between appliances in a single suite.
- Access all suite components from a single location on the GUI.

For example, if you are using one of the Compliance Suites such as the Sarbanes-Oxley (SOX) Edition, you can manage or use all its built-in components from the **Suites** page.

To copy or move groups of components to other appliances, use the import and export functionality. For more information, see Import/Export of Entities Between Appliances.

Note: You have the same manageability of components when accessed in a suite as you do when accessing them as standalone entities.

Managing Compliance Suites

The **Management > Suites** page lists all Compliance Suites in the LogLogic appliance and their corresponding number of alerts, reports, and search filters in each suite.

Note: When you install LogLogic LSP, you can also import the suite of alert rules, custom reports, and search filters if available with LogLogic LSP. However, you cannot install LogLogic LSP from the **Management > Suites** page. You must follow the installation procedure in the *TIBCO LogLogic® Log Source Packages Installation and Upgrade*.

You can list the reports of a suite in the Reports and Search menus. When viewing the reports of a suite from the Reports or Search views, you can run each report ad-hoc or update the report. When you update a report, the changes affect all suites containing that

report. For more information about how to list reports under the Reports and Search menus, see Creating a Suite.

You can perform the following tasks with Compliance Suites:

Task	Steps
View existing suites	Go to the **Management > Suites** page to view the following information: • Number of associated alerts, reports, and search filters • Description • The suite owner • If the suite report is marked as shared • If the reports are listed under the Reports and Search menus as a suite. See Viewing Existing Reports Listed as a Suite.
To view or modify the suite components: • Alerts • Custom reports • Search filters	Click the number in the corresponding column: • Alerts column • Reports column • Search Filters column See Adding or Removing Components in a Suite.
Create a suite in the appliance	See Creating a Suite.
Edit an existing suite	To edit the suite details : 1. Click the required suite name. 2. Update the fields (name, description, sharing, and listing under Reports and Search menus) as required. If required, you can add or remove components in the suite. 3. Click the **Update** button to save your changes.
Delete a suite	If you no longer need a specific suite, you can completely remove it and all

Task	Steps
	its components from the appliance. **Note:** To maintain the suite as a reference point, consider renaming the suite, removing it as a listed item in the navigation menu, and removing shared access for other users. 1. Select the check box for one or more suites and click the **Remove** icon. 2. To confirm removal, click **OK**. After deleting a suite, the reports, alerts, and search filters from that suite are not part of any suite, but are available under the following menus: • **Reports > All Saved Reports** • **Alerts> Manage Alert Rules** • **Search > All Saved Searches**

Creating a Suite

To create a suite, specify the details of the suite and then add components to the suite.

A suite need not contain all three components—alerts, reports, and search filters. For example, you can create a suite containing only alerts; only reports; or alerts, reports, and search filters.

Procedure

1. Go to **Management > Suites**.
2. Click the **Add New** icon.
3. On the **Suites** tab, provide the following information:

Field	Description
Name	Name of the suite
Description	Description of the suite
Share with Other Users	Identifies whether the suite can be accessible by other users
List under Reports/Search Menu	Identifies whether the reports in the suite are listed as a suite under **Reports** and **Search** menus

4. Click the **Add Suite** button.

   The main **Suites** page is displayed, with the new suite listed in the table. You can also click the **Alerts**, **Reports**, or **Search Filters** tabs to create the suite and to go directly to adding components to your suite. You do not have to click the **Add Suites** button.

Viewing Existing Reports Listed as a Suite

From the Reports menu, you can view all saved reports for which the **Share with Other Users** option has been selected.

Procedure

1. From the **Reports** menu, click the required report type and report sub-category. For example: **Reports > Access Control > User Access**.

   The page displays all saved reports for which the **Share with Other Users** option has been selected.

2. From the Actions column, you can **Run** or **Edit** the saved report.

Adding or Removing Components in a Suite

You can add or remove components such as alerts, reports, and search filters in a suite.

Before you begin

Before adding components to a suite, you must define them in the LogLogic appliance.

Procedure

1. From the **Suites** page, click the suite name to which you want to add components.

   The **Suite** tab appears, with the **Name** and **Description** fields filled in. Check boxes for **Share with Other Users** and **List under Reports/Search Menu** also appear as selected if you have so designated when the suite was created. You can edit these fields if required.

2. Click the appropriate tab - **Alerts, Reports**, or **Search Filters** - corresponding to the component you want to add or remove.

   Any components you already added appear in the table.

3. To add a component:

   a. Click the **Add New**  icon.

      The **Add *component-name*** tab appears, where *component-name* is Alerts, Reports, or Search Filters.

   b. Select the check box for the entities you want to add and then click the **Add**  icon.

      The added components are displayed on the ***component-name*** tab.

4. To remove a component:

   a. Select the components that you want to remove and click the **Remove**  icon.

   b. To confirm the removal, click **OK**.

5. Go back to the **Suites** tab and click **Update** to save your changes.

Import or Export Entities Between Appliances

To share certain configured information from one appliance to another, you can import or export the information as entities.

To repeat configuring the same entities on multiple appliances, you can configure them on one appliance, export them, and then import to other appliances.

List of Entities That Are Imported or Exported from the GUI

You can import and export the following entities from the GUI:

- Alert Templates
- Alerts
- Custom Reports
- Devices and Device Groups
- Search Filters
- Suites
- Users

List of Entities That Are Imported or Exported from the CLI

You can import or export the following entities only from the CLI:

- Advanced Dashboard
- Aggregation Rules
- Bloks
- Data Models
- Enrichment lists
- Groups and nested groups

- Triggers

To export these entities, the Advanced Features must enabled on the **Administration > System Settings > General tab > Advanced Features Settings** section. See Advanced Features. You can import entities when Advanced Features is disabled; however, you cannot view them until you enable Advanced Features.

Important Considerations

By default, imported triggers and aggregation rules are disabled. After importing, perform the following actions as applicable:

Item	State after importing	Required action item
Triggers	Disabled and not synchronized	Enable and synchronize the required triggers from the GUI. See Managing Triggers in *TIBCO LogLogic® Log Management Intelligence User Guide*.
Aggregation rules	Disabled	Enable the required aggregation rules from the GUI. See Managing Aggregation Rules in *TIBCO LogLogic® Log Management Intelligence User Guide*.

Importing Entities using the GUI

From the **Administration > Import/Export** page, you can import configured entities that are exported from another appliance.

Before you begin

Before importing an alert that includes an alert receiver, you must create an alert receiver with the same name. Otherwise, the import fails.

Procedure

1. Click the **Browse** button next to the **File Name** field.

2. Browse to and select the file. The file name appears in the **File Name** field.

   > ⚠ **Warning:** Files must be in valid XML format. Attempting to import other formats results in an error message.

3. Click **Load** to fetch a list of entities available for import.

   The entities available for import appear in the **Available Entities** list.

4. From the **Available Entities** list, click an entity or hold the Ctrl key while clicking to highlight more than one entity.

5. Use the arrow buttons to select one or more entities.

   The selected entities move to the **Selected Entities** section on the right side.

6. Click **Display Info** to display the information about the selected entities in XML format in the lower text area.

7. Click the **Import** button to import the selected entities to the appliance.

   > ⚠ **Caution:** If the number of devices to be exported is more than 60,000, the import operation might fail.

Result

The imported entities are available under the following menus:

Entity	Menu
Alert templates	Alerts > Manage Alert Templates
Alerts	Alerts > Manage Alert Rules
Custom reports	Search > All Saved Searches
Devices	Management > Devices
Search filters	Search > All Search Filters

Entity	Menu
Suites	Management > Suites
Users	Management > Users

Exporting Entities using the GUI

From the **Administration > Import/Export** page, you can export configured entities into an XML file format.

Note: Check Point interfaces can be exported and imported, but not Check Point firewalls and Check Point servers.

Procedure

1. Go to the **Administration > Import/Export > Export** tab.
2. From the **Entities** list, select the required entity:
   - Alerts
   - Users
   - Search Filters
   - Custom Reports
   - Devices
   - Suites
   - Alert Templates

   The entities available for export appear in the text area in the **Available Entities** section.
3. Select the appropriate **Export Mode** option. This controls how exported entities are handled after they are imported on another appliance.

Mode	Description
Insert	If the entity exists in the target appliance, the entity is rejected on import.
InsertOrUpdate	If the entity exists in the target appliance, the entity is updated on import; otherwise a new entity is inserted. **Note:** The InsertOrUpdate option does not work on devices with a Collector Domain ID. For these devices, only the description is updated.
Delete	If the entity exists in the target appliance, the entity is deleted on import.

4. From the **Available Entities** section, select an entity to export. Click an entity or hold the Ctrl key while clicking to highlight more than one entity.
5. Use the arrow buttons to select one or more entities.

   The selected entities move to the **Selected Entities** section on the right side.
6. Click the **Export** button to export the selected entities.
7. In the File Save As dialog box, specify the location to downloaded the file in XML format.

Note: If the number of devices to be exported is more than 60,000, the export operation fails and the XML file is not downloaded.

Import or Export Configurations using the CLI

You can import or export the following entities only from the CLI:

- Advanced Dashboard
- Aggregation Rules
- Bloks
- Data Models

- Enrichment lists
- Groups and nested groups
- Triggers

Note: This feature is available only if the Advanced Features option is enabled in the System Settings. For more information on enabling advanced features, see General Settings.

Make sure that the LogLogic LMI system is running before you import or export any configurations.

The configuration of Advanced Features is saved in a JSON file. Starting from LogLogic LMI 6.4.0, the entities are stored in groups. The value of the group name and path of the entities are stored in the `groupName` and `parentPath` parameters.

For example, if a data model is stored in the **LSP > Active_Directory** group, then the following parameters are stored for that data model in the JSON file:

```
"groupName": "Active_Directory",
"parentPath": "/LSP/Active_Directory"
```

Export Commands

Use these commands when exporting configurations from the CLI.

Create a symlink in `/loglogic/bin/.` to the llconf file: `/loglogic/logu/configurator/bin/llconf`.

To do this	Run this command	Result
export all configurations	`./llconf export`	The default `loguconfig.json` file is automatically created in the same directory.
export configurations into the specified file	`./llconf export -f <path_to_file>` For example:	The file is saved at the defined location. Only entities are included; not

To do this	Run this command	Result
	`./llconf export -f logu.json --verbose`	groups.
export configuration into yaml format	`./llconf export -y`	The default `loguconfig.yml` file is automatically created in the same directory.
export configuration into a zip file	`./llconf export -z`	The default `loguconfig.zip` file is automatically created in the same directory.

export the selective configuration from the appliance	`./llconf export --configlist <path_tofile>`
	The following expressions can be used in the `configlist` file:
	`{`
	`"sources" : [ "abc", "xyz" ], // This will export Data Models named "abc" and "xyz" only`
	`"bloks" : ["blok*"], // Export all Blok configurations with name starting "blok"`
	`"alertTriggers" : ["*"], // Export all triggers`
	`"smartlists" : ["*"], // Export all Smartlist configurations`
	`"dashboards" : ["*"], // Export all Dashboard configurations`
	`"aggrRules":["*"], // Export all aggregation rules`
	`"triggerGroups":["*"] // Export all trigger groups`
	`}`	

To do this	Run this command	Result
	`// To skip all Advanced Data Models "source": []}` **Note:** If you are using this sample, make sure you remove the comments (shown by //) at the end of the line	
export a JSON file with only the groups structure.	`./llconf export -f logu.json --verbose --groupsOnly`	Entities other than groups are not exported.

Import Commands

Use these commands when importing configurations from the CLI.

To do this	Run this command	Result
import all configurations	`./llconf import`	Imports configurations from the default configuration file `loguconfig.json` or `loguconfig.yaml` or `loguconfig.zip`.
import configurations from a specific file	`./llconf import -f <path_to_file>` **Attention:** The file path must not contain spaces.	Imports configurations from the specified file.
overwrite the existing configurations into the system	`./llconf import -o`	Imports configurations from the default file and overwrites the existing configurations on the appliance.

To do this	Run this command	Result
import a JSON file with only the groups structure.	`./llconf import -f logu.json --verbose --groupsOnly`	Imports only the groups structure from the specified JSON file. Entities are not imported.

Alert Receivers

Use the **Administration > Alert Receivers** tab to define your SNMP Traps or Syslog receivers, for example, a network monitoring and ticketing system.

After you set up a trap or syslog receiver, you can define an alert in the **Alerts > Manage Alerts** tab.

Using this tab you can:

- Add a new Alert Receiver in the system, see Adding a New Alert Receiver.
- Modify an existing Alert Receiver from the system, see Modifying an Alert Receiver.
- Remove existing Alert Receivers from the system, see Removing an Alert Receiver.

Viewing existing Alert Receivers

Note: To view existing Alert Receivers, you must have access to all devices in the appliance. Make sure that the **Manage Users > Devices tab > Access all devices in appliance** check box is selected to view all Alert Receivers.

From **Administration** menu, select **Alert Receivers**. A list of all alert receivers is displayed. The Alert Receivers page displays the following information:

Alert Receiver Information

Element	Description
Name	Name of the SNMP Trap or syslog you designate.
IP Address and port	IP address and port of the SNMP Trap or syslog.
Type	The alert receiver type.
Enabled	Indicates whether the SNMP Trap or syslog is activated for your appliance.

Element	Description
Description	Description of the SNMP Trap or syslog.

Note:
- Currently, only the latest updates are saved in this file.
- One collecting port is applicable to all receivers.

Adding a New Alert Receiver

The **Add Alert Receiver** page lets you add a new alert receiver.

Procedure

1. Click **Administration > Alert Receivers** from the home page.
2. Click **Add New**.
3. In the **Name** field, enter the name for the alert receiver.
4. In the **IP Address:Port** field, enter the IP address and port for the alert receiver.
5. Under **Enable**, select the **Yes** radio button to enable the receiver after you complete this tab.
6. Under **Receiver Type**, select the radio button to indicate whether this is an **SNMP Trap** or **Syslog** receiver.
7. (Syslog only) In the **Port** field, enter the port number for the syslog receiver.
8. (SNMP trap only) In the **Outbound Traps Community String** field, enter the community string used between the SNMP manager and SNMP proxy to identify the sender.

   The SNMP receiver is configured with a community string, which is like a password, obtained from the SNMP receiver system administrator. All appliances sending traps to a single receiver must use the same string. Appliances sending traps to that receiver with any other string are ignored.
9. (SNMP trap only) In the **Optional Text** field, enter a description to appear in your

trap message. For example, use this field to distinguish between two machines on the same trap server that sends messages. This optional text string is included in every SNMP trap.

10. In the **Description** field, enter a description for the alert receiver.
11. Click **Add** to add the alert receiver to the appliance.

Result

The appliance returns you to the **Alert Receivers** tab, which now includes the new alert receiver.

Modifying an Alert Receiver

From the **Modify Alert Receiver** tab you can edit an existing alert receiver.

Procedure

1. Click **Administration > Alert Receivers** from the home page.
2. Click on the Name of the existing Alert Receiver.
3. Make the appropriate changes.
4. Click **Update** to update the alert receiver on the Appliance.

   The appliance returns you to the **Alert Receivers** page.

Removing an Alert Receiver

You can remove existing Alert Receivers at any time.

Procedure

1. Click **Administration > Alert Receivers** from the home page.
2. Click in the Alert Receiver name’s check box to select.
3. Click **Remove** to remove the selected alert receiver from the system.
4. From the **Remove Alert Receivers** tab, confirm the removal of the selected alert

receivers from the system, click **Confirm Remove**.

Configuring Advanced Features

As an administrator, you can configure Advanced Feature settings that are not available from the GUI or to users without administrator access rights.
To configure such settings, perform the following procedure.

> ⚠ **Caution:** Be careful when changing the configuration parameters. Deleting a parameter might cause the corresponding features to stop working or not work as expected.

Procedure

1. In the `/loglogic/bin` folder, create your own temporary configuration file.
2. Copy all existing configuration parameters into your file, and then change the values of the required parameters. Add new parameters, if required.
3. Run the following command in the `/loglogic/bin` folder:

   ```
   ./llconf -f <filename>.conf --verbose
   ```

   where <*filename*> is the name of your temporary configuration file.

Result

After you run the command, all parameters from your file are overwritten in the configuration settings of the appliance.

Sample Configuration Files

For the configuration settings and sample configuration files of Advanced Features, see the following sections:

- Configuration of Advanced Alerts
- Configuration of Scheduled Queries
- Configuration of Aggregation Rules

- Configuration of the Monitoring Console
- Setting Up the Geographical Database

Configuration of Advanced Alerts

As an administrator, you can change the configuration settings of Advanced Alerts.

Triggers describe what action should be taken when a correlation Blok is triggered. If several triggers are associated with the same correlation Blok, all of them are triggered. When a trigger is activated, an alert is sent out in the form of an email or syslog notification.
Advanced Alerts are configured in LogLogic LMI. The alert configuration includes the following settings:

- A list of categories for acknowledgement and triggers: such as Attack on Third Party or Security Alert
- A list of severities
- The working hours, in 24-hour time format: for each day, you can specify up to two time ranges. For example, if you specify 08:30-12:30 and 13:30-18:00, it indicates a break of one hour between 12:00 PM and 6:00 PM.
- The SLA for each severity, in minutes: for example:`"medium": "240"` indicates that the SLA for a medium severity alert is 240 minutes.

  Example: Suppose the working hours are set as per the following sample configuration file, the SLA for a medium alert is 240 minutes, that is, 4 hours, and the alert arrives on Monday at 11:30 AM. The SLA time is calculated by considering breaks in the working hours, if any. Therefore, in this case, the alert expires at 16:30 hrs or 04:00 PM, considering there is a one-hour break in the working hours of Monday.
- The alert purge interval in hours, when the alert purge task must be triggered. The default value is 24.
- The alert retention policy, in days.

Sample Configuration File

The following is a sample file that includes all settings for advanced alerts. In a browser, type the URL: `<appliance_IP>/monitoring` and then navigate to

`/unity/system/config/alerting/tenant1`. The path mentioned in *nodePath* must be `/unity/system/config/alerting/tenant1`. The values shown in the following sample file are the default values stored in LogLogic LMI.

```
{
  "configurations": [
    {
      "nodePath": "/unity/system/config/alerting/tenant1",
      "data": {
        "Categories": [
          "Attack on third party2",
          "Authorized Activity",
          "Authorized security testing",
          "Emergency changes",
          "False positive",
          "Known error",
          "LogLogic Event",
          "Network Noise",
          "Security Alert",
          "Suspicious Activity",
          "Unauthorized Activity",
          "Unknown"
        ],
        "Severities": [
          "Info",
          "Low",
          "Medium",
          "High"
        ],
        "workHours": {
          "monday": {
            "morning": "08:00-12:30",
            "afternoon": "13:30-17:30"
          },
          "tuesday": {
            "morning": "08:00-12:30",
            "afternoon": "13:30-17:30"
          },
          "wednesday": {
            "morning": "08:00-12:30",
            "afternoon": "13:30-17:30"
          },
          "thursday": {
            "morning": "08:00-12:30",
            "afternoon": "13:30-17:30"
          },
```

```
          "friday": {
            "morning": "08:00-12:30",
            "afternoon": "13:30-17:30"
          }
        },
        "slaForSeverity": {
          "Info": "1440",
          "Low": "720",
          "Medium": "240",
          "High": "60",
          "None": "10"
        },
        "alertPurgeSchedulerInterval" : 24,
        "alertsDaysOfRetention" : 90
      }
    }
  ]
}
```

Configuration of Scheduled Queries

As an administrator, you can change the configuration of scheduled query parameters.

You can modify the following parameters related to queries:

- Email attachment size (`SMTPAttachmentLimit`): Queries can be scheduled to run at the specified days and time. The report can be sent as email attachments to the specified recipients. You can configure the maximum size of the email attachments (in MB) in the `SMTPAttachmentLimit` parameter. The default value is 10 MB. You might need to add this parameter in the your configuration file if it is not present in the current appliance file.
- Limit of concurrent queries (`maxConcurrentQuery`): You can set the maximum number of queries that can be concurrently processed on the appliance in the `maxConcurrentQuery` parameter. This is a global setting for queries running on the appliance, and not specific to a user. The default value is 25.

Sample Configuration File

The following sample file includes settings for email attachment and maximum concurrent queries. The path mentioned in *nodePath* must be `/unity/system/config/querynodes/querynode-0000000000`. The values shown are the default values stored in the LogLogic LMI appliance.

```
{
  "configurations": [
    {
      "nodePath": "/unity/system/config/querynodes/querynode-
0000000000",
      "data": {
        "services": {
          "rest": {
            "host": "0.0.0.0",
            "port": 9681
          },
          "query": {
            "host": "0.0.0.0",
            "port": 9620
          }
        },
        "storage": {
          "cache": "data/qcache",
          "maxSplitH2fileSize": 31
        },
        "SMTPAttachmentLimit" : 10,
        "maxConcurrentQuery" : 25,
        "version": 1
      }
    }
  ]
}
```

Configuration of the Monitoring Console

As an administrator, you can change the configuration of the Monitoring Console.

In the Monitoring Console configuration files, you can modify advanced settings such as defining the number of alerts that can be retained by default, or the storage options. The values in the following sample files are the default values stored in the LogLogic LMI

appliance. You can modify the values as required and then perform the steps described in Configuring Advanced Features.

Sample Configuration File for Monitoring Console Parameters

The following sample file includes the Monitoring Console settings:

```
{

  "configurations": [
    {
      "nodePath":
"/unity/system/config/hawkconsolenodes/hawkconsolenode-0000000000",
      "data": {
        "services": {
          "rest": {
            "host": "0.0.0.0",
            "port": 9687
          }
        },
        "alerts": {
          "notificationRetentionCount": 100000,
          "lowRetentionCount": 100000,
          "mediumRetentionCount": 100000,
          "highRetentionCount": 100000
        },
        "domainConfig": {
          "defaultDomainName": "lmi_domain",
          "defaultClusterPort": 9688,
          "domainTransportConfigPath":
"conf/DomainTransportConfigPath.yml",
          "repositoryPath": "conf/repository"
        },
        "storage": {
          "cache": "/loglogic/data/.hawkconsole/hcache",
          "type": "mysql",
          "host": "127.0.0.1",
          "port": 3306
        },
        "version": 1
      }
    }
```

```
  ]
}
```

Parameter Description	Default value
Maximum number of alerts	100,000 alerts `"alerts": { "notificationRetentionCount": 100000, "lowRetentionCount": 100000, "mediumRetentionCount": 100000, "highRetentionCount": 100000 }`
Configuration of the appliance that hosts the Monitoring Console	`"domainConfig": { "defaultDomainName": "lmi_domain", "defaultClusterPort": 9688, "domainTransportConfigPath": "conf/DomainTransportConfigPath.yml", "repositoryPath": "conf/repository" }`
Storage medium of alerts	MySQL database, its IP address, and port `"storage": { "type": "mysql", "host": "127.0.0.1", "port": 3306 }`

Configuration of Aggregation Rules

As an administrator, you can change the configuration of aggregation rules.

You can modify the following aggregation parameters:

- `purgeFrequencyInMinutes`: The aggregated data of all aggregation rules in the appliance is purged every 60 minutes (default value). If required, you can change the default frequency by changing the value of the `purgeFrequencyInMinutes` parameter (in minutes).

Sample Configuration File

The following sample file includes the settings for aggregation rules. The values shown are the default values stored in the LogLogic LMI appliance.

```
unity/system/config/aggregationnodes/aggregationnode-0000000000
Node data:
{
 "services": {
   "rest": {
     "host": "0.0.0.0",
     "port": 9685
   },
   "query": {
     "host": "0.0.0.0",
     "port": 9626
   }
 },
 "storage": {
   "type": "mysql",
   "path": "/loglogic/data/.aggregation/acache",
   "indexed": true,
   "host": "127.0.0.1",
   "port": 3306
 },
 "purgeFrequencyInMinutes": 60,
 "version": 1
}
```

The changes take effect only after you run the `llconf` command, as per the steps provided in Configuring Advanced Features.

Setting Up the Geographical Database

You can use the `geoiplookup()` function in LogLogic LMI to search logs that originated from a particular geographical area such as location, country, city, and postal code. You can use the function within SQL and EQL queries, in Advanced Search, and in a Geomap widget on the Advanced Dashboards.

The Geomap widget gives you a unified view of your data visualization and its geographical distribution. For example, you can plot VPN connection logs and the IP addresses from which they originate. The widget displays the IP addresses as points or a bubble chart on the geographical map.

To use the `geoiplookup()` function to fetch the geographical information of a specified IP address, an administrator must download the appropriate MaxMind database file (`.mmdb`) to the `/loglogic/data/geoIP` directory on the appliance.

Note: When you obtain third-party software or services, it is your responsibility to ensure you understand the license terms associated with such third-party software or services and comply with such terms.

Important Considerations

- The function can use only one MaxMind database file (`.mmdb` file) at a time. Therefore, ensure that the `/loglogic/data/geoIP` directory includes only one file.
- In a Management Station setup, you must copy the `.mmdb` file to each Remote Appliance.
- In a High Availability environment, the `.mmdb` file on the active appliance is automatically replicated on the standby appliance.

Note: If you replace the `.mmdb` file, you do not need to restart the appliance.

Related Topics

For more information, see the following topics in the *TIBCO LogLogic® Log Management Intelligence User Guide*:

- Description of the `geoiplookup()` function and its parameters: Predefined Functions
- Search query examples: Search Examples
- Configuration and usage of the widget: Geomap widget

Artificial Intelligence Queries

LogLogic LMI includes artificial intelligence and machine learning capabilities to classify information from unknown log sources. This is achieved by using the TensorFlow trained model to automatically classify logs of access or audit types. By using an additional classifier in the advanced search query, you can view the additional information in the search results.

You can use the training model classification in advanced search queries, correlation, aggregation rules, and Bloks.

As an administrator, ensure that:

- Advanced Features are enabled on the appliance from the **Administration > System Settings > General** tab.
- The TensorFlow training model is present at the following location on the appliance:

```
/loglogic/logu/classifier/ll_tax_v1
```

The directory contains the model and its classification dimensions.

Note: You can use the built-in training model provided in LogLogic LMI. You cannot create a new model or modify the built-in model.

Creating or Customizing EQL Functions

LogLogic LMI provides built-in EQL functions. If required, as an administrator, you can customize existing functions or create your own EQL functions, package them into a library, and use them in LogLogic LMI.

Existing (built-in) EQL functions are available in the `lmi-querylanguage` library. To create your own package, you can:

- Modify the functionality of an existing EQL function
- Use an existing function as the parent class of your new function

By packaging your functions into JAR files and saving the package at `/loglogic/logu/customfunctions`, the functions are loaded into LogLogic LMI. Then users can use the functions in the following features:

- Advanced Search with EQL or SQL
- Column expression of a data model
- Definition of aggregation rules

Limitation

Customized functions are not supported in the following features:

- Within a GROUP BY clause
- Definition of correlation Bloks

Before you begin

Before creating your function library, ensure that the following requirements are met:

- Your Java development and build environment has Java 8 or 11.
- The Java version on the PC matches the one on the LogLogic LMI appliance.

Procedure

1. Create a Java or maven project.
2. Include the following JAR files as dependencies in the project:
   - `lmi-api-<LMIversion>.jar`
   - `lmi-querylanguage-<LMIversion>.jar`
   - `ll-utils-<version>.jar`

   These files are stored on the appliance at `/loglogic/logu/lib`.
3. Write the classes and functions required. For more information, see the following sections:
   - Guidelines
   - List of Base Classes
   - Base Class Methods
4. Export the project as a JAR file, excluding the project metadata, and save the JAR file at the following directory on the LogLogic LMI appliance:

   ```
   /loglogic/logu/customfunctions/
   ```

   Note: Do not include the dependencies in the exported file.

Result
Your functions are loaded into the LogLogic LMI appliance and are ready for use.

What to do next
Use your new or modified functions in an advanced search, a data model, or an aggregation rule.

Guidelines

When creating your own functions and classes, you must consider the following guidelines and limitations:

- The implementation of a new class must be a child class of `ScalarFunction` or

`AggregateFunction` - either directly or indirectly.

> **Warning:** Do not use any other class as a parent class.

- Objects must be converted into data types defined as an Enum type in the DataType class (com.tibco.apollo.DataType). You can use the conversion methods defined in the following utility classes:
  - com.tibco.apollo.querylanguage.functiondefs.Utils
  - com.tibco.apollo.querylanguage.evaluators.impl.utils.DataConversionUtils
- Any new function must be dependent on either Java internal functions or the following libraries:
  - `lmi-api-<LMIversion>.jar`
  - `lmi-querylanguage-<LMIversion>.jar`
  - `ll-utils-<version>.jar`

> **Note:** External libraries are not supported.

List of Base Classes

Define your classes directly or indirectly based on one of the following classes:

Type of Usage	Base class	Location of base class
For aggregation operations	AggregateFunction	In the `lmi-querylanguage.jar` library, as `com.tibco.apollo.querylanguage.functiondefs.AggregateFunction`
For operations other than aggregation	ScalarFunction	In `lmi-querylanguage.jar` library, as `com.tibco.apollo.querylanguage.functiondefs.ScalarFunction`

Base Class Methods

Your new class that extends the ScalarFunction or AggregateFunction class must override the following methods of the base class. The method must return the value as specified in

the following tables.

Methods in both classes

Method	Return value
String getName()	Must return the function name. The same function name must be used in the search query.
DataType getReturnType(List <*DataType*> args)	Must return the expected return type of the function **Note:** • If extended from ScalarFunction and for using the function in a WHERE clause, the return type must be DataType.BOOLEAN. • If extended from AggregateFunction, the return type for this method must be the same as the type of the input argument of the AggregateFunction class.
boolean isParameterListValid (List <*DataType*> args)	Must return `true` or `false` based on the validation of input arguments of the function **Tip:** `com.tibco.apollo.DataType` contains Enum values of different data types. You can use those values to match the data type of the input arguments.

Methods - only in the ScalarFunction class

Method	Return value
String getEvaluationExpression()	Must return the function as an expression with input arguments
Object evaluate(Event event, List <*Object*> list)	This method evaluates every log event that is valid as per the criteria in the search query. The `event` object refers to the log event and the `list` object contains the arguments that are passed into the function in the search query.

Methods - only in the AggregateFunction class

Method	Return value
String getTotalAggregationExpression (List <*DataType*> paramTypes)	Must return the string `partialResult`.
String getResultUpdateExpression(List <*DataType*> paramTypes)	Must return the string `case when partialResult is null then ? else partialResult+? end`.
Aggregator getAggregator(List <*DataType*> paramTypes)	Must return a `com.tibco.apollo.querylanguage.functiondefs.Aggregator` object The calculation is done by using the `add()` function and the result is obtained by using the `getResult()` function. These functions must be defined in your customized Aggregator class. See Customizing the Aggregator class.

Customizing the Aggregator class

You can customize the Aggregator class by defining it as either an inner class or within the same package. You must define the following methods of the Aggregator class.

Method	Description
Object getResult()	This method must return a result object that is converted to the appropriate data type.
void add (Object object)	This method is invoked for all logs in the search query. • The input argument is the object passed as an argument of the function. • You must process the input object and store the result in any class member object so that the stored object can then be returned from the `getResult()` method.

Processing Jumbo Messages

In LogLogic LMI, messages that exceed the default processing limit (64 KB) are termed as jumbo messages. You can configure LogLogic LMI to process jumbo messages.

By default, LogLogic LMI cannot process messages larger than 64 KB according to the UDP syslog standard. However, you can enable this feature for TCP syslog and ULDP protocols.

Enabling or disabling jumbo messages does not affect receiving and processing of replayed data and the functioning of the file collector. However, you must read the following Important Considerations section before enabling this feature.

Important Considerations

- Enabling this option might impact search performance.
- Jumbo messages can be processed by TCP syslog and ULDP, but not by UDP syslog.
- Advanced Search can process jumbo messages. However, the classic search methods (such as index search and regular expression search) cannot.
- For jumbo messages, the complete message content cannot be displayed on the Triggered Alerts, real-time reports, and Log Source Status pages.

To start processing jumbo messages, modify the `/loglogic/conf/tcpcoll.conf` file as follows:

Procedure

1. Set the value of the `UseTcpCollectorQueue` property to 1.
2. Specify the maximum message length. When the `MaxMsgLength` limit is exceeded, the extra characters of the messages are truncated. The valid values for the `MaxMsgLength` property are as follows:

UseTcpCollectorQueue	Valid range of MaxMsgLength (in bytes)
1 (jumbo messaging is enabled)	5 to 1048576
0 (jumbo messaging is disabled)	5 to 65535

If the value of `MaxMsgLength` is invalid or is not specified, the value is set to the maximum limit of message length.

What to do next

After updating the value of `UseTcpCollectorQueue` or `MaxMsgLength`:

- If you are running this setup on a single node that is not in an HA setup, you must restart `mtask` by running the following commands:
  1. `$ mtask stop`
  2. `$ mtask start`
- If you are running this setup in an HA setup, you must restart both active and passive nodes.

Report Settings

All generated reports except Advanced Search can be downloaded in multiple formats, for example, CSV, PDF, or HTML. Some report settings can be configured from the command line interface.

Maximum rows in the downloaded report files

The number of lines in the downloaded file is limited to 5,000 even if the generated report has more number of lines. For CSV files, you can change this limit:

1. Edit the following file:

```
/loglogic/tomcat/webapps/logapp20/WEB-
INF/config/realtime/ReportRuleMasterConfig.xml
```

2. Change the value 5000 to the required value:

```
<csvMaxRows>5000</csvMaxRows>
```

System Settings

The **System Settings** tabs let you manage the overall system configuration of a LogLogic appliance.

The following sections describe the settings on the various tabs of the **System Settings** page.

General Tab

Use the **Administration > System Settings > General** tab to configure system-wide settings.

The system settings automatically take effect after you click **Update**.

Note:
- Changes to certain settings might require restarting the appliance.
- Not all options are available for LogLogic ST Appliances. Options that apply only to LogLogic ST Appliances are noted where applicable.

- SNMP Trap Collector Settings
- Syslog Forwarding Settings
- System Performance Settings
- Usage Count
- Theme for Rebranding LogLogic LMI
- Custom Logo Upload Settings
- Build Details

General Settings

Field	Description
Originating Email	The email address that the appliance uses for the return address email notifications in alerts and scheduled reports. **Note:** If this field or any of the SMTP settings field are changed, you must re-synchronize triggers.
Concurrent Login Sessions	Indicates the maximum number of concurrent login sessions allowed for each LogLogic LMI user. After the permitted number of concurrent connections is reached, a message is displayed to the user on the next attempted login. The message indicates that the limit has been reached and requests the user to close one of the active sessions. The default number of concurrent login sessions allowed per user is 100.

Field	Description
	Note: • This value is applicable to all users in the system. For example, setting the value to 10 implies each user in the system can have at the most 10 concurrent sessions. • A full application restart is required for the changes to take effect. Follow the system prompts. • The limit on concurrent sessions is not applicable to the REST API used for Advanced Features. • After reaching the maximum number of concurrent login sessions, or when users abruptly end their login session, you can remove a user session from the **Management > Users > User Sessions** tab.
Multiline Delimiter	Sets a character string to be used as line delimiter in multiline logs.
Purge stale devices	Enables the option to purge stale devices based on the last time data was received. For more information, see Purge Stale Devices.
Purge Stale Devices Period	Specify the number of days after which stale devices must be purged.
SSH Daemon at Startup	The SSH Daemon provides access to the appliance's Command Line Interface (CLI) from SSH clients. By default, the SSH Daemon is turned on in the appliance. Click the slider to **No** to disable the SSH Daemon when you reboot the appliance. For details, see Command Line Interface (CLI).
Auto-identify Log Sources	Automatically detects any syslog log sources connected to the appliance. This includes: • Multiple log sources sharing the same IP address. LogLogic LMI considers multiple sources using the same IP

Field	Description
	address as a single host, because LogLogic LMI uses the IP address to uniquely identify them. • Log sources whose log data is converted to syslog during collection To view all identified log sources, use **Management > Devices**. If you do not enable this option, you must manually add the following log sources. **Note:** If the Auto-identify Log Sources option detects a log source but does not recognize the exact type, the appliance adds it to the **Management > Devices** list as a general syslog log source. To manually change a general Syslog log source type: 1. In the **Management > Devices** tab, click the log source name. 2. From the **Device Type** list, select the specific device type. 3. Click **Update**. The **Type** column displays the device type you associated with the auto-identified log source. If you enable **Auto-identify Log Sources** and you have several thousand devices configured that need to be auto-identified, routing rules and alerts can slow the auto-identify process down.
DNS Resolve All Device Names	Updates the DNS Resolve Flag for multiple devices. **Note:** If you select **No**, the **Management > Devices > Apply this update to all devices, not just to those on this page** check box overrides your General settings **No** option.
Full Text Indexing	Allows indexing of data. You can set this for your appliances independently. If enabled, all data is indexed.

Field	Description
	Note: Indexing uses additional storage. To retain the index data, configure rules from **Administration > Data Retention Rules**.
Allow Disabling of admin Super User	Allows an admin user to disable the admin super user account. For more information, see Managing Users.
UI Verbose Logging	Enables or disables logging detailed error messages on the GUI. The default value is **Yes**. If you select **No**, detailed logging is disabled, and a generic message is displayed instead of error or exception details. You can view the detailed information in syslog files by accessing the machine through SSH.
Log parsing for reports	Enable or disables parsing in LogLogic LMI, so that no content is added to the Real Time database-driven reports. The default value is **Yes**. This feature is only available on LogLogic EVA, LogLogic LX Appliances, and LogLogic MX Appliances, and in effect makes them similar to a LogLogic ST Appliance.
Manage Device	Enables or disables the ability for users to configure or add devices under **Management > Devices**. This setting overrides the Manage Devices privilege provided to a user or role, except to the admin super user account.
Accept Detail	Allows drill down for the Real Time and Summary detail reports. **Note:** • This option is not available on LogLogic ST Appliances. • You must enable this option to view **Reports > Network Activity > Accepted Connections > Network Activity > Application Distribution**, and detail reports. This might require additional time and storage in downloading these reports.

Advanced Feature Settings

Field	Description
Advanced Features	Enables or disables the Advanced Features. The default is **Off** (disabled). Select **On** to enable the following Advanced Features: • Advanced Search • Advanced Dashboards • Bloks: Filter Bloks, Correlation Bloks, and Time Bloks • Data Models • Enrichment Lists • Exporting and Importing Configurations Only an administrator with CLI access using root can use this feature. For more information on exporting and importing configurations, see the "Import or Export Entities Between Appliances" section in *TIBCO LogLogic® Log Management Intelligence Administration*. • Monitoring Console • Monthly index • Queries ◦ Search queries ◦ Scheduled queries ◦ Tail queries ◦ Distributed Advanced Search queries • REST API support for Advanced Search • Rule Management: ◦ Triggers ◦ Aggregation rules ◦ Distributed aggregation rules

Field	Description
	Important Considerations After enabling the advanced features, all sessions to the WebUI are disconnected for the period when the Tomcat engine restarts, after which users can login again. • You cannot access advanced features on a standby node in a high availability setup. However, you can access them from the public IP address or the IP address of the active node. Before configuring HA, you must disable Advanced Features on both active and standby appliances. • Use caution when enabling advanced features on LX4025, ST4025, MX4025, ST2025-SAN, or LX1025R1 models, because the memory requirements of these features when in use might cause performance issues. Also, continuous use of Advanced Features on these models can cause the appliance to run out of memory and lead to engine restarts or failure. • On an appliance where the memory is less than the minimum required 32 GB, you cannot enable Advanced Features from the GUI. You must enable it from the CLI, by running the system logu enable command. • For information about the behavior of Advanced Aggregation and Monitoring Console features after upgrading to LogLogic LMI 6.4.0, see the Upgrade Considerations section in *TIBCO LogLogic® Log Management Intelligence Configuration and Upgrade*. For information about enabling advanced features using the CLI, see the logu command.
Monitoring Console	Enables or disables the Monitoring Console and displays the **Monitoring > Console** menu. You can enable or disable this feature using the CLI by running the monitoring_console command.

Field	Description
Advanced Aggregation	Enable or disable the Advanced Aggregation features. After enabling the Advanced Aggregation option, the **Management > Rules > Aggregation** tab is visible to users and they can use the Advanced Aggregation features. By default, Advanced Aggregation is switched off. This feature can be enabled only if the Advanced Features option is enabled. If Advanced Aggregation is switched on before upgrading your setup, it remains switched on after the upgrade. **Important:** Before disabling Advanced Aggregation, ensure that you delete or disable any advanced aggregation rules to avoid storing unnecessary aggregated data. To enable or disable this feature using the CLI, see the advanced_aggregation command.
Monthly Index	Enables or disables the monthly index feature. The default is **No**. This feature can be enabled only if the Advanced Features option is enabled. For information about enabling monthly index using the Command Line Interface (CLI), see the monthly_index command.
Monthly Index Load Divisor	When the monthly index is so large that the memory issues occur when loading the terms during an Advanced Search and when Advanced Search and Monthly Index are both enabled, this parameter controls what fraction of the Monthly Index terms are loaded into memory during an Advanced Search. The default value is 1 (load all terms) with a possible range of 1-5, where 5 indicates 1/5th, which means load only 20% of the terms in memory.

Field	Description
	Note: Reducing the fraction of terms loaded into memory helps with memory issues. However, it has a performance impact on searches where no memory issues exist, because the frequency of disk access during an Advanced Search increases. For information about configuring the monthly index using the CLI, see the monthly_index_load_divisor command.

Index Search Settings

- **Tag Search**—The default is **Yes**. Allows to select keywords within log messages and to use the Tag Search feature.
- **Timeout**—The default is **10** minutes (600 seconds), allows user to set the timeout limit to retrieve Index Search results. Traditionally, this value was kept in a configuration file (`/loglogic/tomcat/webapps/logapp20/WEB-INF/web.xml`), now users can update this value from the GUI from the **Admininistation > System Settings** page. The value as seen on the GUI takes precedence over the `web.xml` file. Changing timeout values through the GUI requires no further action by the user, all necessary engines get the new values.
- **Message Encoding**—The default is **Windows-1252**. Allows user to choose different message encoding formats. The options are: Windows-1252, ISO-8859-1, ISO-8859-15, ISO-8859-2, and UTF-8. Select UTF-8 to process logs from Japanese Windows 2008 R2, Windows 2008, or Windows 2003.

  Note: You should not change the encoding to UTF-8 if you wish to parse non-UTF-8 logs.

Multithreading Parameters

Index search is multithreaded, and the behavior of multithreading is governed by the following configuration file:

`/loglogic/tomcat/webapps/logapp20/WEB-INF/spring/index-search.properties`

Independent of the number of threads, each thread works on a one-hour data bucket. The user cannot change the bucket size of 60 minutes.

If you change any value in the `index-search.properties` file, you must restart `mtask` by running the following commands:

1. `$ mtask stop`
2. `$ mtask start`

To take advantage of multithreading, the search range specified must be equal to or greater than the value of the `index.search.parallel.minHours` setting.

Note: These settings are not available on the GUI.

The following table lists the default values of the multithreading parameters:

Parameter	Description	Valid values	Default value
`index.search.parallel.enabled`	Enable/disable multithreaded index search	true, false	true
`index.search.parallel.minHours`	Minimum search time range in hours to trigger a multithreaded search Applicable only when `index.search.parallel.enabled=true`	value in hours	6
`index.search.parallel.numCountThreads`	Number of threads allocated for counting in a multithreaded search. Used in search to quickly navigate to the beginning offset of the page Applicable only when `index.search.parallel.enabled=true`	number of threads	3

Parameter	Description	Valid values	Default value
`index.search.parallel.numSearchThreads`	Number of threads allocated for searching in a multithreaded search Applicable only when `index.search.parallel.enabled=true`	number of threads	3

Syslog Port Settings

Set the Syslog UDP port and collector port numbers in this section.

If you want to add your own port numbers, ensure that you do not use the system reserved ports. For the complete list of reserved port numbers, see Port Assignments.

Field	Description
Syslog UDP Port	The syslog UDP port used for incoming messages. The default is 514. If you choose to use non-default ports and the firewall is enabled, then these ports must be added to the list of allowed ports in the **Administration > Firewall Settings** tab. To accept traffic on multiple ports, you must separate each port with a space. A maximum of 32 syslog UDP listening ports can be set. **Note:** If the firewall is enabled, run the `system firewall` command to add the customized syslog UDP port to the firewall rule.
Custom Syslog TCP Collector Ports	You can add a maximum of 14 port numbers separated by comma or space. After clicking Update, these port numbers are updated in the `/loglogic/cong/tcpcoll.conf` file. Any existing port numbers in the file are overwritten.

For more information, see Configuring TLS Syslog.

Secure ULDP Settings

From the **Administration > System Settings > General tab > Secure ULDP Settings** section, you can set the secure Universal Lossless Data Protocol (ULDP) settings.

By default, the **Secure ULDP** field is disabled. If you enable it, you can change the **Secure ULDP Port** and the **Minimum Secure ULDP Protocol** version fields. If you want to change the default secure ULDP port and if the firewall is enabled, then you must add the port to the firewall rule by running the firewall command.

The default value of **Minimum Secure ULDP Protocol** is TLSv1.2, and the other available option is TLSv1.3

TLS version	Supported with	Not supported with
TLSv1.3	Secure ULDP running on port 5515 or a non-default port	• FIPS with ULDP • TIBCO LogLogic® Universal Collector or TIBCO® Operational Intelligence Agent
TLSv1.2	• Secure ULDP running on port 5515 or a non-default port • FIPS with ULDP when FIPS is enabled • TIBCO LogLogic® Universal Collector or TIBCO® Operational Intelligence Agent	None

Note:

- TLSv1.1 and lower versions are not supported.
- To set up a secure connection with ULDP, perform the steps described in the Setting Up a Secure Connection With ULDP section in the *TIBCO LogLogic® Log Management Intelligence User Guide*.

Changing the TLS version

Before selecting the TLS version on the GUI in the **Minimum Secure ULDP Protocol** field, you must perform the following steps:

1. Generate the certificates required for establishing a secured connection with LogLogic® Universal Collector or TIBCO® Operational Intelligence Agent.
2. Install the certificates in LogLogic LMI.
3. Restart `engine_uldpcollector` by running the following commands:

```
$  mtask -s engine_uldpcollector stop
$  mtask -s engine_uldpcollector start
```

4. Select the required TLS version from the GUI.

Sending Logs from LogLogic Universal Collector

To send logs from LogLogic Universal Collector to LogLogic LMI you cannot use TLSv1.3; you must use TLSv1.2.

However, if you had set the **TLS version** field to TLSv1.3 earlier and then change it to TLSv1.2, then you must perform the following steps:

1. Stop `engine_stunnel` by running the following command:

```
$  mtask -s engine_stunnel stop
```

2. Get a list of all running `stunnel` processes by running the following command:

```
$ps -ef | grep stunnel
```

3. Terminate each running process by running the following command for each process ID obtained from the output of the `$ps` command in the earlier step:

```
$  kill -9 <process_ID>
```

4. Start `engine_stunnel` by running the following command:

```
$  mtask -s engine_stunnel start
```

5. Test the connection from LogLogic Universal Collector to LogLogic LMI.

Sending Logs from TIBCO® Operational Intelligence Agent

To send logs from TIBCO Operational Intelligence Agent to LogLogic LMI you cannot use TLSv1.3; you must use TLSv1.2. Follow the steps described in Forwarding Logs in the *TIBCO® Operational Intelligence Agent User Guide*.

Data Privacy Settings

Data privacy is about sharing data while protecting identifiable personal information from searches and reports.

Data privacy concerns exist wherever identifiable personal information is collected and stored in digital form. The legal protection for the right of data privacy varies around different regions/continents.

⚠ **Warning:** When using data privacy in Management Station, ensure that:

- All appliances are configured to use certificate authentication. This ensures that all queries and connections are authenticated and secured. This prevents an unauthorized Management Station from connecting to the appliance.
- All appliances have the following data privacy configuration:
  - The data privacy mode must be enabled on all Management Stations as well as the remote appliances that are connected to Management Station.
  - You must export the Column Manager settings that you configured on one appliance and import them for use on another appliance. This saves you from configuring the settings again on the other appliance.

The following sections describe how the data can be hidden for data privacy using the LogLogic appliance:

Enabling Data Privacy Mode

Procedure

1. Determine which data to hide from Searches, Reports and Alerts; for example, user name, social security number, host or domain name. Verify with the authorized legal representative about maintaining the data privacy compliance laws for your region.
2. Use the **Management > Column Manager** menu to define which columns to hide in the Data Privacy Mode. For more information on how to define columns, see Column Manager - Overview.

   Note: By default, some columns are hidden when Data Privacy mode is enabled. However, you can choose to hide more columns.
3. Ensure that two representatives are available to lock or unlock the Data Privacy Mode.
4. Navigate to **Administration > System Settings**.
5. On the **General** tab, under **Data Privacy Options**, click the **On** radio button to enable the data privacy mode to hide the columns from Index Search results and Reports menus.
6. Each representative must enter a security key and email address, and click **Apply**.

   Note: The key length must be between 6 and 64 characters.

Result

After enabling the data privacy mode, a lock icon is displayed on the upper right corner of the screen, and the following menu items are disabled:

Main Menu	Disabled Menu Items
Search	• Regular Expression Search • Advanced Search
Management	• Device Types • Column Manager • Advanced Features
Administration	• Message Routing • Data Files • Replay This menu item is available only on ST models.
Dashboards	Advanced Dashboards

Note: When data privacy mode is enabled, the following types of alerts are not displayed on the Show Triggered Alerts page:

- VPN Connection Alert
- VPN Statistic Alert
- VPN Message Alert
- Pre-defined Search Filter Alert
- Cisco PIX/ASA Messages Alert
- Network Policy Alert

Disabling the Data Privacy Mode

You can disable data privacy mode permanently or for the current user session.

To permanently disable the data privacy mode:

Procedure

1. Go to the **Administration > System Settings > General** tab.
2. Under the Data Privacy Options section, click the Off option button. The Data Privacy Mode is disabled.

To disable the data privacy mode for the current user session:

Procedure

1. Click the lock icon on the upper right-hand corner of the screen.
2. On the Authenticate Disabling of Data Privacy Mode window, both representatives must enter their Security Keys.
3. Click **Apply**.

When disabled for the current user session, the unlock icon is displayed on the upper right-hand corner of the screen.

Reset Security Key

Note: When the Data Privacy mode is enabled, the **Reset Security Key** link is displayed on the **General** tab next to the **Data Privacy Mode** buttons. The key length must be between 6 and 64 characters.

- **Reset Security Key**—Click the **Reset Security Key** link to reset the security key of the representative that was set for the Data Privacy Mode. Enter the **Old Security Key, New Security Key**, and **Retype New Security Key** fields. These are mandatory fields. Click **Apply** to use the new security key.
- **Forgot Security Key**—Click the **Forgot Security Key** button if the representative has forgotten the security key for the Data Privacy Mode. A new password is generated and emailed to the requesting authenticating user. The confirmation window appears. Click **Yes** to receive an email for the newly generated password.

  Note: To receive an email with a new password, you must enter an email address in the Originating Email field; otherwise an error message appears.

Global Retention Settings: Raw and Indexed Data

You can specify the global data retention time for raw and indexed data.

These settings are global to the appliance, whereas the default retention rule values under the **Administration > Data Retention Rules > Default Data Retention Rule** section apply to all rules except custom rules.

> ⚠ **Warning:** Changing the values in this section and clicking **Update** to save the changes immediately updates the default retention rule values under the **Administration > Data Retention Rules > Default Data Retention Rule** section.
>
> However, the reverse does not happen - changing the default retention rule values under **Administration > Data Retention Rules > Default Data Retention Rule** does not affect the values under **Administration > System Settings > Global Retention Settings: Raw and Indexed Data**.

Additionally, the following settings control index retention purging. The archiver is responsible for initiating the index purging process.

- **Index Retention Purge Interval**: this value controls how often the system checks for both the removal of expired index folders and the removal of expired documents from indexes. It is set in units of seconds. The maximum permitted value is 86400 seconds.
- **Skip Index Purge on Start**: When true, this flag suppresses the index purging for the first pass when the archiver is started.

Default Values

The following table lists the default values of the fields. The default values of the **Raw Data** and **Indexes** fields are different for different appliance families.

Appliance Family	Raw Data	Indexes	Index Retention Purge Interval	Skip Index Purge on Start
LogLogic ST Appliance	10 years	10 years	7200 seconds	No

Appliance Family	Raw Data	Indexes	Index Retention Purge Interval	Skip Index Purge on Start
	(3650 days)	(3650 days)	(every two hours)	
LogLogic LX Appliance	3 months (90 days)	3 months (90 days)	7200 seconds (every two hours)	No
LogLogic MX Appliance, MX Virtual	1 year (365 days)	1 year (365 days)	7200 seconds (every two hours)	No

For other archive settings, see Archive Config Tab.

Scheduled Report Settings

- **Bundle Report Emails**—Allows concatenation of multiple emails into a "bundle" to reduce the number and frequency of email messages sent to Administrators and other users.
- **Compress Attachments**—Permits compressing of large email attachments when they exceed a predetermined file size (in KB) specified in the **When Attachments Exceed** field.

SNMP Settings

Option	Description
SNMP Daemon	By default, the SNMP Daemon is disabled to maintain system security. This option must be enabled if the appliance uses an SNMP Trap Sink. Select **Yes** to enable this option when you reboot the appliance. See SNMP Trap Sink Settings.
Inbound	Type a private or customized community string for your appliance. The

Option	Description
Polling Community String	default is 20 characters. **Note:** It is good practice to change the default value of Inbound Polling Community String because multiple instances of LogLogic EVA must not have the same value.

SNMP Trap Sink Settings

The SNMP agent generates startup and shutdown traps to help monitor the appliance.

These two traps are sent to the SNMP manager configured in the SNMP Trap Sink.

Starting from version 6.2.0, LogLogic LMI supports SNMP Trap versions 1, 2c, and 3. Depending on the version, you must specify the following fields:

Field	Description	Applicable to SNMP versions
IP Address:Port	Type the IP address and port number of your SNMP trap receiver.	Versions 1, 2c, and 3
Outbound Traps Community String	Type the SNMP trap receiver community string (public or private)	Versions 1 and 2c

This feature is active only if you set SNMP Daemon to On.

SNMP Trap Collector Settings

The SNMP Trap Collector supports SNMP versions 1, 2c, and 3. The port number for trap collection must be specified in the **Ports** field. The default port is 162. You can enter up to five port numbers, separated by a space or comma.

After entering port numbers in the **Ports** field and clicking **Update**, you must select the option to restart the appliance for these settings to take effect.

You must add the engine ID, user name, and password of each SNMP v3 trap sender in the file `/loglogic/conf/snmpv3security.conf` on the receiving LogLogic LMI appliance in the following format:

```
<EngineId> < User Name> <Password>
```

Run the following command to restart `engine_trapcollector`:

```
mtask -s engine_trapcollector restart
```

Note: When sending v3 traps with authentication, the LogLogic LMI trap sender uses the following credentials:

- engineId: 0x80001f888058edf1347e148a5900000000
- user name: snmpv3user
- password: ELDo2GE10o2g5I7c

Syslog Forwarding Settings

If log data that includes IPv6 addresses with the IPv6 prefix are forwarded from LogLogic LMI to Syslog devices, sometimes the logs might not be parsed correctly at the receiving log source. To avoid this, you can remove the IPv6 prefix from logs before forwarding them.

Procedure

1. On the **Administration > System Settings > General** tab, enable the **Remove IPv6 From Logs** field.
2. Click **Update** to save your changes.

System Performance Settings

System performance settings are as follows:

- **Remove PIX Active IP Connections**—Select the duration of time the appliances retains messages for IP connections. Connections that do not terminate properly are

stored in the database until the expiration time is reached. This relates to the Real-Time, Active Connections report. To free more space on the appliance, set this threshold low.

- **Concurrent Regular Expression Searches**— (applies only to appliance models over the 1000 series) Select the number of concurrent searches to perform. The default and maximum number of concurrent searches possible are specified in the following table.

  The more the number of regex searches that are concurrently running, the longer it takes to execute them because they share the same resources.

	Model Name	Default Searches	Maximum Searches
MXVirtual (LogLogic EVA)		1	2
H4 R1 Models	LX1025R1	1	2
	LX4025R1	12	12
	ST4025R1	12	12
	ST2025-SANR1	12	12
H4 R2 Models	LX1025R2	1	2
	LX4025R2	12	12
	ST2025-SANR2	12	12
	ST4025R2	12	12

	Model Name	Default Searches	Maximum Searches
H5 Models	LX1035	1	2
	LX4035	12	12
	ST2035-SAN	12	12
	ST4035	12	12

- **System Maintenance Start Time**—Select the time to start system maintenance. The default is 2:00 AM. This activity is logged in the General Syslog.
- **Refresh Auto-Identified Device Interval**—Select the refresh time, in days, that the appliance checks for new auto-identified log sources that you add.
- **Enable Daily AD User Cleanup Task**—Select this check box and specify the time when the daily task of cleaning up Active Directory (AD) users must run. When an AD user account is used to login to LogLogic LMI, a corresponding user is created in LogLogic LMI. The task runs if this check box is selected and if the user credentials are specified under the **Administration > System Settings > Remote Servers** > Setting up Active Directory section, and deletes these corresponding LogLogic LMI users from the **Management > Users** tab in any of the following scenarios:
    - The AD user is disabled or deleted on the AD server
    - The AD user is removed from all associated roles or groups on the AD server
- **Sort the Management Station Status by**—Select the default sort order for displaying remote appliances (RAs) in the Management Station on the **Dashboards > Management Station Status** page. The default sort order is **ID**. Other options are:
    - Model
    - IP Address
- **Optimize Device Selection List**—**Show all Source Devices** is the default. If the appliance has more than 4,000 devices, selecting **Show Only Device Groups** improves display performance on many GUI pages. **Show Only Device Groups** limits device selection lists to device groups; individual devices do not appear in the device selection lists. This selection affects:

- All LogLogic LX Appliance or LogLogic MX Appliance Real-Time and Summary report filter pages
- Search Archived Data, Real-Time and Scheduled Search tabs
- Devices tab for Alerts
- Devices tab for Manage Users
- Message Routing
- Import/Export

Usage Count

You can view the amount of data that has been ingested and indexed over time, so that you can charge a subletting customer accordingly. Navigate to **Dashboards > Advanced Dashboards > Advanced System Status** to view a graph of ingested versus indexed data.

Note: If you change any of these values, restart `mtask` either by following the instructions on the screen, or manually by running the following commands:

1. `$ mtask stop`
2. `$ mtask start`

Configuring Data Usage Count

Field	Description	Default Value	Permitted value
Ingest Count Update Period	The frequency at which to calculate the total ingested bytes of data since the previous calculation	5	Between 5 and 60 minutes
Count Retention	Retention period in days,	365	Between 90 and 540 days

Field	Description	Default Value	Permitted value
	for storing the ingest and index count information. The default value is 365 days.		
Index Size Update Period	The frequency at which to calculate the change in disk space of indexed data	5	Between 5 and 60 minutes

Theme for Rebranding LogLogic LMI

From the **Administration > System Settings > General tab > Theme** section, you can configure the background colors of the header bars and font color of the header text.

Changing the theme includes setting the following colors on the GUI:

Field	Default value (hex)	Description
Primary Header Color	#0080cb	Color of the title bar
Secondary Header Color	#3498db	Color of the breadcrumb bar
Font Header Color	#aaeaff	Font color of the menu items and the text in the breadcrumb bar **Note:** The color of menu items changes only when you hover over the menu item.

You can change the theme immediately after upgrading or at any other time. Perform the following steps:

1. Select the color from the color palette or type the hexadecimal value of the color.
2. For the changes to take effect, either refresh the page, or log out and log in again.

Custom Logo Upload Settings

Select the image logo to be used on the GUI. The image format can be .bmp, .gif, .jpeg, or .png.

Field name	Location of logo	Maximum image size (pixels)
Login Screen Logo	• Login page • Change password page that is displayed when the user is prompted to change the password	225 x 86
Screen Logo	Upper-left corner of the GUI	225 x 86
Report Logo	Upper-right corner of a report after saving the report	300 x 105

If you are using custom logos, you can now restore default logo images at any time. This can be done for all logos or selected logos.

To do this, go to the **Administration > System Settings > General** tab. In the **Custom Logo Upload Settings** section, restore the Login screen, Screen, and Report logos to the default logos.

- To restore selected logos, select the corresponding **Restore to default** check box.
- To restore all logos to default images, select the **Restore all to defaults** check box.

Build Details

- **Build Number**—Displays the currently running build number for the appliance software.
- **Log Source Package (LSP)**—Displays version information of the installed LogLogic LSP.

Remote Servers Tab

On the **Remote Servers** tab, identify and define the SMTP and Remote Authentication servers you wish to use with your appliance.

Enter the server information in the **SMTP** and **Remote Authentication Server** sections as applicable.

Settings automatically take effect after you click **Update**. A message displays on your screen after the remote server is successfully set up.

Note: Changes to certain settings might require a reboot to take effect. A message is displayed if a reboot is required.

Related topics

- Configuring SMTP
- Setting up a Remote Authentication Server

Configuring SMTP

The SMTP settings are used to specify an outgoing email server for the appliance.

Each appliance typically uses one email server to send alert notifications and scheduled reports. If you do not define a server, the appliance cannot send email.

Procedure

1. Go to **Administration > System Settings > Remote Servers**.
2. In the SMTP section, enter the following information.

Field	Description
Server	Enter the IP address of the mail server.
Port	Enter the SMTP port for the mail server if you want to change the default value. The default port number depends on the protocol selected in the **Protocol** field.
User ID	Enter the user name for the mail server, if it requires one
Protocol	Select the protocol to be used by the SMTP mail server while sending notification emails. Depending upon your selection, the port number also changes. The available options are: • None (port 25) • TLS (port 465) • STARTTLS (port 587) **Important:** If you select the protocol as **TLS** or **STARTTLS**, ensure that: • You also configure TLS on the SMTP server • On the LogLogic LMI appliance, you use the Java library that supports the highest TLS version supported in LogLogic LMI.
Password	Enter the password for the mail server, if it requires one.
Verify Password	Enter the password for the mail server again to confirm the password.

3. Click **Update** to save your changes.

What to do next

If any of the SMTP settings or the Originating Email field are changed, you must re-synchronize triggers.

Setting up a Remote Authentication Server

The remote authentication server settings let the LogLogic LMI appliance participate in a centralized login authentication implementation.

TACACS, RADIUS, or Active Directory can be used by the appliance as authentication servers to verify user credentials. Active Directory also allows defining roles for groups of users, to which you can assign specific user privileges and access to specific log sources. You can define up to eight remote authentication servers.

If you have configured multiple authentication servers, then the appliance attempts to authenticate a user using all servers one after another until the user is finally found on a server and authentication is successful. For example, If the authentication attempt fails on server1 (for example, the user does not have an account on server1), then the appliance attempts to authenticate the user on server2, and if necessary and applicable, on server3 and server4, and so on. An error is logged for each server on which the user is not found.

Before you begin

- Add the appropriate users to the remote authentication server or ensure that their login IDs already exist.
- For TACACS or RADIUS:
    - Add the IP addresses of the LogLogic LMI appliance to your remote authentication server.
    - If you have a failover configuration, you must add the private IP addresses from both appliances to the remote authentication server. For details, see Failover.

Procedure

1. Go to **Administration > System Settings > Remote Servers**.
2. Select the appropriate option for the **Remote Authentication Servers**:

Option	Description	Default Port
RADIUS	Configure a remote server using the RADIUS protocol. RADIUS is used only for access control, and not for RADIUS accounting.	1812
TACACS	Configure a remote server using the TACACS+ protocol.	49
Active Directory	Configure a remote Active Directory server. The appliance roles associated with groups on the Active Directory server are used for controlling appliance authentication, user privileges, and access to log sources.	389

What to do next

Continue with the procedure for the remote authentication server option that you selected.

- Setting up TACACS or RADIUS
- Setting up Active Directory

Setting up TACACS or RADIUS

Procedure

1. In the **Server Name** text field, enter the name of the remote authentication server.
2. Select the **Enable** check box to enable this remote authentication server for the appliance after you click **Update**.
3. In the **Server IP** text field, enter the IP address for the remote authentication server.
4. In the **Port** text field, enter the port number for the remote authentication server, unless you want to use the listed default.
5. In the **Secret** text field, enter the associated password for the remote authentication server.
6. In the **Verify Secret** text field, enter the associated password again for the remote authentication server.

7. In the **Timeout** text field, enter the amount of time (in seconds) the Appliance waits for a reply from the authentication server.
8. From the **Method** radio buttons, select an authentication method.

   The selections vary depending on your original selection for the Remote Authentication Server. For the RADIUS and TACACS authentication methods, options include:

   - PAP
   - CHAP
   - MSCHAP
   - MSCHAP2 (RADIUS only)
9. Click **Update** to save your entries or changes.

Setting up Active Directory

Procedure

1. In the **Auth Type** field, select the type of authentication for the Active Directory server to perform:
   - Kerberos
   - Simple Authentication (username/password)
2. Depending on the type of authentication you selected, enter the information in the relevant fields.

Field	Description
Server Name, Server IP	Name and IP address of the remote authentication server **Simple Authentication** • Enter either the server name or the IP address. • If server name is not entered, it is treated as an empty string, and not as NULL. **Kerberos**

Field	Description
	• The server name cannot be empty. • Add the server name and IP address in the `/etc/hosts` file in the following format: `<IP_address> <ServerName>`
Enable (check box)	By selecting the check box, this remote authentication server is enabled for the appliance after you click **Update**.
Port	Enter the port number for the remote authentication server, if you want to change the default value.
Enable SSL (check box)	By selecting this check box, secure connection is established on the AD server. Ensure that you have the certificate file of each AD server and that the certificate is added to the trust store. See step 5.
Realm	The realm for the remote authentication server For example: `SQA2008R2a.lab`
NT Domain	Domain name of the remote authentication server For example: `SQA2008Ra` Applicable only to the Simple Authentication type.
User, Password	Credentials of any user who has access to the Active Directory server. This is required so that when the daily AD user cleanup task runs, if the users with remote authentication are removed from all associated roles or groups, or are disabled or deleted from the AD server, the corresponding users are also removed from the **Management Users** tab.

3. Click the **Test** button to test the connection to the specified Active Directory server.

   a. When prompted, enter a login name and password of any user for the server and click **Test Connection**.
   b. The pop-up remains open to display the status of the test. If the connection test times out (after fifteen seconds), a time-out message appears in the **Connection Status** box on the pop-up.
4. Click **Update** to save your entries or changes.
5. Import the AD server certificates:
   a. Go to the **Administration > SSL Certificate > Trusted Certificates** tab.
   b. Each server has its own certificate. Paste each server's certificate in the **Import Trusted Certificate** box. Each certificate must begin on a new line.
   c. Click **Import**.
   d. Click **Yes** to confirm restarting the GUI and wait for the GUI to restart.
6. After adding the certificate to the trust store, in the file `/loglogic/tomcat/bin/setenv.sh`, disable endpoint verification by setting the value of the `JAVA_OPTS` parameter:

   ```
   JAVA_OPTS="$JAVA_OPTS -
   Dcom.sun.jndi.ldap.object.disableEndpointIdentification=true"
   ```

   Note: If a different value of `JAVA_OPTS` is already configured in the file, add this line in the file after the existing line.
7. Restart Tomcat for the `JAVA_OPTS` settings to take effect.

Database Table Retention Tab

Use the **Administration > System Settings > Database Table Retention** tab to configure the length of time the appliance stores database contents before purging.

The **Database Table Retention** setting is applicable to LogLogic LX Appliance, LogLogic MX Appliance, LXvirtual, and MXvirtual appliances. The setting for ST appliances can be configured in the MySQL configuration table.

When is data purged?

Data stored in the appliance database is purged in the following scenarios:

- On expiration: All database data eventually expires and is deleted. The expiration is specified in the Duration field on the **Administration > System Settings > Database Table Retention** page.
- On emergency: When the disk usage of `/loglogic` reaches a threshold value defined in `LxDiskUsagePercent`. See Changing the Database Purging Threshold. The same parameter is used by the scheduler for both LogLogic LX Appliances and LogLogic ST Appliances.

The Scheduler

The scheduler is involved in data purging and handles emergency purging of MySQL tables. A table with the oldest data is removed first, regardless of log type or table type. However, eventually all event data tables are purged if disk usage for `/loglogic` does not fall below the threshold setting value.

The scheduler deletes all data stored in the database tables and the `engine_archive` process does not start performing emergency purging as long as the scheduler is still purging database tables.

The archiver

The archiver handles purging of raw and index data. A file with the oldest expiration time is deleted first. To set the retention time of raw logs and index data, see Data Retention Rules.

The archiver begins purging raw and index data when the scheduler sets a flag to notify the archiver that the scheduler has deleted all the event data it can from the database. If the disk usage of `/loglogic` is still above the emergency purge threshold and the scheduler sets the flag, `engine_archive` begins the deletion it is responsible for. It is best practice to adjust the emergency purging threshold of the scheduler to match the emergency purging threshold of the `engine_archive` process.

The archiver threshold depends on the type of appliance and whether failover is configured or not.

The default value of the table threshold is 90% without HA and 80% for HA pair, which matches the lowest of all platform models for this threshold. This value results in balanced data removal in most cases. To change the default value of the table purging threshold, see Changing the Database Purging Threshold.

Retention settings

From the **Administration > System Settings > Database Table Retention** page, you can select the number of days or disable retention for:

- File Transfer History Settings
- Database contents, by table category (indexed Alert, Report, and Statistic Tables)

The available numbers of days to select varies for different table categories. Settings made in this tab override the default settings for this appliance model.

Disabling retention settings immediately purges the data collected for that file or database category.

To set the retention time period for raw and index data for LogLogic LX Appliance, LogLogic MX Appliance, and LogLogic ST Appliances, see Creating a New Retention Rule.

It is good practice to simplify retention settings by using consistent retention durations throughout to avoid confusion. However, you can use data retention settings to ensure the data you need most is retained longer than data you don't need as much. For example, if you need Authentication data and don't need SEC Deny IP data, set Authentication to a higher duration and SEC Deny IP to a lower duration.

For archival purposes, confirm that the appliances are receiving log data. It is best practice to have a copy of the older data, if you reduce retention settings. For details, see Forwarding Logs to Other Appliances (Routing) and Log Source Management.

For each table category, the **Database Table Retention** tab displays the **Table Category**, **Duration** in number of days, the number of **Database Entries**, and **Data Size** for each category.

To deactivate data or file retention, instead of letting the appliance purge files as disk space fills up, you may select **Disabled** from the **Duration** list for that file or database table category.

> **Warning:** Selecting **Disabled** sets the file retention period to zero, and the data is immediately purged from the appliance for that file or database category.

Click **Update** to save your changes.

Changing the Database Purging Threshold

You can set or change the threshold when data in the database tables are purged for emergency purposes. This is the emergency purging threshold.

Procedure

1. Login to the command line interface.
2. Create a file `/loglogic/conf/sched.conf` if it does not exist.
3. Add the following line in the file `/loglogic/conf/sched.conf`:

   `LxDiskUsagePercent=X`

   The valid range for X is 0 ~ 100.
4. Restart `engine_lx_scheduler` by running the following command:

   `mtask -s engine_lx_scheduler restart` (On an LogLogic ST Appliance the engine name is `engine_st_scheduler`.)

   > **Warning:** Restarting `engine_lx_scheduler` triggers a failover. If it is fine in your setup, then change the threshold on the passive appliance first and then on the active appliance. Doing so causes only one failover.
   >
   > To avoid a failover in your setup, perform the following steps:
   >
   > a. Stop the `engine_cluster_membership` service on the passive appliance.
   >
   > b. Change the threshold on the active appliance and restart the scheduler.
   >
   > c. Apply the change on the passive appliance, restart the scheduler, and then start the `engine_cluster_membership` service.

Setting the Time Zone and Time

From the **Administration > System Setting > Time** tab, you can modify the time for the appliance or configure a Network Time Protocol (NTP) server.

Note: This section is not applicable to LogLogic® EVA - Container Edition.

If you do not set the correct time for your appliance, the system does not function correctly.

Procedure

1. Go to **Administration > System Settings > Time**.
2. From the **Time Zone** list, select a time zone for your network.
3. To enable synchronizing your local time, select the **Update Time** check box. Then select how to update the appliance time in one of the following ways:
   - (Recommended) To synchronize your local time with that of an NTP server, select **NTP Server** and enter a host name or IP address for the NTP server.
     - If you have multiple appliances connected together, you must set up a common external NTP server for all appliances to ensure that the time on all appliances is synchronized. Ideally, this is the same NTP server used by the entire data center.
     - If you do not have access to a common external NTP server, you can use the appliance running as a Management Station as the common NTP server. The time settings of all appliances must be in sync.
     - It is important to have an NTP server set up for a single appliance, too.
   - To manually define the system time, select **Specify Time** and enter the system time (MMDDhhmmYY.ss).

     Note: When configuring LogLogic EVA on cloud platforms, a static time configuration is not supported. Use the **NTP Server** option instead.
4. To save your changes, click **Update**.

   The appliance prompts you that an immediate reboot is required.

5. To let the appliance reboot for changes to take effect, click **OK**.

Customizing the Login Page

Use **Login Page** to customize text, enable, disable, or preview the look of the appliance's login page.

Login customizations also appear in the CLI login.

Customizations made on this page can be copied over to other appliances in a Management Station cluster.

Procedure

1. Select **Administration > System Settings > Login Page**.
2. Click a radio button to **Enable Login Page Content** allowing users from viewing the login content on your login page. The default is **No**, disable.
3. In the **Login Page Title** text field, enter a title for the login page. The maximum number of characters and spaces is 35.
4. In the **Login Page Content** text box, enter the text users view when logging in to the appliance. The maximum number of characters and spaces is 3000.
5. The tab displays the number of characters remaining as you type. The counter might not accurately reflect the word count if you copy and paste text into the field.
6. Click **Preview** to view the login page with your changes, before updating the page.
7. Click **Update** to update and save the appliance login page.

> **Note:** If the appliance is in a Management Station cluster, clicking **Update All** updates all appliances in the cluster with the changes.

Password Control Tab

Use the **Password Control** tab to enable or disable the requirement of strong passwords for users who log in to the GUI of this appliance.

Strong passwords provide greater security for appliance access by requiring more complicated passwords and imposing time constraints on them.

You can choose from several password controls to require of users. The more password controls you enable, the stronger the password security is on the appliance.

With strong passwords enabled, you can specify requirements for any or all of the following:

- **Enable Password Control**—Selecting this check box requires all user passwords to use strong passwords. This setting controls only the GUI passwords. Strong passwords for CLI accounts are controlled through the `show strong_passwd` command.
- **Enforce Password Length Rules**—The minimum and maximum number of characters required and allowed in a password (default is15). This is not strictly enforced if the length field is NULL.
- **Enforce Password Character Rules**—The minimum number of characters required in a password for different character types: uppercase letters, lowercase letters, digits, and non-alphanumeric characters (default and minimum for each = 1)
- **Enforce Password Expiration Rules**—The number of days after which a password expires (1 through 99999)
- **Enforce Password Management Rules**—The number (3 to 100) of a user's previous passwords that a new password cannot match. Three previous passwords is the default value.
- **Enforce Account Lockout Rules**—The number of failed login attempts, after which the user's account is locked, and the duration of the lockout time interval (default = 1440 minutes, i.e. 1 day)

  Locked out accounts are automatically available again after the time interval expires. To unlock an account sooner, reset this value to a low setting. The account is unlocked after this lower setting expires. After that time, reset this value back to its original higher setting.

When you enable strong passwords on the appliance, each existing user is prompted for a password change upon their next login if the existing password is not a strong password.

Archive Configuration Page (LogLogic ST Appliance and LogLogic EVA Only)

The Archive Configuration option in the Administration menu appears only if the appliance model supports it. The LogLogic ST Appliance or LogLogic EVA instance support archiving to an NFS remote server, and some LogLogic ST Appliances support archiving to SAN storage.

The purpose of archive mapping is to allow the user to manually migrate archived data to a larger disk, and then replace the existing disk with the larger one. For example, the user would do this when all of the possible 32 mounts points are used up, but more space is still needed. However, this is a feature for advanced users only, since it requires a lot of manual work.

Archive Config Tab

The Archive Config tab on the **Administration > System Settings** page is available only on LogLogic ST Appliance and LogLogic EVA.

When active and disk space is needed, the archiver performs repeated archiving loops, checking which files should be archived first, if any, for each pass. The following settings are available for controlling this process:

File counts to archive per pass

Setting	Description	Default value
Data files by time (oldest first)	Controls how many log data files should be archived per pass of the archiver, giving the highest priority to the oldest files.	10
Data files by size (largest first)	Controls how many log data files should be archived per pass of the archiver, giving the highest priority to the largest files. The default value 0 indicates that archiving by size is turned off.	0

Setting	Description	Default value
Index files (0 disables Index archiving)	Controls how many index folders are archived per pass of the archiver. Setting this value to 0 disables index archiving.	1

When using both of the data file settings at the same time, the archiver first moves the largest files, after which it moves the oldest files. The default is to always move the 10 oldest files per archiving pass.

As the index folders contain the indexed log data for many log data files, the setting for index files to archive per pass is typically much smaller than the settings for log data. The default ratio is 10 log files to 1 index folder per archiver pass.

Index Archiving

By default, both raw and indexed data is archived. Under some circumstances, an appliance might archive all indexed data to remote storage. If you want to optimize index searches and you have archiving enabled, you can disable index archiving to ensure that indexed data remains on local storage.

User Accounts

Administration of a TIBCO LogLogic® appliance includes creating, modifying, and deleting user accounts.

User accounts are configured with specific appliance privileges, and users can be associated with specific log sources attached to the appliance.

From a Management Station, you can manage users on remote LogLogic appliances. You can also replicate users and roles to remote appliances, so the user accounts and role configurations on the Management Station are automatically configured on the remote appliances as well.

If you use an Active Directory server for remote authentication, you need to configure roles on the appliance. These roles map to Active Directory (AD) groups, letting you allow AD groups and users appliance access and control their privileges and log source access like any appliance user.

Note: User privilege on the remote appliance or product is required in order for the user to remotely control the appliance or product. See Adding or Modifying a User.

Managing Users

Administrators can manage the users in the system, set user privileges, and control user sessions.

Setting the **Allow Disabling of admin Super User** option under **Administration > System Settings > General** tab indicates that disabling the default super user admin is allowed.

Important:

- Only an admin super user can view or change the value of this option. Other users cannot view this option under **System Settings > General** tab even if they have the full set of privileges.
- Selecting the option Yes indicates that the admin user can be disabled by a user who has the full set of privileges. After selecting Yes, the admin user can be disabled using the Enable option from the **Management > Users > Edit User** page.
- When the admin user is enabled, any user, regardless of privileges cannot access the **Management > Management Station** page. However, after the admin user is disabled, users can access this page.
- There is always one administrator user in the system that cannot be deleted or disabled.

Users Tab

The **Users** tab lists all the user accounts on the appliance. You can access the **Users** tab from **Management > Users**.

If users are listed but not selectable, they are Active Directory (AD) users who have logged on to the appliance before. AD users have access only if AD servers are configured as remote authentication servers for the appliance and roles are defined on the appliance.

AD users are controlled on the AD server. From the appliance GUI you cannot add, modify, or remove users from the AD server itself. Once an AD user logs in to the appliance, that AD user automatically appears in the **Users** tab.

- To add a new user to the appliance, click **Add New** button. The **General** tab of the Edit User page opens.
- To modify an existing user on the appliance, click the required User ID.
- To remove a user from this list, select the user's check box and then click the delete icon .

> **Note:** Removing a user with remote authentication from the **Users** tab removes the user entry only from the appliance users list, and not from the AD server itself.

However, when a user with remote authentication is removed from all associated roles and groups, or is disabled or deleted on the AD server, it is removed from the appliance users list on the Users tab, if both the following conditions are true:

- The user name and password of a user who has access to the AD server are configured in the **Administration > System Settings > Remote Servers >** Active Directory section.
- The daily AD user cleanup task is scheduled to run.

After the user is deleted, an entry is logged in the `sys.log` file.

User Devices/Privileges Report Tab

The **User Devices/Privileges Report** tab allows the LogLogic admin to see the privileges assigned to all users of the appliance.

All appliance users are listed by default on this page. Users without User Admin privilege cannot access the **User Devices/Privileges Report** tab.

Advanced Options provide a way to report on a selected set of users and privileges. The **Advanced Options** function on the **User Devices/Privileges Report** tab opens four select boxes; the two boxes on the left contain all Available Users and all Available Privileges, and the two boxes on the right are for Selected Users and Selected Privileges. Arrows provide a means to move a single entry (single arrow) or all entries (double arrows) between the left and right select boxes.

After selecting the required items, click the **Run** button to see the **Privileges** Report. As with other reports, the **Privileges Report** can be saved in CSV, HTML, or PDF format.

User Sessions Tab

From the **User Sessions** tab, the LogLogic Administrator can view the currently active user sessions.

The user ID and time of last request is displayed in a table. Users without User Admin privilege cannot not access this tab.

In the following situations, the LogLogic Administrator might need to remove active user sessions by clicking the remove icon :

- The number of concurrent login sessions for a specific user has reached the maximum number configured in the **Concurrent Login Sessions** field.
- A user who abruptly ended the login session and has reached the maximum concurrent session limit is unable to login again.

Adding or Modifying a User

To add or modify a user, use the **Management > Users > Users** tab.

- To add a new user, click the **Add New User** icon. When adding a new user, the tabs must be accessed and completed in sequence.
- To modify an existing user, click the user name in the **User ID** column. When modifying a user, you can access any tab because they are already filled from when the user was added.

When filling out the tabs, you can switch between the tabs without clicking **Add** or **Update**. Your entries are retained until you click **Add**, **Update**, or **Cancel** on one of the tabs. For example, when adding a user:

- Complete the **General** tab.
- Complete the **Privileges** tab.
- Complete the **Devices** tab.

Click **Add** on any of these tabs.

The **Users** tab appears, displaying the new user in the list.

Procedure

1. Go to **Management > Users > Users**.
2. Start from the **General** tab and enter information or edit the required fields.
3. Click the **Privileges** tab.

   You can select or edit user privileges for the new or existing user. You can assign any

or all privileges visible on the tab to a user, including that of top level Administrator.

4. Click the **Devices** tab and select the desired devices and device types to be monitored by the user. By default, Device Type is set to “All” on the **Devices** tab.
5. Click the down arrow and scroll through the list of **Available Devices** to select the desired device and device type for monitoring. Alternately, hold down the Shift key or the Control key to select multiple devices and types from inside the **Available Devices** pane.
6. Move the highlighted devices and types to the **Selected Devices** pane by clicking the single right-pointing arrow (**>**). All highlighted devices and types move to the right as a group. Clicking the double arrows (**> >**) causes all available devices and types to move to the Selected Devices pane.

Defining or Updating Users

Use the **General** tab to define or update the user on the appliance.

Procedure

1. In the **User ID** field, enter a user name. (`null` is not a valid value for this field.)
2. In the **First Name** field, enter the first name of the user.
3. In the **Last Name** field, enter the last name of the user.
4. In the **Email** field, enter the contact email of the user.
5. In the **Phone** field, enter the contact phone number of user.
6. Select the appropriate **Authentication** radio button to choose local or remote authentication. The default is local. If you select remote, the **Password** fields disappear and you can proceed to step 9.

   A remote user is authenticated from a server such as a TACACS or RADIUS server. Information related to authentication, such as passwords, is stored on the remote server, not on the local appliance.
7. In the **Password** field, enter a password used for user login.

Note:
- Users without administrator privileges can change their passwords from the top navigation bar by navigating to the > **Your LogApp Account > Change Password** dialog.
- By default, the password must be at least six characters, containing at least one non-alphabetical character, and cannot be the same as the user ID. Spaces at the beginning and end of the password are removed before being stored in the system. The password requirements might differ if password control has been enabled in LogLogic LMI.

8. In the **Verify** field, enter the password again to confirm the password.
9. In the **Enable** field, select whether to enable the user's account to provide appliance access. If you disable a user account, the account information remains on the appliance but the user has no access.
10. Click **Save** after you have completed all **Users** tabs. (Your settings are retained as you move between these tabs while adding or modifying a user.) To cancel adding or modifying the user and return to the Users tab, click **Cancel**. Changes made on any **Users** tabs for this user are not retained.

User Privileges

Use the **Privileges** tab to grant various privileges to users when adding or modifying them.

User privileges control access to different parts of the appliance. Navigation menu items display according to the level of privileges granted to a specific user.

When adding a user, the User tabs are dependent on each other. You must complete the **General** tab before accessing the **Privileges** tab. After you complete the required fields in the **Privileges** tab, you can continue to the **Devices** tab.

When modifying a user, you can access any tab because they are all filled from when the user was added.

Note: Depending on your appliance selection, you might see only a subset of these options in the navigation menu of your appliance.

Who Can Grant Which Privileges

When creating a user on the appliance, you can grant privileges only less than or equivalent to the privileges you have. For example:

- An **Administrator** can grant any privileges to other users.
- Only user 'admin' can access **Management > Management Station** menu to configure cluster. Any other users with Administrator privilege cannot view this menu.
- To grant the **Administrator** privileges, you must have the **Manage Administrators** privilege.
- To grant all the **User Admin** privileges, you must have both the **Manage User** and **Manage Administrators** privileges.
- To grant the **Manage User** privilege, you must have either the **Manage Administrators** or **Manage User** privilege.
- To grant the **Manage Administrators** privilege, you must have the **Manage Administrators** privilege.
- To grant the **Replicate User** privilege, you must have the **Replicate User** privilege.

See List of User Privileges.

List of User Privileges

Select the check boxes to grant specific privileges under each category, or clear the check boxes to remove privileges.

The following privileges are available:

- Administrator
- User Admin
- Report Admin
- Config Admin

Important Considerations

- If privileges are greyed out, the user might be created on a remote authentication server (Active Directory) and not on the LogLogic LMI appliance.
- An appliance administrator can modify the privileges of remote users from the **Management > Users > Users** page. To modify privileges for an AD-based user: On the **Directory Roles** tab, click the role name and click the **Privileges** tab. For more information, see User Roles.
- All appliance options are listed in the following tables. However, depending on the appliance selection, you might see only a subset of these options in the navigation menu for your appliance.

Administrator

Administrator privileges for all categories.

User Admin Privileges

Privilege Option	Description
Manage User	Lets users create, modify, and remove users or roles on the appliance.
Manage Administrators	Lets users create, modify, remove user accounts with administrator privileges.
Replicate User	Lets users replicate user accounts on the Management Station appliance to attached remote appliances.

Report Admin Privileges

Privilege Option	Description
Real-Time Reports	Lets users create, modify, remove, and run Real-Time reports.
Search Archived Data	Lets users search log data captured by the appliance. Also, allows users to Replay archived data from its archived location.
Access Custom Reports	Lets users access Custom Reports. This option controls the Add/Modify/Delete Custom Reports and the Run/Schedule Custom Reports menus. If you enable or disable this privilege, both Add/Modify/Delete Custom Reports and the Run/Schedule Custom Reports options are automatically enabled or disabled. **Note:** For any user to view **Search > All Saved Searches > All Index Reports**, and **All Search Filters** menus, this privilege must be enabled.
Add/Modify/Delete Custom Reports	Lets users add, modify, and delete Custom Reports.
Run/Schedule Custom Reports	Lets users run and schedule Custom Reports.

Config Admin Privileges

Privilege Option	Description
Manage Devices	Lets users add, modify, and remove devices and device groups.
Port Configuration	(LogLogic LX Appliance, LogLogic MX Appliance only) Lets users add, modify, and remove the port definitions on the appliance.
Message Routing	Lets users manage the appliance message routing configuration. Users

Privilege Option	Description
Configuration	can add, modify, and remove upstream devices and routing filters. **Note:** User with “Access all devices” privilege should be given the Message Routing Configuration privilege.
Manage Alerts	Lets users add, modify, and remove alerts.
Manage Check Point Devices	Lets users add, modify, and remove Check Point devices.
Manage PIX/ASA Codes	Lets users manage the categorization of incoming messages based on the PIX/ASA severity.
System Configuration	Lets users manage system settings for the appliance. Users have full access to configure general settings, remote servers, database table retention values, appliance network settings, time settings, and archive config settings.
Firewall Settings	Lets users define access rules for TCP or UDP packets accessing the appliance.

Backup/Archive Configuration	Lets users define the backup and archive configuration for the appliance:
	• Backup configuration:
	  ◦ NFS or SCP on any appliance
	  ◦ SAN on ST SAN appliances
	• Archive configuration: Archiving to the NFS remote server is available only on LogLogic ST Appliances and LogLogic EVA.

View Management Station Status	Lets users access the **Dashboards > Management Station Status** menu if the Management Station is configured.
Manage File Transfer Rules	Lets users add, modify, and remove file transfer rules for devices.
Import/Export	Lets users import and export components such as alerts, reports, search filters, and suites

Privilege Option	Description
Manage Suites	Lets users add, modify, and remove suites on the appliance.
Manage SSL Certificate	Lets users manage SSL Certificates for the appliance. Users can manage LogLogic signed certificates, import certificates, and import private keys.
Manage Device Types	Lets user add, modify, remove, import, and export device types.
Column Manager	Lets users define which columns to hide from the Searches and Reports when the Data Privacy mode is enabled.

Associating Users with Log Sources

Use the **Devices** tab to define a user's access privileges for specific log sources configured for the appliance.

For example, to give users access to logs from certain firewalls and feeders to run reports, you can assign them using the **Devices** tab. You can also select the **Access all devices in the appliance** check box to grant access to all log sources associated with the appliance.

Note: To associate users with log sources, you must have log sources added to the appliance from the **Management > Devices** tab. For details, see Device Management.

Procedure

1. From the **Device Type** list, select a specific log source to associate with the user.
2. Do one of the following:
   - Select the **Access all devices in the appliance** check box to allow access to all log sources on this appliance.
   - Highlight a log source in the **Available Devices** list box and click the appropriate arrow.

   After you add a log source, it displays in the **Selected Devices** list box.

Note: All groups defined in the **Management > Devices > Device Groups** tab appear in this list.

3. To disassociate a log source from the user, highlight the log source in the **Selected Devices** list box and click the appropriate arrow.
4. If you are finished entering user information on all tabs, click **Add**.

Devices list box Add/Remove arrows

Arrow	Description
	Adds all the available log sources associated with the Device Type.
	Adds the selected available log sources (The system associates the log source you select with the Device Type).
	Remove (The system no longer associates this log source with the Device Type).
	Deletes all available log sources associated with the Device Type.

Warning: To exclude a log source from a user, you must exclude the All *device-type* group as well. For example, to exclude a user from accessing the server “Linux 1", you must exclude both “Linux 1" and “All Linux”.

Removing a User

Use the **Remove User** tab to delete a user from the appliance database.

Procedure

1. Select a check box for the User ID from the **Management > Users** tab.
2. Click **Remove**.

3. The **Remove User** tab displays.
4. Click **Confirm Remove** to delete the user.

Adding or Modifying Users on a Managed Appliance

Adding or modifying users on managed appliances involves the same tabs as on a single appliance, plus the **Appliances** tab.

The **Appliances** tab automatically creates or updates the user settings and privileges from the current appliance on the managed appliances selected in this tab.

You can also add a user to a managed appliance by:

- Adding a user through direct access to the appliance
- Replicating the user information to other managed appliances from the **Administration > User Replication** tab (see Replicating Users on a Managed Appliance)

The **Appliances** tab displays only on Management Stations. If you access another appliance through the Management Station, you cannot see the **Appliances** tab.

> **Note:** When removing a user from the management station, LogLogic LMI automatically removes the user from all managed appliances to which the user had been granted access, without the need for an administrator to initiate a user replication action.

Procedure

1. In the **Management > Users > Users** tab, click **Add New**.

   To update an existing user, click a user ID in the User ID column on the **Management > Users > Users** tab.
2. Fill out the **General, Privileges**, and **Devices** tabs as described in Adding or Modifying a User.
3. Select the **Appliances** tab.

4. Highlight an appliance IP address from the **Available Appliances** section and click the single arrow button to move it to the **Selected Appliances** section.

   To move all available appliances to the **Selected Appliances** section, click the double arrow button.

5. Click **Add** to add the new user to the selected appliance(s).

   When appliances are selected for a user, any change to that user on the Management Station is automatically replicated to the selected appliances.

Replicating Users on a Managed Appliance

User replication executes a system-wide replication of its user database from the current appliance (displayed in the bottom right corner of the screen) to the selected managed appliances.

Replicated data includes the rights and privileges of all users and roles on the current appliance.

The list box lists each managed appliance.

Procedure

1. Select the appliances in the Management Station cluster in the **Administration > User Replication** tab list box.

   > **Note:** To select multiple appliances, hold the Ctrl-key for Windows or the Command key for Apple Macintosh.

2. Click **Confirm** to update the selected appliances.

Result

If you created a user with the same user name on a remote appliance as well as a Management Station appliance, and the remote appliance is in the Management Station cluster when you replicate users, the shared user is not replicated. To replicate the shared user, you must delete the user on the remote appliance before replicating users. If you do not first remove the user on the remote appliance, the shared user is not replicated and you might not be able to successfully run searches and reports.

If you do not see the target appliance in the cluster, verify that it is in the list from the **Management Station > Configuration** tab. If needed, you can add an appliance to the cluster from that tab.

User Roles

Creating a role on the appliance lets you use an Active Directory server to remotely control authentication and access permissions on the LogLogic appliance.

The appliance role corresponds to a group on the Active Directory server.

After an appliance is created, it uses the role's corresponding AD groups for login authentication, permission settings, and access to specific log sources just as if the user was configured directly on the appliance. Users can be configured with multiple roles.

When defining a role on the appliance, you map it directly to an existing Active Directory group. If you want to create an appliance role for which an AD group does not exist, you must first create the group on the AD server.

The LogLogic appliance permission settings you assign to that role are automatically applied to all users in that AD group who log in to the appliance. You assign users to roles by including the users in the corresponding AD groups on the AD server.

The **Directory Roles** tab lists all the Active Directory roles defined for use on the appliance. You can access the **Directory Roles** tab from **Management > Users**.

Directory Roles is available in the GUI only if an Active Directory remote authentication server is enabled in the **System Settings > Remote Servers** tab. If you create roles then disable the AD server in the GUI, the roles are retained for whenever the AD server is re-enabled.

- To add a new role to the appliance, click **Add New**. The **General** tab appears.
- To modify an existing role on the appliance, click the role’s **Role Name**.
- To remove a role from the appliance, check the role’s check box and then click **Remove**.

If you configure Active Directory use on a Management Station, the managed appliances also display the **Directory Roles** tab.

Adding or Modifying a Role

To add or modify a role, use the **Management > Users > Directory Roles** tab.

- To add a new role, click **Add New**  .
- To modify an existing role, select the **Role Name** from the role list.

When adding a role, the **Role** tabs are dependent on each other. You must complete the **General** tab before accessing the **Privileges and Devices** tabs.

When modifying a role, you can access any tab because they are all filled from when the role was added.

When filling out the tabs, you can switch between the tabs without clicking **Add** or **Update**. Your entries are retained until you click **Add**, **Update**, or **Cancel** on one of the tabs. For example, when adding a role:

Procedure

1. Complete the Configuring General Settings for a Role tab.
2. Complete the Role Privileges tab.
3. Click **Add** or **Save** on any of these tabs.

Result

The **Edit User Roles** section is closed, and the **Directory Roles** tab displays the new role in the list.

Configuring General Settings for a Role

Use the **General** tab to define or update the general settings for a role.

Procedure

1. In the **Role Name** text field, enter a role name.
2. In the **Description** text field, enter a description for the role.
3. In the **Directory Group DN** text field, enter the distinguished name of an existing group on the Active Directory server.

For example:

```
cn=group-name,cn=Users,dc=shared,dc=directory-server-
domain,dc=local
```

If you are creating roles on the appliance before connecting to the AD server, you can leave this text field empty. Otherwise, this field is required for enabling any role.

4. Select the appropriate **Enable** radio button. This enables the role to control Active Directory users accessing the appliance.

What to do next

Complete the Role Privileges tab.

Role Privileges

Use the **Privileges** tab to grant various privileges to users with this role when adding or modifying them.

User privileges control access to different parts of the appliance. Navigation menu items display according to the level of privileges granted to a user through their role.

Note: Depending on your appliance selection, you might see only a subset of these options in the navigation menu.

Select the check boxes to grant specific privileges under each category, or clear to remove privileges. The privileges you can grant depend on the privileges you have on the appliance. For more information, including a complete list of privileges, see User Privileges.

Associating Devices

The **Device** tab lets you associate users with specific log sources on the appliance.

To let this user access all log sources on this appliance, including log sources added in the future, select the **Access all devices in the appliance** check box.

Procedure

1. From the Device Type list, select the type of available log sources (or all log sources) to be displayed.
2. Move log sources between the **Available Devices** and **Selected Devices** lists by using the arrow buttons:
    a. To select a single log source, highlight its name under **Available Devices** and click  .
    b. To select all listed available devices, click  .
    c. To remove a single log source, highlight its name under **Selected Devices** and click  .
    d. To remove all listed selected devices, click  .
3. After you finish configuring all **User** tabs, click **Save** or **Update**.

    Note: Your settings are retained as you move between these tabs while adding or modifying a user.

    All groups you define in the **Manage Devices > Add Device Group** tab appear in this list. If you selected **Show Only Device Groups** from the **System Settings > General** tab, **Available Devices** lists only device groups.
4. To cancel adding or modifying the user and return to the **Users** tab, click **Cancel**. Changes made on any User tabs for this user are not retained.

Configuring the Appliance Settings

The **Appliance** tab lets you automatically create or update a user on different remote managed appliances.
This tab is applicable only to Management Stations or Management Appliances. It is available for LogLogic appliance-based users, and not for Active Directory users.

When appliances are selected for a user, any change to that user is automatically replicated to the selected appliances.

Note: User privilege on the remote appliance is required for the user to remotely control the appliance.

- When adding a user on a remote appliance, the **User** tabs are dependent on each

other. You must complete the **General**, **Privileges**, and **Device** tabs before you can access the **Appliances** tab.

- When modifying a user, you can access any tab in any sequence, because the data in the tabs already exists and is fetched from when the user was added.

You can also add a user to a managed appliance by:

- Adding a user through direct access to the appliance
- Replicating the user information to other managed appliances from the **Administration > User Replication** tab.

To specify managed appliances that this user can access and to add user privilege for the remote appliance, perform the steps in the following procedure.

Procedure

1. Move log sources between the **Available Appliances** and **Selected Appliances** lists by using the arrow buttons:
   - To select a single appliance, highlight its name under Available Appliances and click  .
   - To select all listed Available Appliances (or Products), click  .
   - To remove a single appliance, highlight its name under **Selected Appliances** and click  .
   - To remove all listed Selected Appliances, click  .
2. Click **Add** or **Update** after you have completed all **Users** tabs.

   Your settings are retained as you move between these tabs while adding or modifying a user.

   To cancel adding or modifying the user and return to the **Users** tab, click **Cancel**. Changes made on any **Users** tabs for this user are not retained.

Removing a Role

Use the **Remove Role** tab to delete a role from the Appliance.

This does not change any group or user settings on the Active Directory server itself.

Procedure

1. Select the check box for the Role Name from the **Management > Users > Directory Roles** tab.
2. Click **Remove**.

   The **Remove Role** tab is displayed.
3. Click **Confirm Remove** to delete the role.

Network Settings

Use the **Network Settings** menu to configure the appliance on your network.

From the Network Settings screen, you can add or update the IP addresses that are used to access your appliance.

Configuring your Network Settings

From the **Administration > Network Settings** page, you can configure the network settings for your appliance.

Note: In an AWS environment, changing the value of the fields on the **Administration > Network Settings** page might cause the appliance to become unresponsive. Do not change the value in these fields.

Procedure

1. Go to **Administration > Network Settings**.
2. On the **Configuration** tab, enter the following information:

Field	Description
Hostname	Enter the hostname for your machine. **Note:** Do not change the hostname for LogLogic EVA running in an AWS environment.
Default Gateway	Select a default gateway for your machine. Your gateway can be either IPv4 or IPv6. IPv4 uses 32 binary bits to create a single unique address on the network. An IPv4 address is expressed

Field	Description
	by four numbers separated by dots. Each number is the decimal (base-10) representation for an eight-digit binary (base-2) number, for example: 10.114.75.2 IPv6 uses 128 binary bits to create a single unique address on the network. An IPv6 address is expressed by eight groups of hexadecimal (base-16) numbers separated by colons, as in `2001:cdba:0000:0000:0000:0000:3257:9652` Groups of numbers that contain all zeros are often omitted to save space, leaving a colon separator to mark the gap (as in `2001:cdba::3257:9652`).
Primary DNS server, Secondary, Tertiary,	(Optional) If you have DNS servers, type their IP addresses in the **DNS Server** fields. Enter the primary, secondary, and tertiary IP addresses in the respective fields.
<x> v4 static routes <y> v6 static routes	Click the number of static routes link or go to the **Static Routes** tab to view a list of static routes.

3. In **Interface Pairing**, choose the interface associated with the appliance. Depending on the type of appliance you have and the selection you make in the **Default Gateway** fields, the available options are displayed in the **Interfaces** section.

   For physical connections, see the *TIBCO LogLogic® Log Management Intelligence Hardware Installation Guide.*

Note:

- The eth0 port must be connected to the network because that is the only port configured on the appliance by default at startup. The other ports do not work unless the eth0 port is connected to the network.
- Do not change any network settings after two appliances are paired in an HA setup.
- When multiple interfaces are configured, it is recommend to use different networks for each interface.

The following interface options are available:

Check box	Action	Description
pair eth0 and eth1 into bond0 (1U/2UR1 Appliances)	By default, the check box is selected, configuring the appliance to use bond0. Clear the check box to allow access to both eth0 and eth1.	Clearing the check box lets the two physical interfaces, eth0 and eth1, become one logical interface. For example, to connect the appliance to two separate networks, you must clear the check box and configure eth0 and/or eth1, remembering to set your Gateway Interface as appropriate.
pair eth1 and eth2 into bond0 (2U Appliances)	By default, the check box is selected, configuring the appliance to use bond0. Clear this check box to allow access to both eth1 and eth2.	This option allows up to five NICs.
pair eth2 and eth3 into bond0 (For TIBCO	By default, the check box is selected, configuring the appliance to use bond0. Clear this check box to allow access	Clearing the check box lets the two physical interfaces, eth2 and eth3, become one logical interface. For example, to connect the appliance

Check box	Action	Description
LogLogic® ST2025-SANR1 Appliance models)	to both eth2 and eth3.	to two separate networks, you must clear the check box and configure eth2 and/or eth3, remembering to set your Gateway Interface as appropriate.
pair eth4 and eth5 into bond1 (For TIBCO LogLogic® ST2025-SANR1 Appliance models)	By default, the check box is selected, configuring the appliance to use bond1. Clear this check box to allow access to both eth4 and eth5.	This option allows up to five NICs.
pair eth4 and eth5 into bond2	Select this check box to allow access to both eth4 and eth5 as bond2. To connect the appliance to two separate networks, clear the check box and configure eth4 and/or eth5, and remember to set your Gateway Interface, IP address, and Netmask as appropriate.	Selecting the check box makes the two physical interfaces, eth4 and eth5, become one logical interface.

4. Enter the following information in the **Interfaces** section:
    - For IPv4:
        - **Network IP Address:** IP address for the interface. IP address MUST be set.
        - **Netmask:** Netmask for the IP address. Netmask MUST be set.
        - **Speed:** Select the connection speed from the list.
    - For IPv6:
        - **Network IP Address:** IP address for the interface. IP address MUST be

set.

- **Prefix:** Prefix for the IP address. Prefix MUST be set.
- **Speed:** Select the connection speed from the list.

> **Note:** The speed field is not available in LogLogic EVA.

5. Click **Save** to save your changes.

What to do next

If your changes require a reboot, click **Reboot Now** or **Reboot Now** as applicable for the changes to take effect.

Configuring for a Multi-homed Network

If the appliance is configured to operate in a multi-homed network setting, the Linux kernel parameter of `rp_filter` must be changed from 1 to 0 for all the NIC interfaces.

Static rules do not work if `rp_filter` is set to 1 or 2. There are multiple `rp_filter` settings and all of them should be modified. They are located at `/proc/sys/net/ipv4/conf/*/rp_filter`.

> **Note:** This operation could make the appliance vulnerable to DDoS attacks. See Red Hat documentation at https://access.redhat.com/solutions/53031.

Procedure

1. Find all network interfaces on the machine using the following command:

```
$> ls /proc/sys/net/ipv4/conf/ | grep -v all | grep -v lo
```

2. With the list of network interfaces, edit the conf file `/etc/sysctl.conf` by appending the configuration into it, one line per interface:

```
net.ipv4.conf.<interface_name>.rp_filter = 0
```

   For example:

```
net.ipv4.conf.default.rp_filter = 0
net.ipv4.conf.eth0.rp_filter = 0
net.ipv4.conf.eth1.rp_filter = 0
```

Reboot the machine for the settings to take effect.

Note: If you already have static routes configured, you must configure the multi-home network and additionally change the values in `/proc/sys/net/ipv4/conf/*/rp_filter`. This is because though both have the same effect, changing the settings in `/proc/sys/net/ipv4/conf/*/rp_filter`, does not require rebooting the system.

Adding or Removing Static Routes

You can view the existing static routes, add new routes and delete specified static routes using the **Static Routes** tab. You can add one or multiple static IPv4 or IPv6 routes at a time.

Procedure

1. Click the appropriate button:
   - To add one static route, click the **Add** button. A new empty row is added to enter information.
   - To add multiple static routes, click the **Batch Add** button. The **Add Multiple Static Routes** window appears.
2. Enter the following information for the IPv4 or IPv6 routes, as applicable. For multiple entries use space to separate each entry. You can copy and paste the multiple entries.

IPv4	IPv6
• **Network IP** – IP address for the destination network	• **Network IP** – IP address for the destination network

IPv4	IPv6
• **Netmask** – Netmask for the destination network • **Gateway** – Gateway for the destination network	• **Prefix** – Prefix for the destination network • **Gateway** – Gateway for the destination network

3. Click **Save** to add the new routes in the Static Routes list.
4. To remove a static route:
   a. Click the ⊗ button from the **Action** column for the static route you want to remove.
   b. Click **Yes** in the confirmation window to delete the static route.

Network Access Control

From the **Administration > Firewall Settings** page, you can restrict network access based on source IP address and destination port, similar to access lists used by routers or firewalls.

For example, you can allow TCP port 443 access from a specific host or subnet.

Firewall Settings

You can define rules for Source IP address or subnet, and destination TCP or UDP ports.

Your input rules can allow and deny access based on the following criteria of the machine requesting authentication or access:

- IP address:
  - All
  - Single IP / CIDR
- Protocol:
  - TCP
  - UDP
- Port:

  For a list of ports see Port Assignments.

  You can change the Syslog UDP or TCP port number from 514, or specify additional port numbers. See General Settings.

  Note: If non-default Syslog UDP or TCP Port is specified (using the **Administration > System Settings > General** tab) and the **Enable IP Firewall** check box is selected, then you must add an additional Input Rule for each non-default Syslog UDP or TCP port to allow this port to receive syslog messages. For instructions, see Adding an Input Rule.

- Action:
  - Accept
  - Deny

Adding an Input Rule

Use the **Administration > Firewall Settings** tab to add input rules and define your firewall settings.You can also use the system firewall CLI command to add or delete a firewall rule, or to turn the firewall on or off. For more information about the CLI command, see .

> ⚠ **Warning:** Turning off **Firewall Settings** means any IP address can access the services on the appliance.

New input rules are added to the bottom of the rule list. Input rules are processed in descending order. Therefore, if you add a rule that might be superseded by one of the higher rules in the list, you must first delete the higher rule for your new rule to be effective.

For example, a default input rule accepts all IP addresses with UDP port 514. If you add a rule denying access to a particular IP address (for example 180.22.21.5) using UDP and port 514, that rule is superseded by the higher default rule that accepts all input using UDP and port 514. To make your added rule effective, you must:

1. Add a new rule denying 180.22.21.5 using UDP on port 514.
2. Delete the default rule that accepts all IP addresses using UDP on port 514.
3. To still accept all other IP addresses using UDP and port 514, add another new rule accepting all IP addresses using UDP on port 514.

   Because this new “accept all” rule appears after the “deny 180.22.21.5” rule, both rules are executed. The appliance accepts input from all IP addresses using UDP on port 514 except 180.22.21.5.

Procedure

1. Select **Administration > Firewall Settings**.
2. Select **Enable IP Firewall** to activate the **Input Rule** box.

3. In the **Input Rule** section, define the rules:
   a. Define an **IP Address**.
   b. To accept all IP addresses with the **Protocol** and **Port** you define, select **All**.
   c. Specify an IP address, or subnet mask, or both for the rule.

      The IP address indicates which hosts are allowed to communicate with the appliance. The format for this field is `IP-address/subnet-mask`. For example:

      - For a 24-bit subnet mask: 192.168.2.0/24
      - For a 16-bit subnet mask: 192.168.0.0/16
      - For an 8-bit subnet mask: 192.0.0.0/8
      - For a 72 bit IPv6 subnet: fd0f:c4al:e456:0000:5200::/72
4. Select the **Protocol** (TCP or UDP) to associate with the port you specify.
5. Select a **Port** from the list of ports active on the appliance.

   For a list of ports refer to Port Assignments. To add a custom port by using the CLI, run the firewall command.

Protocol	Port number
HTTP	80
HTTP Collector	4433
HTTPS Remote Control	4443
HTTPS	443
Loglogic Tunnel	11965 (this port is deprecated and may not be available)
MCAGENT	2508
MCAGENT	2098

Protocol	Port number
MCAGENT	2099
NTP	123
SSH	22
SNMP	161
SNMP-Trap	162
SYSLOG	514
Loglogic Receiver	5514
NetFlow	2055
NetFlow	9555
NetFlow	9995
ULDP	5515
ULDP	5516

6. Select an **Action** to indicate whether your appliance accepts or denies a packet that meets the rule requirement. The default is **Deny**.
7. Click **Add** to add the rule to the **Input Rule Table**.
8. Click **Apply** to activate the rules.

   The **Input Rule Table**, beneath the **Input Rule** section, lists the currently active rules.

Column	Description
IP Address	IP address or subnet you typed for the Input Rule.
Port	Port you selected for the Input Rule.
Protocol	Protocol you selected for the Input Rule.
Action	Action to take if the packet meets your rule requirements.
☒	Deletes the access rule from the list.

Deleting an Input Rule

Use the **Administration > Firewall Settings** tab to delete input rules and define your Firewall Settings.You can also use the system firewall CLI command to add or delete a firewall rule, or to turn the firewall on or off. For more information about the CLI command, see .

Procedure

1. Select **Enable IP Firewall**.
2. Click the ☒ icon for the rule to delete.
3. Click **OK** to exit the confirmation window.
4. Click **Apply** to accept the changes.

> **Note:** The appliance treats port 443 (HTTPS) differently; you cannot delete the last rule for port 443. This prevents you from losing browser access to your appliance; at the same time letting you restrict access to port 443.
>
> For a list of ports, see Port Assignments.

SSL Certificate Management

You can generate and activate SSL certificates for an appliance in any of the following ways:

Certificate type	Description	For more information, see...
LogLogic signed certificate	The appliance creates and activates its own certificate and private key.	Creating and Activating LogLogic Signed Certificates
Certificate Authority (CA) certificate	A CA is used to provide the root and signed (reply) certificates; the private key of the appliance is used.	Signing the Certificate Using a CA
Other certificate	The appliance uses certificates and a private key that you generate and provide from a source other than the first two options.	Importing Certificates and Keys

Important: When changing the system time of your appliance, ensure that the validity period of the certificates imported into the appliance matches the new system time.

Creating and Activating LogLogic Signed Certificates

Use the **LogLogic Signed Certificate** tab to create and activate a LogLogic signed certificate for your appliance IP.

Warning: Completing this procedure before changing the Certificate Signing Request changes the private key. You must maintain consistency between your keys.

Procedure

1. Click **Administration > SSL Certificate**.

   The **LogLogic Signed Certificate** tab opens.

2. Enter the following information:

Field	Description
Common Name	The DNS name (recommended) or IP address of the appliance
Organizational Unit	The name of your department
Organization	The name of your company
City	The name of your city
State Name	The name of your state
Country	Select the abbreviation of your country
Validity Period	The number of days for which the certificate remains valid

3. Click **Activate**.

Result

The appliance creates its own certificate and private key using the LogLogic root CA to sign the certificate, and then activates the new certificate and private key. Tomcat automatically restarts to apply the changes.

Signing the Certificate Using a CA

Use the **Certificate Signing** tab to generate a signing request for the current certificate used by the appliance.
After the CA signs the request, use this tab to import the CA root certificate and the signed certificate.

If your organization's PKI policies dictate the use of one or more intermediate certificate authorities, an additional step is required to ensure that LogLogic LMI can properly verify the entire certificate trust chain.

The simplest configuration involves the LogLogic LMI certificate and a root certificate. In such a scenario, the root certificate is not just the root CA, but also the signing CA. This distinction is important when determining the sequence in which to paste the certificates:

- When using multiple CA certificates, consider the Root Certificate text box as a field that is used for the Issuing Certificate, and then all other CA certificates are added in the same box (see the step for multiple certificates).
- In situations involving a single CA certificate that serves as both the root and issuing CA, the intermediate CA is the signing CA and the root is still the root CA (see the step for single certificates). There can be any number of intermediate CAs. The exact situation depends on your organization's policies.

> **Note:** Some certificate authority software do not provide all the necessary files when returning your signed certificate, and some do not require any intermediate certificates. However, LogLogic LMI does require them if they are used in the certificate signing process. Therefore, if you do not know how many CA certificates to expect to be given when your signed host certificate is returned to you, then verify with your organization's PKI administrators whether your organization uses multiple CA certificates. There is no limit to the quantity of intermediate certificates that LogLogic LMI can use.

Before you begin

You must first complete the **LogLogic Signed Certificate**. See Creating and Activating LogLogic Signed Certificates. The **Certificate Signing** tab uses information from the **LogLogic Signed Certificate** tab, as well as the appliance private key that it activates.

Procedure

1. Click **Administration > SSL Certificate > Certificate Signing**.
2. Click **Generate**.

   A certificate signing request (CSR) from LogLogic is generated based on the information in the **LogLogic Signed Certificate** tab.

   > **Note:** By default the GUI generates CSRs with keys that are 2048 bit in size.

3. On the certreq.csr dialog box, open the `.csr` file.
4. Copy the text from the `.csr` file to the website of your trusted CA. The CA returns a root and reply certificate.

   > **Note:** Generating a root and reply certificate from the CA might take time.

5. Perform any **one** of the following:
   - If your organization uses a single CA certificate that serves as both the root and issuing CA, paste the text from the root certificate of the CA in the **Root Certificate** text box. Then proceed to step 6.
   - If your organization uses multiple CA certificates, perform the following steps:
     a. Paste in the certificate that was given to you by your PKI system that was used to sign the LogLogic LMI appliance's certificate.
     b. Paste in the contents of each CA certificate in reverse order of the certificate chain, so that the top root CA certificate is pasted as the last certificate.

        > **Warning:** If the order is random, backwards, or incorrect in some way, the certificates do not work correctly after the LogLogic LMI web server is restarted.

        For example, if you have the following certificates with Common Names (CN= values) of `toprootCA`, `interCA`, `signingCA`, and `myLMI`, and with the certificates in that sequence within the trust chain, then they must be pasted into the Root Certificate text box using the following bottom up order:

i. signingCA.cer

ii. interCA.cer

iii. toprootCA.cer

The `myLMI` certificate is required during step 6.

If you are unsure which sequence the CA certificates exist in the chain, contact your organization's PKI administrator ,or follow the certificate chain by looking at the **Issued By** and **Issued To** fields for each certificate. The first certificate pasted into the **Root Certificate** text box must be the one that signed the LogLogic LMI appliance's certificate. The last certificate that is pasted in the **Root Certificate** text box must be the top-level root CA certificate.

> **Warning:**
> - Do not separate the certificate contents with blank lines.
> - Make sure that the contents of each certificate start on a new line.

6. Paste the text from the root certificate of the CA in the **Root Certificate** text box.
7. Paste the text from the CA-generated certificate in the **Reply Certificate** text box.
8. Click **Import** to import the certificate.

Result

The certificates are imported and Tomcat automatically restarts to apply the changes.

Importing Certificates and Keys

Use the **Certificate Import** tab to activate your own certificates and a key that you generate from another source beside the appliance itself.

> **Warning:** Completing this procedure before changing your Certificate Signing Request, changes your private key. You must maintain consistency between your keys.

Procedure

1. Click **Administration > SSL Certificate > Certificate Import**.

2. Paste the root certificate into the **Root Certificate** box.
3. Paste the generated private key into the **Private Key** box.
4. Paste the certificate generated for this Appliance IP address into the **Certificate** box.
5. Click **Activate**.

Result
Tomcat automatically restarts to apply the changes.

Multiple Intermediate Certificates

If there is only one subordinate certificate, you must paste the contents of the subordinate (intermediate) certificates into the **Root Certificate** field right after pasting the certificate for the host appliance itself. However, if there are multiple subordinate certificates, then they have to be pasted in a certain order. The order is based on the order they exist in the certificate chain, starting with the certificate that signed the host certificate.

Importing a Trusted Certificate

You must create a trusted relationship between the Management Station and its Remote Appliances.
From the **Administration > SSL Certificate > Trusted Certificate** tab, you must import the current certificate of the Management Station into the Remote Appliance, and vice versa.

Important: You must perform this procedure in the following scenarios:

- For each Remote Appliance in the setup
- Every time the certificate of a Remote Appliance changes

Procedure

1. On the Management Station, copy the text from the **Current Certificate** field.
2. On the Remote Appliance, paste the current certificate of the Management Station

into the **Import Trusted Certificate** field and click **Import**.

The new trusted certificate is imported from the Management Station and installed on the Remote Appliance. Though the Remote Appliance web server automatically restarts, log collection on the Remote Appliance is not affected.

3. Because the Management Station uses two-way authentication, you must copy the text from the Management station and paste the current certificate of the Remote Appliance into the **Import Trusted Certificate** field of the Management Station.

   The new trusted certificate is imported from the Remote Appliance and installed on the Management Station. Though the Management Station web server automatically restarts, log collection on the Management Station is not affected.

Result

Two-way authentication is established between the Management Station and its Remote Appliances.

Security Settings

Security settings for internal and external communication between different nodes of LogLogic LMI are stored in the `/loglogic/conf/llsecurity.conf` file.

For enhanced security, LogLogic LMI provides the following features:

- Two-factor authentication
- Ability to prevent cross-site request forgery

Two-factor authentication

By default, two-factor authentication is disabled in LogLogic LMI. However, you can enable this feature by configuring parameters in the `/loglogic/conf/llsecurity.conf` file.

Before you begin

Before enabling two-factor authentication, ensure that:

- SSL certificates have been issued to all users.
- The Distinguished Name (DN) in the user's SSL certificate matches the user name in LogLogic LMI.
- Users have imported the SSL certificates in the browser to access LogLogic LMI.

If a user removes or deletes a certificate from the browser, the browser cache must be cleared and the browser restarted.

Procedure

1. Connect to the appliance by using SSH and edit the `/loglogic/conf/llsecurity.conf` file.
2. Configure the following parameters in the file:

Parameter	Default Value	Description
CLIENT_AUTH_ENABLED	false (disabled)	Enables or disables client authentication.
CA_CERT_FILE_PATH	(empty)	Path to the extra client CA certificate file. For example, `/loglogic/conf/<CA_CERT_FILE>` Applicable only if two-factor authentication is enabled.

3. After making any changes to these parameters, restart the LogLogic LMI application by running the following commands:
   a. `$ mtask stop`
   b. `$ mtask start`

Cross-site Request Forgery

Cross-site Request Forgery (CSRF) is a method of malicious access to a website, in which unauthorized commands are transmitted via a user that is trusted by the website. LogLogic LMI includes the ability to prevent CSRF attacks.

Configuring TLS Syslog

TCP connections between log sources and LogLogic LMI can be secured by TLS.

TCP syslog feeders, rsyslog feeders, and ULDP clients can connect to LogLogic LMI via TLS using a certificate.

Note: TLS TCP connections are supported for log data collection. However, encrypted log data forwarding to LogLogic LMI appliances uses SSH encryption.

By default, a TCP collector supports two ports 514 and 6514, and up to 14 more custom ports. If the firewall is enabled on the LogLogic LMI appliance, you must run `system firewall` to add these ports to the firewall rule. The rules take effect automatically.

Note: The iptables rules are created automatically.

Every TCP collector's port can be used for collecting logs - either unencrypted or by using TLS. The TCP collector automatically detects secure TLS TCP connections.

The port numbers and other information can be configured using a configuration file `/loglogic/conf/tcpcoll.conf`.

Procedure

1. Create the file `/loglogic/conf/tcpcoll.conf`, if it does not exist.

   Note:
   - Ensure that the configuration key file has `0644` permission.
   - In an HA configuration, you must save this file on both the active and the standby nodes.
   - It is best practice to place certificate files in a sub-folder under `/loglogic/conf/` so that the certificate files are backed up automatically.
   - The TCP collector runs with lowered privileges. The following files mentioned as parameters in the configuration file must be readable by the `logapp` user:
     - `TLSCert=<file_name>`

       For example: `TLSCert=/loglogic/conf/certs/cert1`
     - `TLSKey=<file_name>`

       For example: `TLSKey=/loglogic/conf/certs/key1`
     - `TLSDefaultClientCA=<file_name>`

       The default client file is `TLSDefaultClientCA=/loglogic/conf/certs/client179.ca`

2. Add the following information in the file:

   a. Specify the TLS version in the `TLSVerFlags` parameter:

The `TLSVerFlags` parameter is a 4-bit number, in which each bit represents one TLS version. Setting the corresponding bit to 1 indicates that the version is enabled. By default, TLSv1.2 is enabled, and so the default value of `TLSVerFlags` is 4, that is, 0100 (hex). The default value and some examples are explained in the following table.

	Bit 3	Bit 2	Bit 1	Bit 0
TLS version represented by the bit:	TLSv1.3	TLSv1.2	TLSv1.1	TLSv1.0
TLSVerFlags=12 TLS versions 1.3 and 1.2 are supported	1	1	0	0
TLSVerFlags=4 Only TLS versions 1.2 is supported	0	1	0	0

b. By default, client certificate validation is disabled. To enable it, add this line:

```
TLSClientVerify=1
```

0 indicates disabled; any other number indicates enabled

c. An SSL certificate is generated whenever the IP of the appliance changes. The SSL certificate is for receiving logs from Blue Coat Proxy appliances over HTTPS. This certificate is the default to be used by LogLogic LMI (TLS server). To configure the certificate and the key path:

```
TLSCert=/loglogic/conf/certs/cert1

TLSKey=/loglogic/conf/certs/key1
```

> **Note:** The external keys and certificates cannot be synchronized between HA nodes, and must be configured manually on the active and standby nodes.

d. For client certificate validation, specify the client's CA file.

   For example:

```
TLSDefaultClientCA=/loglogic/conf/certs/client179.ca
```

> **Note:**
> - You must manually put the CA file of the client certificate in this directory, and it must have the same name as specified in this configuration file.
> - Only one client CA is supported. If you specify multiple lines for the `TLSDefaultClientCA` parameter, only the last one is considered. If you use a certificate chain to sign the client certificate, the chain must be stored in the same file that is mentioned in this parameter.

e. To have LogLogic LMI listen on any additional ports for receiving TCP syslog data, you must explicitly specify the port numbers using the `ListenOnPort` parameter:

   In the following example, the port number is 4321:

```
ListenOnPort=4321
```

   If you want to add more ports, specify each port on a separate line. For example:

```
ListenOnPort=4322
ListenOnPort=4323
ListenOnPort=4324
ListenOnPort=4325
```

> **Note:** This step is required for both TLS TCP syslog and octet framing features (except for the default ports 514 and 6514).

f. (Optional) To enable octet-counting framing on the additional ports, you must set the `FrameOnPort` parameter for those ports. To specify multiple ports, each port must be listed on a new line:

```
FrameOnPort=2513
FrameOnPort=2514
FrameOnPort=2515
```

The port listed for octet-counting framing must be a listening port, which means the `ListenOnPort` parameter for the port must appear before the `FrameOnPort` parameter. For example, if you configure the port 2345 to receive octet-counting framing messages, you must configure the parameters as:

```
ListenOnPort=2345
FrameOnPort=2345
```

Once a port is set to receive octet-counting framing messages, the port rejects messages that are not of the octet-counting framing type.

> **Note:** Because the ports 514 and 6514 are open by default, it is sufficient to specify the `FrameOnPort` parameter for these ports and omit the `ListenOnPort` parameter.

3. (Optional) The password associated with TLSKey can be set by using the CLI command:

```
> set tls syslog key password
```

The password in encrypted when stored. It is stored in MySQL in the `Settings2` table of the `logappconfig` database, with a setting called `GlobalTlsKeyPwd`.

Follow the instructions to type a password and confirm the password. To erase the password, run the same command, but press the Enter key when prompted for the password.

4. Save the file and restart the `engine_tcpcollector` using the following command:

```
$ mtask -s engine_tcpcollector restart
```

Failover

The LogLogic failover feature allows high availability (HA) of a LogLogic appliance by providing:

- Real-time replication of logs
- A fast and reliable failover mechanism

Failover Architecture

You can configure LogLogic appliances in a one-to-one HA pair failover configuration.

There is a single active appliance with a single standby appliance. The standby appliance continually stays in sync with the active appliance and automatically takes over for the active appliance if a problem arises.

Each appliance has its own private IP address. The HA pair itself has a public IP address, which serves as the single entry point for the HA pair. All external systems point to the public IP address so data goes to the active appliance in the HA pair regardless of which appliance it is at a given time.

LogLogic failover is built on three software layers that together ensure failover membership, database replication, and recurring node resynchronization.

Public and Private IP Addresses

IP addresses can be used to access appliances virtually.

LogLogic failover is based on an active/standby architecture and provides a single system image composed of two appliances of the same type. The two appliances can be treated as a virtual appliance accessed using a single public IP address.

Warning: It is strongly recommended that in an HA pair, the two appliances must have:

- the same model
- the same LogLogic LMI and LogLogic LSP versions (the hot fix versions need not match)

You must configure each of the two appliances using a private IP address. To configure the private IP address, use the **System Settings > Network** tab on the GUI or `set ip` command on the CLI.

The public IP address is the single entry point of the failover pair. All log sources for logs collected by the failover appliances must use the public IP address as their target IP.

The public IP address is specified during failover configuration. At any instant, only one appliance owns the public IP address and acts as a primary node for collecting and processing logs. The other node is the exact replica, in real time, of the primary node and acts as a standby. In the event of a primary node failure, the public IP address defers to the standby to guarantee availability and minimal loss of logs.

The following diagram is an example of a failover solution for a LogLogic LX Appliance.

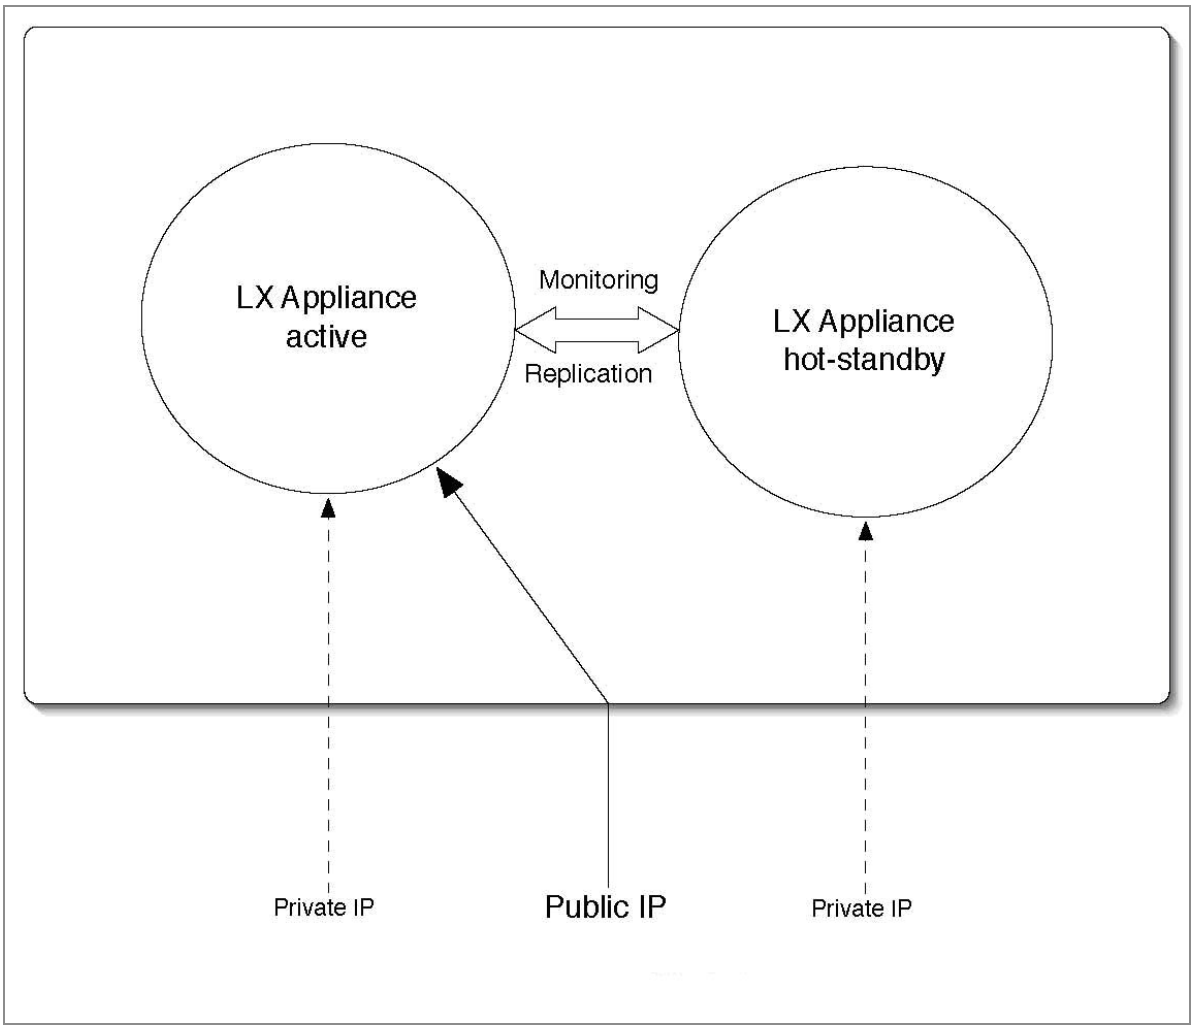

Failover and External Storage

External storage for LogLogic ST Appliances, such as the attached volumes on NFS remote servers, is not replicated as part of the failover configuration.

Be aware that in an HA pair, you must include the public IP addresses and the private IP addresses of the exported volume to the accepted IP address list on the NFS remote server.

When failover happens on a LogLogic ST Appliance HA pair, the configured NFS archive volumes are automatically mounted on the new active node. For example:

1. Appliance 1 is the active LogLogic ST Appliance, and appliance 2 is the standby

LogLogic ST Appliance, NFS remote storage is in use by appliance 1.

2. Appliance 1 fails, appliance 2 becomes the new active appliance.
3. Appliance 2 immediately starts using the NFS remote storage directly, just as appliance 1 had before.

Failover and Backup/Restore

Backup is performed on an HA pair similarly to how it is run for a single appliance. Configure the backup to use the public IP address.

Restoring from a backup to an HA pair involves disabling and re-enabling failover. For more information, see Backup and Restore in an HA Pair.

Failover Software Layers

There are three software layers in the LogLogic failover implementation:

1. Failover Membership
2. Real-Time Replication
3. Node Resynchronization

Each layer represents a building block responsible for a specific functionality.

Failover Membership

The failover membership layer monitors and detects a node failure in real time.

It provides a fast and reliable election mechanism to dynamically failover the public IP address from the node that just failed to an active partner node.

The appliances in the failover configuration are organized in a ring architecture, with two members, or partners, in the ring. Each node monitors its successor through a TCP connection. A small network packet is sent at regular intervals to the successor node.

Failure to receive consecutive heartbeat packets triggers a recovery operation. A recovery operation consists of moving the public IP address to another node. After a node failure detection, a new active node is elected and it assumes the public IP address. In a normal

situation, the total time for a failure recovery is 3 seconds. The public IP address is added to a network interface chosen by the user when configuring the HA feature. The HA feature also uses a network interface to monitor a partner node and to carry data replication traffic. The interface is also chosen by the user among those that already have an IP address assigned.

Note: The user is prompted to choose only if there is more than one choice.

In order to properly monitor that the public IP address is accessible, both interfaces need to be the same.

Real-Time Replication

The real-time replication of the data between the active and standby nodes is done using MySQL replication, as most appliance data is stored in MySQL database tables.

Each node is configured to replicate its databases as defined by the failover membership. The software layer responsible for the real time replication of the data dynamically configures MySQL according to the current state of the failover membership and continuously monitors the progress of the replication recovering any non-fatal error.

MySQL replication provides real-time latency while being very reliable. An SQL replication is usually done in less than two seconds and the protocol correctly handles temporary network disconnection. The volatile data stored in memory and the configuration files stored on disk of the current active appliance are replicated on the standby node every minute.

Node Resynchronization

An operation is periodically started in the background to check and resynchronize the archived data. The time needed to resynchronize both nodes depends on the available network bandwidth and the data size difference. It can be a time-consuming operation, from 5 minutes up to multiple days in the case of an appliance replacement.

An appliance has two types of data:

- Archived data (read-only)
- Active data (currently being modified)

Archiving is based on a threshold of disk usage. Node resynchronization is a background mechanism triggered by a failover membership event or activated periodically to resynchronize the standby with the active node, and guarantees that both appliances have exactly the same data in time.

The implementation is based on an open source utility that provides fast incremental file transfer. A wrapper software on top of the utility provides an online Check Point mechanism of database tables.

When a node that is configured to be part of an HA pair connects with its partner for the first time, an automatic data migration takes place. This operation makes one node a copy of the other. The node that keeps its original content is known as the Source node. The node that loses its original content is known as the Destination Node. The destination node is designated by the user when configuring the HA feature on each node. This operation resumes after disconnections or shutdown of any of the pair members, until completed. Only once this is completed is the pair able to provide the fail-over feature. Until then the Source node assumes the public IP and acts as the Active node. The Destination acts as the standby node but is not allowed to become Active.

After the data migration is complete, and as long as both nodes remain connected, the standby node contains the same data as the active node with at most one minute of latency (LogLogic LX Appliance) or less than 3 seconds of latency (LogLogic ST Appliance). If a node is temporarily disabled and later rejoins the cluster, it becomes the standby node and starts an operation to resynchronize with the active node. During this initial data migration that occurs when a node joins the cluster, the following data is removed:

- Data collected on the standby node before it rejoins the cluster
- Data that does not already exist on the active node

Failover Recommendations

LogLogic LMI provides several best practices for its failover feature, related to maximizing HA pair performance and identifying certain limitations of the failover offering.

Failover Performance

Failover performance metrics include failover time and maximum message rates.

- The failover time is 3 seconds in a normal situation. The failure detection time is 1 second, and the error recovery period is 2 seconds. Heartbeat monitoring is done over the TCP/IP protocol using a heartbeat packet of 64 bytes and a latency of 100ms.
- The failover configuration supports the same maximum message rate as a single LogLogic LMI appliance with the same limitations.

> ⚠ **Warning:** Make configuration changes only to the active appliance, and not to the standby appliance. To avoid mistakes, always use the cluster's Virtual IP address to interact with the configuration GUI.

Failover Limitations

There are a few limitations in the LogLogic LMI failover feature.

- The Real-Time report for Active VPN Connections (under Connectivity) is not available on the standby appliance. It uses a specific shared memory structure that cannot be replicated on the standby appliance.
- The public IP address assigned to the failover function is an alias of the main network interface of the appliance. This is required as part of the mechanism used to update the Address Resolution Protocol (ARP) tables in case failover occurs. Since some routers fail to release the cache of IP and hardware addresses stored in their ARP tables (or store the cache for as long as 10 minutes), the LogLogic appliance sends out an ARP-release packet once per minute. This causes the router to broadcast a discovery request to find the IP address and hardware address of the devices connected to it. When failover occurs, (or when we set up a High Availability (HA) pair), the router ARP tables are updated automatically.
- The virtual public IP address cannot be used for remote authentication (RADIUS, TACACS). Record the private IP addresses of both appliances in the remote server.
- While setting up an HA pair, various configuration files are automatically synced between the active and standby node. If any configuration file is updated after the setup is complete, you must:
    - Manually sync the changes to the standby node by running the command `/loglogic/bin/loadsettings` on the active node.

- Restart the corresponding engine on both nodes, in the following sequence:
    1. Stop the engine on the standby node and then on the active node.
    2. Start the engine on the active nodeand then on the standby node.

  This applies to the files defined in `/loglogic/conf/rsync_conf_files`.

- The SSL certificates for HTTPS and LDAP are not replicated from active to standby when using HA.
- You cannot access advanced features on a standby node in a high availability setup. However, you can access them from the public IP address or the IP address of the active node.

  Before configuring HA, you must disable Advanced Features on both active and standby appliances.
- After a failover, Active Queries tabs are not retained on the new active appliance.
- When logging into the Web GUI for the first time after HA failover, you might see a "Security violation" error when CSRF is enabled. If you see this error, log in again by using the Public IP in a new browser session.

  > ⚠ **Warning:** After you pair two appliances in HA, no network settings can be changed.

- In failover configuration, the administrator must modify archive mounting points on both active and standby appliances.
    1. Using the GUI, configure the archive mounting points for the SAN or NFS remote server through a public IP address.
    2. Using the GUI, on the active node (appliance A), modify the mounting points from the **Administration > Archive Configuration** page.
    3. When the system prompts you to reboot, click **OK**.

  The former standby (appliance B) becomes the new active node after the former active appliance starts the reboot process. The administrator needs to modify the mounting points from the **Administration > Archive Configuration** page for the new active node (appliance B) again, before the former active node (appliance A) finishes the rebooting process.

Failover Installation and Configuration

Installing and configuring failover includes the following:

- Hardware Installation
- Setting up the Failover (New HA Pair)
- Setting up the Software to Replace an HA Pair

Data migration is automatically performed when you set failover on an appliance. Data migration gets the standby appliance data in sync with the active appliance before setting up the failover relationship. This migration ensures that useful data on the active appliance is not accidentally overwritten during failover by data from the standby appliance.

For more information on how data migration works, see Data Migration Between Appliances. The procedures in that chapter are performed automatically during failover configuration.

Hardware Installation

You need to install hardware before you can do failover configuration.

The following diagram shows optimal high availability hardware setup for the failover configuration. The eth0 NIC of each appliance is plugged in the same switch, and the eth1 NIC on a different switch. The appliances are configured to use the fault-tolerant network interface bond0.

Failover Hardware Setup

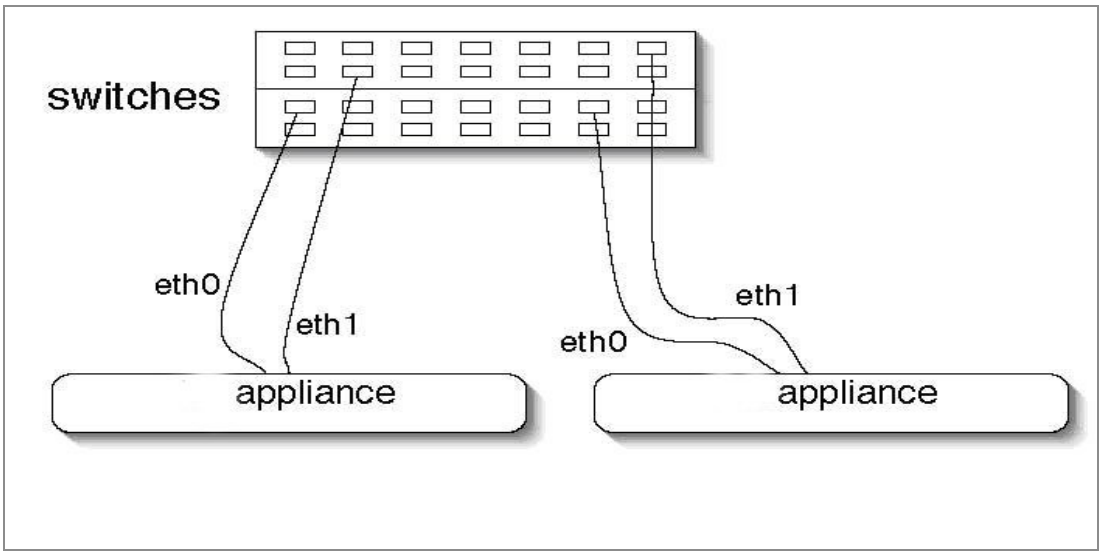

Other configurations also work. Each appliance can be connected to the network using the same or a different switch, using bond0 or not, and either or both NICs (eth0 and eth1). Network configuration that uses eth0 and eth1 independently is supported. As long as the failover membership can be established between the two nodes to guarantee the real-time replication of data, the failover works.

When choosing the IP addresses:

- Each appliance must be able to access the other appliance using its private IP address.
- Each appliance must be able to own the public IP address.
- The last digit of each IP address must be different.

Setting up the Failover (New HA Pair)

Both appliances need to be installed with compatible versions of the LogLogic software.

Before you begin

The failover membership cannot be established if there is a mismatch in software versions between the two nodes. Both the software need to have the same:

- HA protocol version

- Log storage format version
- LogLogic LSP installed version
- Database schema version

Before HA configuration, you must disable Advanced Features on both active and standby appliances.

Procedure

1. Set the NTP server on both appliances in the HA pair.
    a. Log in to the GUI web server using the private IP address of the appliance.
    b. Change the administration password, if required.

       By default, the password must be at least six characters, containing at least one non-alphabetical character, and cannot be the same as the user ID. Spaces at the beginning and end of the password are removed before being stored in the system. The password requirements might differ if password control has been enabled in LogLogic LMI.
    c. Configure the appliance time through the **Administration > System Settings > Time** tab.

       It is good practice to use the network time protocol (NTP) for failover configurations. Setting the NTP server on one appliance in the HA pair automatically sets it on both appliances and reboots them both accordingly.

       It is essential to set the NTP server time correctly during appliance installation.
2. From the CLI, through a serial console, log in to the appliance that is the active node in the HA pair.

   Note: Use a serial console instead of an SSH connection. Network configuration changes can disconnect an SSH connection during this procedure. In all circumstances, a serial connection is maintained.
3. Configure the appliance with a private IP address:

   ```
   set ip <IP address> <netmask> <gateway> [ifdev] [defaultgw]
   ```

   For example, using 10.1.1.71 as the private IP address and 10.1.1.1 as the gateway:

```
> set ip 10.1.1.71 255.255.255.0 10.1.1.1 bond0
New interface settings:
ip 10.1.1.71 255.255.255.0 10.1.1.1
CHANGES HAVE NOT BEEN SAVED!
> save
>
```

Note: If you configure more than one interface in the `set ip` command, make sure the subnets do not overlap.

4. Configure failover:
   a. Ensure that Advanced Features are disabled on both active and standby appliances.
   b. Provide the public IP address of the failover and the private IP address of the second appliance by using the `set failover configure` command.

Note: Make sure that the subnet in the `set failover configure` command is the same as the one configured for the same interface in the `set ip` command.

Also, if you specify an interface it must be an existing interface configured with a different IP address in the `set ip` command.

For example, using 10.1.1.177 as the public IP address and 10.1.1.72 as the private IP address of the other appliance:

In case of IPv6, use `::` as the value of `<broadcast>`.

```
> set failover configure

CAUTION:

You will be prompted to designate one appliance to be the
destination of the
initial data migration ("destination appliance").

All log data from the source appliance will be copied to the
destination
```

```
appliance. In cases where both appliances contain data for the same
time
periods the data from the destination appliance will be
overwritten. Data
from time periods not present on the source appliance will be
deleted from
the destination appliance.

Should the resulting volume of data exceed the capacity of the
destination appliance,
the oldest data will be deleted or archived, as needed. If the
source appliance is not
configured to use an external storage system, then the destination
appliance will be
configured not to use one either.

To minimize such storage space issues, it is best to use appliances
of the same model
and that have enough space for the entire data set.

Enter the public Ip for the HA partner pair
in the form <ip> <netmask/prefix> <broadcast> <ifdev>, or 0 to
cancel:
>> 192.168.1.248 255.255.255.0 192.168.1.255 eth0
This appliance shall be the destination of the initial data
migration [Y/N]:
>> n
Enter the Ip address of the peer appliance in the form <ip>:
>> 192.168.1.249

CHANGES HAVE NOT BEEN SAVED!
>
```

5. Save the changes to apply the new configuration:

```
> save

Writing changes to disk...
Generating new SSL certificate...
Generating RSA private key, 2048 bit long modulus
.........................+++
........+++
e is 65537 (0x10001)
```

```
Generating RSA private key, 2048 bit long modulus
.........................+++
.....................+++
e is 65537 (0x10001)
Using configuration from /loglogic/conf/certs/llssl_.cnf
Check that the request matches the signature
Signature ok
The Subject's Distinguished Name is as follows
countryName           :PRINTABLE:'US'
stateOrProvinceName   :PRINTABLE:'California'
localityName          :PRINTABLE:'Palo Alto'
organizationName      :PRINTABLE:'TIBCO Software Inc.'
organizationalUnitName:PRINTABLE:'MiniHTTPS'
commonName            :PRINTABLE:'192.168.1.248'
emailAddress          :IA5STRING:'support@tibco.com'
Certificate is to be certified until May 12 01:23:11 2028 GMT (3650
days)

Write out database with 1 new entries
Data Base Updated
Creating mysql replication user repl_ll_601
Granting ALL PRIVILEGES on all nodes to repl_ll_601.
Granting ALL PRIVILEGES on node 192.168.1.245 to root
Granting ALL PRIVILEGES on node 192.168.1.249 to root

STOPPING MASTER TASK.............................................
(ok)

[writing new cluster configuration]

STARTING MASTER TASK...(ok)
done.
```

6. Repeat this procedure on the standby appliance.

> ⚠ **Warning:** On the standby appliance, in Step 4 indicate Y for the appliance to be the destination of automatic migration.
>
> If both appliances are configured as destinations, or if neither appliance is, the HA pair does not form. Both appliances report that they are out of cluster.

7. Confirm that the NTP server settings, and the actual observed time, are identical on both appliances in the HA pair.

It is essential to set the NTP server time correctly during appliance installation.

8. (Optional) Enable, disable, or check the status of the Advanced Features, as required. You must be logged in as the `root` user.

Feature	For more information about the CLI command...
Advanced Features	system logu
Monthly Index	system monthly_index
Monitoring Console	system monitoring_console
Advanced Aggregation	system advanced_aggregation

Important:
- Monthly Index, Monitoring Console and Advanced Aggregation can be enabled only if Advanced Features are enabled.
- In HA configuration, Advanced Features, Monthly Index, Monitoring Console and Advanced Aggregation can be enabled only through the CLI, and the configuration can be made only on the active node.

  Before configuring HA, you must disable Advanced Features on both active and standby appliances.
- The `zookeeper_sync` engine performs the sync from the active to the standby node.
- Enabling advanced features causes the `mtask` engines to reboot.

Result

The failover is now set up and both appliances are synchronizing their data.

You can log into the active node using the public IP address of the failover to finish appliance configuration. During normal operations, it is good practice to use the public IP address of the failover for configuration changes and the private IP address of the standby to run reports as this leaves the active node fully available for collecting and processing logs. The standby is always the same and only changes in case of failover. The **Administration > System Settings > Network** tab always shows the private IP address of an appliance.

Setting up Failover with a different interface

Procedure

1. The interface configured by running the command `set ip` is used to send replication traffic and heartbeats to another appliance and to send heartbeats to the other appliance as well.
2. If multiple such interfaces are to be configured, the user is asked to select which one is going to be used for HA replication and heartbeat.
3. The interface configured by running the command `set failover` is used to assign the public IP address and therefore receives logs. The LogLogic software requires that it be one of the interfaces configured by the command `set ip`.
4. You can select an interface other than any selected in Steps 1 and 2. The disadvantage of this configuration is the interface might still send heartbeats while logs are not received. The advantage is that replication traffic does not interfere with the logs, and may use a back-to-back crossover, avoiding dependency on a switch.
5. `Ip address of the peer appliance` is the destination of the replication traffic and of the heartbeats. It must therefore match the IP address assigned to the interface configured ON THE OTHER NODE during Steps 1 and 2.

> **Warning:** Remember that bond interfaces encapsulate actual network interfaces. Therefore, if a bond interface is configured, then the encapsulated network interfaces may not be configured separately. For that reason, it is not advisable to configure bond interfaces and non-bond interfaces on the same appliance.

Replacing a Single Node

Replacing the active or standby node in an HA pair can be either a planned exercise for replacing hardware, or an unplanned replacement due to a failover situation caused by a node failure.

In either situation, see Recovering the Active or Standby Node.

Setting up the Software to Replace an HA Pair

You can replace an existing HA pair with a new HA pair without interrupting log collection..

Before you begin

All four appliances must be running the exact same version of the LogLogic software. The failover membership cannot be established if there is a mismatch in software releases between the nodes. If you must upgrade an appliance, see the *TIBCO LogLogic® Log Source Packages Log Configuration Guides* for the release to which you are upgrading.

The last digit of the appliances' IP addresses cannot be the same, because the final digit is used as the failover node ID.

In this procedure:

- Appliance A - current active node (10.1.1.71)
- Appliance B - current standby node (10.1.1.72)
- Appliance C - new active node (new 10.1.1.72)
- Appliance D - new standby node (new 10.1.1.71)

Procedure

1. Ensure that appliances C and D use the same NTP server as appliances A and B.
    a. Log in to the GUI web server using the private IP address of the appliance.
    b. Configure the appliance time through the **Administration > System Settings > Time** tab.

    Note: Changing NTP settings on the active appliance in an HA pair reboots both appliances in the pair.

2. Shut down the current standby node, appliance B, without changing the failover configuration on the current active node, appliance A.
3. From the CLI, through a serial console, log in to appliance C.
4. Configure appliance C with the private IP address that appliance B was using:

```
set ip private-ip-address netmask gateway [ifdev] [defaultgw]
```

For example, using 10.1.1.72 as the private IP address and 10.1.1.1 as the gateway:

```
> set ip 10.1.1.72 255.255.255.0 10.1.1.1 bond0
New interface settings:
ip 10.1.1.72 255.255.255.0 10.1.1.1
CHANGES HAVE NOT BEEN SAVED!
> save
```

Note: If you configure more than one interface in the `set ip` command, make sure the subnets do not overlap.

5. Set failover on appliance C so it automatically migrates data from appliance A and then synchronizes with appliance A to become its standby node.

   For example, using 10.1.1.177 as the public IP address and 10.1.1.71 as the private IP address of the other appliance:

```
> set failover configure
Enter the public Ip address for the HA partner pair in the form
<ip> <netmask> <broadcast> <ifdev>:
10.1.1.177 255.255.255.0 10.1.1.255 bond0
Should this appliance be the destination of automatic migration: Y
Enter the Ip address of the peer appliance in the form <ip>:
10.1.1.71
CHANGES HAVE NOT BEEN SAVED!
>
```

6. Save the changes to apply the new configuration:

```
> save
Writing changes to disk...
Generating new SSL certificate...
Generating RSA private key, 1024 bit long modulus...
STOPPING MASTER TASK.............................(ok)
[writing new cluster configuration]
STARTING MASTER TASK...(ok)
done.
>
```

   When the migration and failover configuration completes, the status is displayed in the top-right corner of the top navigation bar:

```
failover: master
appliance-A (ok) - standby
appliance-C (ok)
```

7. After resynchronization completes, shut down appliance A.

   Due to the resulting failover, appliance C becomes the active node, and reports the other node as missing. On the GUI of appliance C, the status is displayed in the top-right corner of the top navigation bar:

   ```
   failover: master appliance-C (wait) - standby appliance-A
   (unavailable)
   ```

8. From the CLI, through a serial console, log in to appliance D.
9. Configure appliance D with the private IP address that appliance A was using.

   ```
   set ip private-ip-address netmask gateway [ifdev] [defaultgw]
   > set ip 10.1.1.71 255.255.255.0 10.1.1.1 bond0
   New interface settings:
   ip 10.1.1.71 255.255.255.0 10.1.1.1
   CHANGES HAVE NOT BEEN SAVED!
   ```

10. Set failover on appliance D so it automatically migrates data from appliance C and then synchronizes with appliance C to become its standby node.

    For example, using 10.1.1.177 as the public IP address and 10.1.1.72 as the IP address of the peer appliance:

    ```
    > set failover configure
    Enter the public Ip address for the HA partner pair in the form
    <ip> <netmask> <broadcast> <ifdev>:
    10.1.1.177 255.255.255.0 10.1.1.255 bond0
    Should this appliance be the destination of automatic migration: Y
    Enter the Ip address of the peer appliance in the form <ip>
    <netmask> <broadcast> <ifdev>: 10.1.1.72
    CHANGES HAVE NOT BEEN SAVED!
    >
    ```

11. Save the changes to apply the new configuration:

```
> save
Writing changes to disk...
Generating new SSL certificate...
Generating RSA private key, 1024 bit long modulus...
STOPPING MASTER TASK.............................(ok)
[writing new cluster configuration]
STARTING MASTER TASK...(ok)
done.
>
```

When the migration and failover configuration completes, the status is displayed in the top-right corner of the top navigation bar:

```
failover: master appliance-C (ok) - standby appliance-D (ok)
```

Result

The failover is now set up and both new appliances are synchronizing their data.

Failover Management

To monitor the current state of the failover setup, the appliance provides information through several mechanisms.

- Alerts can be configured to send email during an error condition; for example, if failover occurs or a resynchronization error occurs.
- Detail events are logged internally to record the history of the failover state.
- Configuration of Alert monitoring of failover events (and others) is performed on the **Alerts > Manage Alert Rules** page. All pre-configured System Alerts are visible on this page.

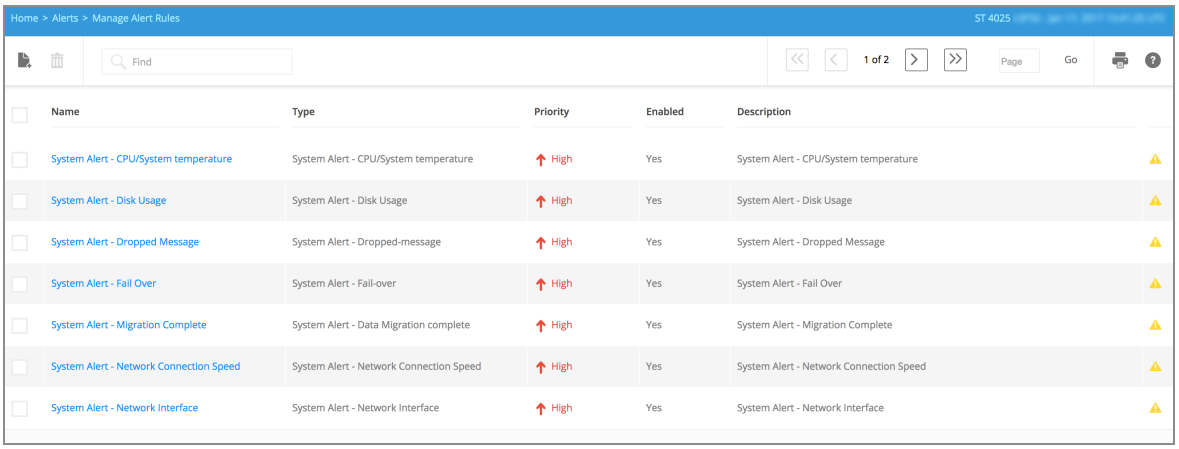

- Click **Alert Name** to open the **General** tab, where you may edit the preconfigured Alert settings, and select other System Alerts.
- The **Alert Receivers** tab appears next to the **Devices** tab (or **General** tab, for a System Alert) when you create a new alert. You must specify an alert receiver for which the alert can be triggered in this tab.

  The appliance can generate an alert to be sent to an alert receiver when the alert rule is triggered. The **Alert Receivers** tab lists all the available alert receivers configured for the appliance.
- Email recipients for System Alerts appear on the **Email Recipients** tab. All admin users are automatically selected to receive System Alerts. Other users can be configured on the **Management > Users > Privileges** tab.
- To search for all operational events on the LogLogic appliance, select **Reports > Operational > System Events** and click the **Create Report** button. Follow the onscreen prompts to generate the report.
- The RAS Warning page of the GUI displays the status of the failover in real time.

Failover Warnings

Failover issues warnings when certain tasks cannot be performed.

Failover Appliance Unavailable

When the failover membership cannot be formed, following message is displayed. The warning message describes the nature of the failover issue.

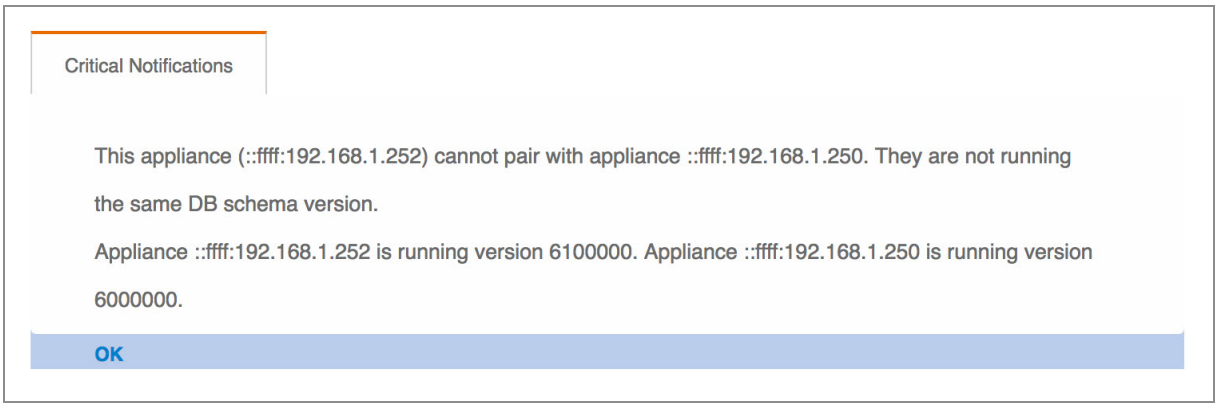

Failover in Degraded Mode

When the failover appliances appear to get out of sync, a message is displayed.

In this situation, a permanent failure of one of the appliances can result in data loss. If both nodes are mostly in sync, this resynchronization operation is relatively fast. However, if the data size difference is large, the operation time depends mainly on the network, CPU, and I/O bandwidth available. There are multiple phases involved in resynchronizing the failover, each with their own status and percentage to completion.

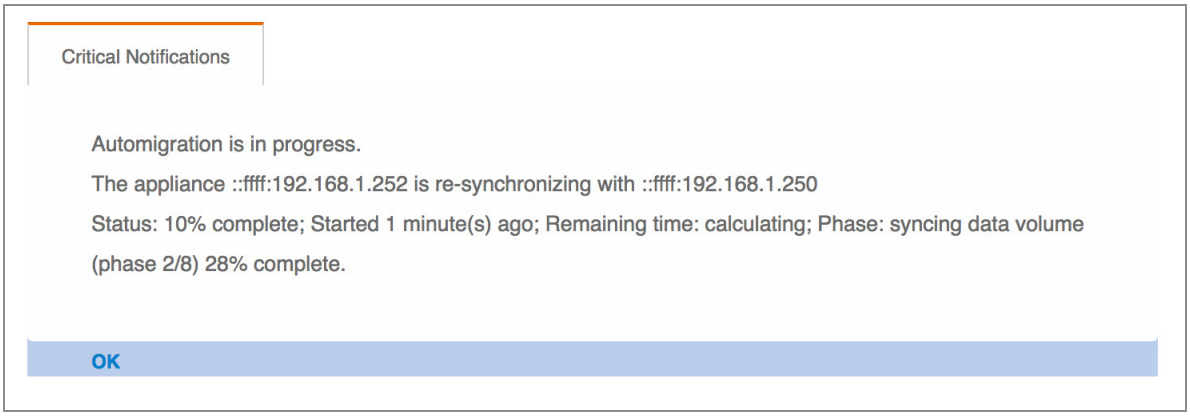

HA Software Upgrade

See the section Upgrading in a High Availability Environment in the *TIBCO LogLogic® Log Management Intelligence Configuration and Upgrade*.

Node Failure and Recovery

Failover can be caused due to node failure.

Three types of failure can occur with failover:

1. Node Failure
2. Recovering the Active or Standby Node
3. Double Failure

Node Failure

Loss of network connectivity is considered a node failure; that is, the miss of a heartbeat and failure to establish a TCP connection.

This can occur due to a hardware failure (the node is down), a network connection failure, or a network partitioning or a software error escalation. Every process running on the appliance is monitored and a repeated software error condition can trigger an escalation that restarts the appliance, hence triggering a node failure in the context of the failover.

Ethernet Disconnection

Unplugging the ethernet cable from a primary appliance in an HA pair triggers a failover.

HA pair (cluster) memberships fail, and eventually the primary appliance enters disabled mode. Plugging in the ethernet cable stops the failure state and allows the appliance to return to the running state, but it rejoins the cluster as the secondary node. What was the secondary appliance now becomes the primary appliance.

If you wish to have the two appliances take their old roles, another failover must be initiated. It is vital that the two appliances are allowed to synchronize before failover is triggered again; otherwise additional log loss occurs. Note that there is always be a small amount of log loss in any failover while the VIP is migrated from the old primary to the new primary appliance.

Run the following command as the `toor` user on the new primary appliance (after the first failover) which, after the second failover is triggered via this command, becomes the secondary appliance:

```
$ mtask -s engine_cluster_membership restart
```

Recovering the Active or Standby Node

If the current active node fails, the standby node takes over and becomes the active node.

When a node joins the failover it becomes the standby and a node resynchronization operation starts. The standby resynchronizes both the missing and existing data from the active node (including the currently modified data set) and the active node resynchronizes the missing data from the standby node.

Replacing the active or standby node in an HA pair can be either a planned exercise for replacing hardware, or an unplanned replacement due to a failover situation caused by a node failure.

The following is a sample Failover configuration using the following parameters:

- Active appliance: 10.20.0.10
- Standby1: 10.20.0.11 (old, to be replaced by Standby2)
- Standby2: 10.20.0.11 (new, uses same IP as Standby1)
- Public: 10.20.0.12
- Subnet Mask: 255.255.0.0
- Broadcasting: 10.1.255.255
- Gateway: 10.1.1.1
- NTP Server: 10.1.1.250

To replace the old Standby1 with a new Standby2, perform the following steps.

Procedure

1. Unplug all the network cables from the old Standby1 appliance (10.20.0.11) and plug them into the new Standby2.
2. Log in to the Standby2 appliance from a serial console, using root/logapp (username/password) to enter the CLI.

> **Note:** Use a serial console instead of an SSH connection. Network configuration changes can disconnect an SSH connection during this procedure. In all circumstances, a serial connection is maintained.

3. Set the IP address on Standby2 with the following command:

```
> set ip 10.20.0.11 255.255.255.0 10.1.1.1 bond0  > save
```

4. Configure failover on Standby2:

```
> set failover configure
```

   a. Enter the public IP address:

```
10.20.0.12 255.255.255.0 10.1.1.255 bond0
```

   b. When prompted about this appliance being the destination of automatic migration, enter Y.

   c. Enter the IP address of the peer appliance:

```
10.20.0.10   > save
```

Double Failure

Failure of a node before the initial resynchronization operation is complete is a double failure and can result in data loss.

As long as both nodes stay up long enough for the resynchronization operation to succeed, there is none or minimal loss of data.

If a failover occurs too frequently, call TIBCO Support. This might be due to an unsupported hardware setup. The failover membership can trigger spurious failover if the network bandwidth is insufficient. It is good practice to use only Gigabit Ethernet connections for the HA replication traffic.

Data Migration Between Appliances

Data migration lets you migrate data and configuration settings from one LogLogic appliance to another.

For a list of supported migration paths between LogLogic LMI appliances, see Supported Data Migration Paths.

The data migration solution is similar to replacing an appliance with a new one in a high availability configuration through failover. Like a failover from an active to a standby appliance, during migration the data is resynchronized while logs are still collected, thus minimizing down time. For more information about failover, see Failover.

Related Topics

When to Migrate Data

When to Migrate Data

The primary purpose of the data migration solution is to migrate log data and configuration settings from one LogLogic LMI appliance or LogLogic EVA to another.

It is particularly useful when upgrading from earlier LogLogic appliance models to newer models.

When migrating data, the following network connection configurations are supported:

- eth0 to eth0
- bond0 to bond0
- split mode (eth0 continues to collect data during migration from eth1 to eth1)

> **Important:** Do not perform backup or restore operations when data migration is in progress.

Important considerations

- Before performing data migration from LogLogic MX Appliance models to an LX4025R1 appliance, you must increase the global data retention on LX4025R1 to match that of the source LogLogic MX Appliance.
- When data migration is in progress, do not run any commands to enable or disable HA.
- With data migration enabled on the source and destination appliances, after data migration is complete, you cannot:
    - Use Advanced Features on the destination appliance.
    - Change the Advanced Features setting (enable or disable) on the source and destination appliances.
- After migrating data from a source appliance that does not support Advanced Features, an error might be displayed when using Advanced Features on the destination appliance. In such a scenario, you must reinstall LogLogic LSP on the destination appliance.

Supported Data Migration Paths

In case of LogLogic EVA, migration can be performed only from one LogLogic EVA to another. The following table lists the supported migration paths for LogLogic LMI appliances:

From: H4R1 or H4R2 appliances	To: H5 appliances
LX1025R1 or LX1025R2	LX1035
LX4025R1 or LX4025R2	LX4035
ST4025R1 or ST4025R2	ST4035
ST2025-SANR1 or ST2025-SANR2	ST2035-SAN

Data Migration on High Availability Appliances

Data migration is automatically performed whenever you configure failover on an appliance. This migration is necessary to ensure both appliances' data is in sync. For more information and procedures for failover, see Failover.

Note: Running Data Retention during data migration might slow down migration processing. Turning off Data Retention during data migration, might reduce the total migration time.

Data Migration From One Appliance to Another

Migration consists of the following tasks:

- Configuring the Source Appliances
- Monitoring the Migration
- Finishing the Migration

Pre-requisites

Before beginning data migration:

- Ensure that both appliances are running the exact same version of LogLogic LMI and LogLogic LSP, including hotfixes.
- If you must upgrade an appliance, see *TIBCO LogLogic® Log Management Intelligence Configuration and Upgrade* for the release to which you are upgrading.

  Note: You must complete the entire upgrade process, including the post-upgrade process to convert legacy data, if needed.

- If Advanced Features are enabled on the source appliance, you must enable them on the destination appliance.

Configuring the Source Appliances

To migrate data from an existing (source) appliance to a new (destination) appliance, you must configure data migration on the source appliance (10.0.20.31, in this example) and then on the new appliance (10.0.20.33).

Note: When migrating data for an HA pair, the source appliance being configured for failover satisfies the requirement. You do not need to disable failover and enable data migration on the source appliance.

Before you begin

If you are migrating data from a LogLogic MX Appliance to a LogLogic LX Appliance, you must run the migration script to match the number of maximum supported devices between LogLogic MX Appliance and LogLogic LX Appliance before actual migration starts. To run the migration script, perform the following steps:

1. Login to the LogLogic MX Appliance through the CLI, as a `toor` user.
2. Run the following command from the `/loglogic/scripts` directory:

   ```
   sh prep_migration_mx2lx
   ```

3. Type `yes` to continue.

Procedure

1. Log in to the source appliance through the CLI as a `root` user, using either a serial console or a remote SSH connection.
2. Enable data migration on the source appliance:

   ```
   > set data migration
   ```

3. Enter the IP address of the other appliance (for example, 10.0.20.33).
4. When the appliance prompts you to identify the direction of data migration, enter `1`.

   ```
   Select the data migration path
   ```

```
0) Do not setup data migration
1) This Appliance -> 10.0.20.33
2) 10.0.20.33 -> This Appliance
```

The appliance warns you that changes have not been saved.

5. Save the changes to apply the new configuration and restart the software:

   ```
   > save
   ```

   The appliance processes the changes, displaying the steps as they occur, and then informs you when it is done.

   Note: If you have multiple IP addresses you are prompted to select the IP to use for High Availability.

6. Log on to the source appliance through the GUI web server.
7. Verify the appliance is correctly receiving and processing logs.

   The dashboard reports an HA error because the new appliance is not yet configured.

Configuring the New Appliance

If you are migrating data to an existing appliance, it is recommended to delete all existing data on that appliance before performing a data migration to it.

Procedure

1. Log in to the new appliance through the CLI as a `root` user, using either a serial console or a remote SSH connection.
2. Enable data migration on the new appliance by running the following command:

   ```
   > set data migration
   ```

   The appliance prompts you to enter the IP address of the other appliance.

3. Enter the IP address (for example, 10.0.20.31).

   The appliance prompts you to identify the direction of the data migration:

```
Select the sense of the data migration
0) Do not setup data migration
1) This Appliance -> 10.0.20.31
2) 10.0.20.31 -> This Appliance
```

4. Enter 2.

   The appliance warns you that changes have not been saved.

5. Save the changes to apply the new configuration and restart the software by running the following command:

   ```
   > save
   ```

   The appliance processes the changes, displaying the steps as they occur, and then informs you when it is done processing changes.

   > **Note:**
   > If you have multiple IP addresses, you are prompted to select the IP address to use for High Availability.

   Data migration now begins. The data migration process includes five phases, all of which are internal and do not require interaction:

   a. Migration of the configuration database.

   b. Initial migration of the BFQ files.

   c. Migration of the parsed database tables.

   d. Verification of the migrated BFQ files.

   e. Final synchronization of the BFQ files, to capture recent modifications during migration processing.

      > **Note:** During migration, the source appliance continues to collect logs. The disk capacity is limited to the lower of the two appliances. The maximum rate of incoming logs is limited to the source appliance because logs are received by the source appliance.

Monitoring the Migration

You can monitor data migration between appliances.

Procedure

1. Log in to the source appliance through the GUI.
2. Choose **Alerts > Manage Alert Rules** from the menu.
3. In the Name column, select **System Alert - Migration Complete**.
4. Configure a Migration Complete alert with the settings you want.

   When you receive the Migration Complete alert, you can remove the source appliance.

   (Optional) You can configure alerts to send email during a failure condition, such as active failover or resynchronization error. For more information about configuring alerts, see*TIBCO LogLogic® Log Management Intelligence User Guide*.

   During migration, each time a user logs in to the GUI, a warning appears providing migration status:

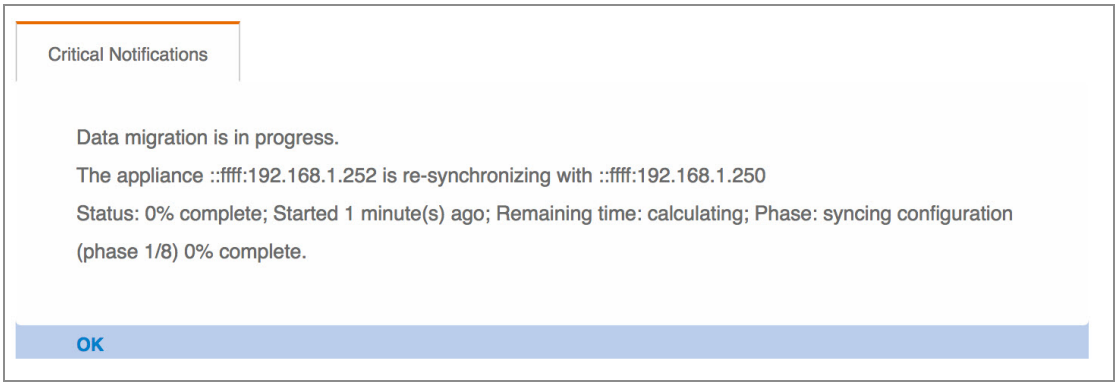

   You can check the status of the migration at any time on the GUI dashboard. Detailed events are internally logged to record system history.

Finishing the Migration

Perform these steps after you receive a System Alert indicating that migration processing is complete, or the GUI dashboard reports completion.

Procedure

1. Log in to the new appliance through the CLI as a `root` user.
2. Disable the data migration mode by running the following command:

   ```
   > unset data migration
   ```

3. Reconfigure the new appliance network settings by running the following command:

   ```
   set ip private-ip-address netmask gateway [ifdev] [defaultgw]
   ```

   For example:

   ```
   > set ip 10.1.1.177 255.255.0.0 10.1.255.255 bond0
   ```

   The appliance warns you that changes have not been saved.

4. Shut down the source appliance. This is required now if the new appliance is using the same network setting as the old appliance.
5. On the new appliance, save the new configuration (in step 3) and restart the software by running the following command:

   ```
   > save
   ```

   The appliance processes the changes, displaying the steps as they occur, and then informs you when it is done processing changes.

Result

The new appliance is now collecting and processing logs instead of the source appliance.

Software Update and Diagnostics

You can update the software on an appliance and get a diagnostics summary.

Update the Appliance Software

From the appliance GUI, use **Administration > File Update** to update the software running on an appliance.

The new software version number displays in the build details in the top-right corner of the top navigation bar, or on the **Administration > System Settings**.

> ⚠ **Warning:** The user must wait for at least one day after a software upgrade before doing a backup; otherwise the backed-up log data is inconsistent with the platform software.

Using File Update

You can use File Update to update your appliance.

Before you begin

- Ensure that you have at least 1GB of disk space available. To verify disk space, log in to the appliance through the GUI and on the **Dashboards > Advanced Dashboards > Advanced System Status** dashboard, check the size displayed on the **Free Disk Space** widget.
- You must download the .tar and .sig files from the TIBCO Support Website and copy them to the following directory (create it if it does not yet exist) of the appliance being upgraded:

  `/loglogic/update`

Procedure

1. Select a file from the **Select File** list.

   All available file updates are listed in the drop-down.

2. Click **Update** to begin the update.

   Note: The system might reboot during the update.

RAID Status

Redundant Array of Inexpensive Drives (RAID) is a method of increasing data reliability by using groups (or 'arrays') of small hard drives to create a single large drive.

There are several RAID configurations; TIBCO LogLogic® appliances use RAID 1, 5, 6, or 10, depending on the appliance model.

From the **Administration > RAID Status** page, you can view a high level of the condition of available drives. The available drives are listed followed by the status for each drive. This page shows only high level information.

This feature is available on the following appliance models:

LogLogic LX Appliances

Model name	RAID drives	RAID type
LX1025R1	2 drives: 2 x 1 TB	RAID 1
LX1025R2	2 drives: 2 x 1 TB	RAID 1
LX4025R1	8 drives: 8 x 4 TB	RAID 10
LX4025R2	8 drives: 8 x 4 TB	RAID 10
LX1035	2 drives: 2 x 2 TB	RAID 1
LX4035	8 drives: 8 x 4TB	RAID 10

LogLogic ST Appliances

Model name	RAID drives	RAID type
ST2025R1-SAN	8 drives: 8 x 1 TB	RAID 6
ST2025-SANR2	8 drives: 8 x 1 TB	RAID 6
ST4025R1	8 drives: 8 x 4 TB	RAID 10
ST4025R2	8 drives: 8 x 4 TB	RAID 10
ST2035-SAN	8 drives: 8 x 1 TB	RAID 6
ST4035	8 drives: 8 x 4 TB	RAID 6

To set up an alert to notify you if a RAID disk failure occurs, see Creating and Managing Alerts in *TIBCO LogLogic® Log Management Intelligence User Guide*.

For more information about the device configuration and hardware details, see the *TIBCO LogLogic® Log Management Intelligence Hardware Installation Guide*.

System Summary for Diagnostics

The **Administration > System Summary** tabs provide a variety of diagnostic information.

These tabs should be used only by, or as instructed by, TIBCO Support.

To update the information on any tab, click **Refresh**.

Process List

You can view a list of all processes running on an appliance.

Use the **Process List** to view a list of processes that are running on the appliance.

Click **Refresh** to update the list.

Network

Use the **Network** tab to view information about system configuration and operations.

This tab lists specific details of the network interfaces on your appliance. It should be used for diagnostic purposes only.

The following types of information about the network is displayed:

- **eth0**—The first ethernet interface.
- **eth1**—The second ethernet interface.
- **eth3**—The fourth ethernet interface

  > **Note:** The number of ethernet interfaces varies with appliance model.

- **bond0**—Two ethernet interfaces combined into a single bonded interface. See back of appliance for bond0 label, or refer to the *TIBCO LogLogic® Log Management Intelligence (LMI)Hardware Installation Guide*.
- **bond1**—Two ethernet interfaces combined into a single bonded interface. See back of appliance for bond1 label, or refer to the *TIBCO LogLogic® Log Management Intelligence (LMI)Hardware Installation Guide*.

Click **Refresh** to update the list.

SAN

You can view the Storage Area Network information for an appliance.

Use the **SAN** tab to view information about Storage Area Network (SAN) settings and details for the ST2025-SANR1 and H5 appliances. This information should be used for diagnostic purposes only.

DB Table Status

Use the **DB Table Status** tab to view information about system operations.

You can use this tab to easily view the status of the database tables used by your appliance. For example, Rows, Avg_row_length, or data being used, which can indicate problems with the database. This tab should be used for diagnostic purposes only.

Click **Refresh** to update the list.

> **Note:** In some cases, the value of `Create_time` may be displayed as `Null`. This is seen when no data is present in the table, and is expected due to the nature of the merge tables in the LogLogic LMI database.

Kernel Ring Buf

Use the **Kernel Ring Buf** tab to view information about system operations.

The information in the **Kernel Ring Buf** tab displays the last lines of the appliances operating system log file. This tab should be used for diagnostic purposes only.

Click **Refresh** to update the list.

Restart/Reboot/Shutdown

Use these tabs for system-wide operations.

- **Application Restart**—Restarts the services. This resets the internal counters used for alert generation.
- **Appliance Reboot**—Reboots the appliance. The appliance returns the login page after rebooting completes.
- **Appliance Shutdown**—Shuts down the system. After shut down completes, the system is accessible only through the direct connect console.

Click **Confirm** to continue with the operation, or **Cancel** to terminate the operation.

Removal of Appliance Log Data

If necessary, you can remove all log data on the LogLogic LMI appliance.

It is best practice to back up all log data on the appliance before clearing log data because after the log data is removed from the appliance, it cannot be recovered.

When the cleanup process starts, all active appliance components such as reports and alerts are brought to a dormant state. All running components such as reports are shut down immediately.

- On the GUI, use the **Administration > Clear Log Data** tab.
- On the CLI, run the following command:

```
/loglogic/tomcat/webapps/logapp20/WEB-INF/cgi/rttime
"type=resetappliancelogs"
```

Important Considerations

- The clear log data operation cannot be performed in an HA environment. To clear log data from a node, first remove the node from the HA environment and then perform the operation.
- An error might be displayed on the GUI in any of the following scenarios:
    - If you have initiated the operation using the CLI command and the operation is in progress
    - When the operation is in progress and the browser window is refreshed
    - While loading the login page
    - If multiple sessions are running simultaneously

  If an error is displayed, wait for the operation to complete, ensure that engines have started, and then launch the Login page in a browser.
- Log data that is archived to an external storage is not removed by the clear log data operation.

Shredding LogLogic LMI Event Data

To comply with the GDPR regulations, LogLogic LMI provides a CLI utility to shred selected data.

You can use the `llshred` utility on the result of an Advanced Search query to shred a list of events from the result. If the utility is run with `dryRun=true`, the log data is only processed. However, with `dryRun=false`, the log data is prefixed with the string `LLSHREDDER` and all characters are replaced with `X`. The number of `X`s is equal to the number of bytes in the original message.

For example, the following log entry includes 55 characters:

```
llfeed,03/23/2018 15:28:06.584,tcp,10252,19208,attitude
```

The utility replaces the log entry with the string `LLSHREDDER` followed by 45 `X`s:

```
LLSHREDDERXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
```

If an event is removed or its length changed, the references to the subsequent events become invalid. Therefore, the utility shreds the log events rather than deleting them. To shred events ingested into the appliance after running the utility, you must run the utility again.

Procedure

1. Back up the data on your appliance.

   > ⚠ **Caution:** The only way to recover shredded events is by restoring a backup. Therefore, you must take a backup before running the utility.

2. Run an Advanced Search query for the events you want to shred. The query must include at least the `sys_eventKey` column in the projection. The `sys_eventKey` column can be in any position in the projection.

   For example:

   ```
   SQL Query:
   select sys_eventKey, sys_body from system
   where sys_collectIP='192.168.56.1' and sys_body contains 'attitude'

   EQL query:
   use system |'attitude'| sys_collectIP='192.168.56.1'| columns sys_
   eventKey, sys_body
   ```

3. Save the search result as a .csv file.

4. Copy the .csv file to the appliance using a file transfer tool such as SCP.
5. From the CLI, run the llshred Command on the data in the .csv file to shred the data.

Impact of llshred Utility on Search and Reports

Advanced Search

If you specify a search term using Regex or `contains`, no result is returned. When you use any of the following fields (exclusively) in a `WHERE` clause or in the source filter of a data model, the search result returns the shredded log events:

- sys_eventTime
- sys_collectIP
- sys_deviceType
- sys_device
- sys_collectorDomain
- sys_eventKey
- sys_filename

The data before shredding is stored in the cache as long as the Advanced Search tab is open. To verify if the data has been shredded, you must close the tab in which the search was executed, so that the cache is cleared.

Index Search, Real-time Reports

Index search and real-time reports return the shredded logs (containing the `Xs`) in the results even after running the `llshred` utility.

Regex Search

Regex search does not return the shredded logs (containing the `Xs`) in the results after running the `llshred` utility.

Health Monitoring Utility

By using the health monitoring utility `hmonlog.py`, you can enable health monitor logging for aggregation node, correlation node, data node, query node, and webapp.

Enabling logging

Health monitor logging is disabled by default, and no results are returned for a search query that uses the following data models. To enable data logging, contact your administrator.

- LogLogic_Monitor_Cpu
- LogLogic_Monitor_Cpu_Load
- LogLogic_Monitor_Diskspace
- LogLogic_Monitor_Memory
- LogLogic_Monitor_Node_Memory

By using the utility, you can enable it on the specified nodes by running the following command:

```
hmonlog.py [-h] [--action {enable | disable | status} ] [--nodes NODES]
```

The following table describes parameters for the hmonlog.py command.

Parameter	Description
-h, -help	Displays the help for the command and exits to the command prompt.
--action {enable | disable | status}	Enables or disables the health monitor log. `status` displays the status of the health monitor log.
--nodes	Specifies that the command should be applied to the nodes listed in the NODES parameter.
NODES	Specifies the list of nodes to which the command must be applied. By default, it applies to all nodes. Enter a comma-separated list of nodes. For example: an, cn, dn, qn, web

After enabling or disabling logging for the webapp node, you must restart the `lldaemon` engine by running the following commands:

```
mtask -s engine_lldaemon restart
```

Note: After restarting `engine_lldaemon`, it might take some time for the GUI operations to be available.

Command examples

Command	Description
`python hmonlog.py --action enable --nodes all`	Enables the health monitoring log for all nodes.
`python hmonlog.py --action enable --nodes cn, dn`	Enables the health monitoring log for correlation node and data node.

For an HA setup

If you are restarting the engine in an HA setup, you must restart the engine on both nodes, in the following sequence:

1. Stop the engine on the standby node and then on the active node.
2. Start the engine on the active nodeand then on the standby node.

IPv6 Support

LogLogic supports IPv6.

About IPv6

Internet Protocol version 6 (IPv6) is the latest version of the Internet Protocol (IP), and provides an identification and location system for computers on networks and routes traffic across the Internet.

IPv6 permits hierarchical address allocation methods that facilitate route aggregation across the Internet, and thus limit the expansion of routing tables. The use of multicast addressing is expanded and simplified, and provides additional optimization for the delivery of services.

IPv6 is based on IP but with a much larger address space and improvements such as a simplified main header and extension headers. The IPv6 address space allows networks to scale and provide global reachability.

The primary motivation for IPv6 is the need to meet the demand for globally unique IP addresses. IPv6 quadruples the number of network address bits from 32 bits (in IPv4) to 128 bits, which provides more than enough globally unique IP addresses for every networked device on the planet. By being globally unique, IPv6 addresses inherently enable global reachability and end-to-end security for networked devices, functionality that is crucial to the applications and services that are driving the demand for the addresses. Additionally, the flexibility of the IPv6 address space reduces the need for private addresses and the use of Network Address Translation (NAT); therefore, IPv6 enables new application protocols that do not require special processing by border routers at the edge of networks.

IPv6 Address Formats

IPv6 addresses are represented as eight groups of four hexadecimal digits separated by colons (:) in the format: x:x:x:x:x:x:x:x.

Following are two examples of IPv6 addresses:

- 2001:DB8:7654:3210:FEDC:BA98:7654:3210
- 2001:0db8:85a3:0042:1000:8a2e:0370:7334

It is common for IPv6 addresses to contain successive hexadecimal fields of zeros. To make IPv6 addresses less cumbersome, two colons (::) may be used to compress successive hexadecimal fields of zeros at the beginning, middle, or end of an IPv6 address (the colons represent successive hexadecimal fields of zeros). Table 1 lists compressed IPv6 address formats.

A double colon may be used as part of the ipv6-address argument when consecutive 16-bit values are denoted as zero. You can configure multiple IPv6 addresses per interfaces, but only one link-local address.

LogLogic Support for IPv6

LogLogic supports IPv6 implementation.

- Native connectivity to log sources.
- Regex and Index searches.
- Parsed reports against IPv6 Log Sources.
- Collection of IPv6 data from LogLogic® Universal Collector 2.6 and later.
- Network configuration:

  IPv6 and IPv4 addresses can be assigned to the same or different interfaces.

  > **Note:** Dual stack IPv6 support is not available for HA VIP interface.

- Collection from IPv4 sources by an appliance with an IPv6 address:

  Direct collection is only possible if the appliance is accessible from an IPv4 address; either by assigning an appropriate v4 address, or by providing network address translation externally. In the case of external translation, the source address reported for the source devices is entirely dependent on the external translation performed.

- Forwarding log messages between IPv4 and IPv6:

  Forwarding from an IPv6 addressable appliance to an IPv4-only appliance is only possible if the forwarding appliance has v4 connectivity; either by assigning an appropriate v4 address, or by providing external network address translation.

Apparent source address of traffic forwarded to an IPv4 address:

Source	Protocol	Description
IPv6 source	any protocol	Always appears as the original source address, but an IPv6 notation is added as a prefix.
IPv6 source	lltcp	In LogLogic LMI 5.6 or later, appears as the original source address.
IPv6 source	lltcp	Prior to LogLogic LMI 5.6, appears as the last 4 bytes of the original source address.
IPv6 source	syslog	Appears as the last four bytes of the original source address.

Apparent source address of traffic forwarded to an IPv6 address:

Source	Protocol	Description
IPv4 source	any protocol	Always appears as the original source address.
IPv6 source	lltcp	In LogLogic LMI 5.6 or later, appears as the original source address.
IPv6 source	lltcp	Prior to LogLogic LMI 5.6, not applicable.
IPv6 source	syslog	Should appear as the original source address, this is not recommended (relies on theoretically illegal v4-mapped IPv6 source address in UDP packets) as some network routers may choose to discard such packets.

- Display of IPv4 source addresses:

  In most cases IPv4 addresses are displayed, emailed, or reported in their traditional dotted-decimal notation. The following exceptions exist (displayed as v4-mapped addresses in V6 notation):

    - recent messages
    - unapproved messages
    - address appearing in the automatically generated names of auto discover devices
    - address appearing in backup file names
- IPv6 addresses for external servers:

  The following external services are supported with IPv6 addresses:

    - NFS and SCP backup
    - NFS and SCP archival
    - NTP
    - DNS (including resolution of names to v6 addresses)
    - SMTP
    - Active Directory
- Static routes:

  LogLogic LMI v5.6 supports both v4 and v6 static routes.

IPv6 addresses are not supported for:

- Check Point LEA
- Parsing of address strings within log messages
- Global groups
- Replay
- Cisco IPS

IPv6 Support Matrix

Specific LogLogic LMI versions provide support for IPv6.

The following example shows the IPv6 support matrix:

Log Source Address	LogLogic LMI	Supported LogLogic LMI version
IPv4	IPv4	v5.5.0 and earlier
IPv6	IPv6	v5.6.0 and later
IPv4	IPv4 + IPv6	v5.6.0 and later

Configuring Oracle JDBC Driver for IPv6 Support

You can configure Oracle JDBC driver for IPv6 support.

Procedure

1. Download a supported Oracle JDBC driver (11g or 12c).
2. Rename the driver to `oracle-10gr2-ojdbc14.jar`.
3. SSH to LogLogic LMI.
4. Rename the `oracle-10gr2-ojdbc14.jar` file to `oracle-10gr2-ojdbc14.org`.
5. Copy the `oracle-10gr2-ojdbc14.org` file to the `/loglogic/tomcat/webapps/logapp20/WEB-INF/lib` directory on the LogLogic LMI appliance. Make sure that you replace the original file.
6. Restart `mtask` by running the following commands:
    a. `$ mtask stop`
    b. `$ mtask start`

Command Line Interface (CLI)

Using the console command line interface (CLI) you can set up, configure, and maintain a LogLogic Appliance.

The console is the terminal for accessing the appliance. For detailed help on a specific command, type:

```
help cmd
```

or

```
? cmd
```

Connecting to the Appliance

Before you can use any command line options, you must connect to the appliance.

Use a laptop or other terminal device to make this connection. All commands are logged internally to enhance appliance security.

Procedure

1. Use a null modem cable to connect the appliance to COM1.
2. Open a terminal utility.

   > ⚠ **Warning:** It is good practice to connect to the CLI through a serial console, and not using SSH, when issuing network configuration commands such as set failover. Network configuration changes might reconfigure the network card, disconnecting an SSH connection.

3. Set the communication setting from the terminal login dialog. For example:

   9600 baud, Null, 8 bit, 1

baud rate: 9600

data bits: 8

parity: none

stop bits: 1

4. Log in to the appliance in the console mode.
5. From the terminal program, log in as user `root` with the default password `logapp`. There are two passwords for shell and CLI. By default, both have the same password.

> **Note:** Change the default password for the CLI and shell login. To change the password, type system passwd Usage > at the command prompt and follow the prompts.

Result

Ping the appliance. You must be able to successfully ping the appliance.

exit Command

The `exit` command exits from the login shell and system.

This command has no arguments or keywords.

Example

To exit the session:

```
> exit
```

network Command

The `network` command has options that let you activate, deactivate, or restart all network interface(s), or ping a specific system on the network.

Type the following command from your command line.

```
network  [ start | stop | restart | ping ip-address ]
```

network Syntax Parameters

Parameter	Description
start	Activates network interface(s).
stop	Deactivates network interface(s).
restart	Deactivates and then activates network interface(s).
ping *ip-address*	Pings the specified IP address.

The network `ping` command determines network connectivity. When using ping for fault isolation, you initially run it on the local host, to verify that the local network interface is up and running. Then, ping hosts and gateways further and further away.

The `ping` command uses the ICMP protocols mandatory ECHO_REQUEST datagram to elicit an ICMP ECHO_RESPONSE from a host or gateway. ECHO_REQUEST datagrams (ping) have an IP and ICMP header, followed by a strict time value and then an arbitrary number of pad bytes used to fill out the packet.

Examples

To restart all network interfaces:

```
> network restart
Removing default gateway...
[ OK ]
Bringing down the eth0 interface...
[ OK ]
Bringing down the eth1 interface...
[ OK ]
Bringing up the eth0 interface...
[ OK ]
Bringing up the eth1 interface...
                     [ OK ]
                     Setting up default gateway...
```

```
                        [ OK ]
```

To determine network connectivity with the system that has IP address 10.1.1.222:

```
> network ping 10.1.1.222
PING 10.1.1.222 (10.1.1.222): 56 octets data
64 octets from 10.1.1.222: icmp_seq=0 ttl=64 time=2.1 ms
64 octets from 10.1.1.222: icmp_seq=1 ttl=64 time=1.0 ms
64 octets from 10.1.1.222: icmp_seq=2 ttl=64 time=0.8 ms
64 octets from 10.1.1.222: icmp_seq=3 ttl=64 time=1.2 ms
64 octets from 10.1.1.222: icmp_seq=4 ttl=64 time=1.3 ms

--- 10.1.1.222 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 0.8/1.2/2.1 ms
```

plugin Command

The `plugin` command has options that let you register, unregister, activate, or deactivate, new plugins for device.

Type the following command from your command line.

```
plugin [ add | del | enable | disable | view | view all ]
```

`network` Syntax Parameters

Parameter	Description
add<devID> <devFormat> | <prog> <execType> | <progType> | <keyIn> <keyout>	Registers new plugin for device (deviceID, devformat).
del <devID> |	Unregisters plugin for device (deviceID,

Parameter	Description
<devFormat>	devFormat)
enable <devID> <devFormat>	(Re)activates plugin registered for (devID, devFormat)
disable <devID> <devFormat>	Deactivates plugin registered for (devID, devFormat)
view <devID> <devFormat>	Displays status of plugin registered for (devID, devFormat)
view all	Displays status of all registered plugins where: <devID> - numerical device ID <devForamt> - numerical device format ID <prog> - plugin executable path name <execType> - [binary | shell | script |...] <progType> - [0 |1 |2], use 1 for now <keyIn> - plugin input format specific key <keyOut> - plugin output format specific key

raid Command

The raid command shows hard drive information of the hardware RAID.

Note: You must be careful before changing any raid setup.

The raid command displays its submenu as shown next.

```
> raid
//localhost> ?
Copyright(c) 2004, 2005 Applied Micro Circuits Corporation(AMCC). All
rights reserved.
```

```
AMCC/3ware CLI (version 2.00.03.008)
Commands  Description
--------------------------------------------------------------------------
info      Displays information about controller(s), unit(s) and port(s).
maint     Performs maintenance operations on controller(s), unit(s) and
ports.
alarms    Displays current AENs.
set       Displays or modifies controller and unit settings.
sched     Schedules bachground tasks on controller(s)             (9000
series)
quit      Exits the CLI.
          ---- New Command Syntax ----
focus     Changes from one object to another.  For Interactive Mode
Only!
show      Displays information about controller(s), unit(s) and port(s).
flush     Flush write cache data to units in the system.
rescan    Rescan all empty ports for new unit(s) and disk(s).
commit    Commit dirty DCB to storage on controller(s).        (Windows
only)
/cx       Controller specific commands.
/cx/ux    Unit specific commands.
/cx/px    Port specific commands.
/cx/bbu   BBU specific commands.                                  (9000
only)
```

Type `help command` to get more details about a particular command.

save Command

The `save` command saves system configuration settings such as setting IP or failover.

The format to use this command is:

```
save
```

Example

To save your system configuration settings to disk:

```
> save
```

set Command

The `set` command sets up the system IP address, DNS server IP address, Ethernet type, system clock and time zone, NTP server IP address, and failover.

After these tasks are complete you can access the appliance through the GUI.

Warning:

- Set up your failover system before using the appliance(s).
- When using SHA256 instead of MD5 message digests, appliance performance may be reduced by as much as 20%.
- It is good practice to connect to the CLI through a serial console, and not using SSH, when issuing network configuration commands such as `set failover`. Network configuration changes might reconfigure the network card, disconnecting an SSH connection.

Type the following command from your command line, using the appropriate parameter. For more information about a parameter, click the parameter name.

```
set [ clock | data migration | digest | dns | ethn | failover | ip |
ipv6 | ntpserver | regexsearches | reverse_forward | strong_passwd |
timezone | tls syslog key password]
```

`set` Syntax Parameters

Parameter	Description	Options
clock	Sets the system date and time.	set clock takes an option in the format: `MMDDhhmm[[CC]YY][.ss]`
data migration	Configures the appliance for data migration. After entering the command, the appliance prompts you to identify which migration path to use. You must run the command on both the source and destination appliances.	None

Parameter	Description	Options
digest	Sets the appliance SHA Digest. The default is the 128-bit MD5 Digest. If the digest setting is changed, the appliance is restarted to synchronize the log data collection processes to use the new SHA Digest. Usage of the 256-bit SHA2 Digest can reduce the maximum message handling rate of the appliance up to 20%.	The command takes one of the following options: • SHA256 • MD5 • default
dns	Queries the Internet Domain Name System (DNS) for host information. This command helps to convert host names into IP addresses and vice versa.	This command takes one option: `dns-server-ip-address`
eth*n*	Changes network card settings. *n* is the number of the interface (eth0, eth1, and so on).	This command takes one of the following options: • 100baseTx-FD • 100baseTx-HD • 10baseT-FD • 10baseT-HD • 1000baseTx-FD • 1000baseTx-HD • auto
failover	Assigns or resets failover active and standby appliance roles.	This command takes one of the following options: `configure | disable` `configure` - Configures the active and standby appliances

Parameter	Description	Options
	Important: Before disabling failover, you must disable Advance Features by running the system command. After configuring failover, you can reenable the Advanced Features from the CLI on the active appliance.	for a failover. `disable` - Resets the active and standby appliances during a failover, removes the data migration configuration for the appliance, and stops data migration. The command prompts you for several options before taking certain actions. For more information, see Failover.
ip	Configures the kernel-resident network interfaces on the appliance.	This command takes several options: `ip-address netmask gateway [ifdev] [defaultgw]` `ifdev` specifies eth0, eth1, eth2, or bond0. The default is bond0. `defaultgw` specifies the default gateway. Optional for specific NICs, but one NIC must be specified. The last gateway specified in `defaultgw` is in effect. **Note:** The `defaultgw` keyword no longer has any effect and is allowed only for backward compatibility. When asked whether you want a certificate generated for Blue Coat when you also plan to use

Parameter	Description	Options
		the TLS TCP syslog feature, ensure that you choose yes. Alternatively, you can use a custom certificate.
ipv6	Configures the kernel-resident network interfaces on the appliance with IPv6 address.	`ipv6-address ipv6-prefix gateway [ifdev] [defaultgw]` `ifdev` specifies the network interface name or bond interface, like eth0, eth1, eth2, or bond0. `defaultgw` specifies the default gateway. Optional for specific NICs, but one NIC must be specified. The last gateway specified in `defaultgw` is in effect. **Note:** The `defaultgw` keyword does not have any effect and is allowed only for backward compatibility.
ntpserver	Sets the network time server.	This command takes one of the following options: `ipaddress | hostname`
regexsearches	Sets the number of simultaneous regular expression searches that the appliance can run.	This command takes one option: `limit`
reverse_forward [disable | tunnel_init on| tunnel_init off| tunnel_init add <ip_ address>|tunnel_init		`disable` - Disables reverse tunnel. `tunnel_init on` - The appliance tries to initiate tunnels to the partners

Parameter	Description	Options
delete <ip_address>| tunnel_accept [on|off]]		configured. `tunnel_init off` - The appliance does not try to initiate tunnels. `tunnel_init add <ip_address>` - To add a LogLogic appliance IP address to initiate a tunnel to. `tunnel_init delete <ip_address>` - To remove the tunnel to the LogLogic appliance with the specified IP address. `tunnel_accept [on|off]` - The appliance does not accept tunnel connections.
strong_passwd [enable | disable | settings | expiration]	Controls the use of strong passwords for user authentication through the CLI on the appliance. (To set strong passwords for GUI access, see System Settings.) After disabling strong passwords, all settings are retained, but are only effective when strong passwords are enabled.	`enable` - turns on the requirement of strong passwords for appliance users `disable` - turns off the requirement of strong passwords for appliance users `settings` - sets the strong password requirements for the appliance. This command requires five options, as follows: • `lowermin` - Minimum required lowercase letters (default and minimum = 1) • `uppermin` - Minimum

Parameter	Description	Options
		required uppercase letters (default and minimum = 1) • `digitsmin` - Minimum number of numeric digits (default and minimum = 1) • `nonalphanum` - Minimum number of non-alphanumeric characters (default and minimum = 1) • `minlength` - Minimum number of total characters in the password (default = 15; minimum is 6 or the sum of the other four settings, whichever is greater) `expiration` - the number of days after which a user password expires on the appliance (1 through 99999 or never)
timezone	Sets the time zone conversion. A time zone table displays with all possible selections. Enter a selection from this time zone table.	None
tls syslog key password	Sets the TLS key password if you are enabling TCP TLS support for the TCP collector. For information about how to change and clear	None

Parameter	Description	Options
	passwords, see the Configuring TLS Syslog section.	

Examples

To set up a failover configuration for your Appliances:

On the active appliance:

```
> set failover configure
Enter the public Ip address of the cluster
in the form <ip> <netmask> <broadcast>:
CHANGES HAVE NOT BEEN SAVED!
> save
Writing changes to disk...
Removing default gateway...
Bringing down the eth0 interface...
Bringing down the eth1 interface...
Bringing up the eth0 interface...
Bringing up the eth1 interface...
Setting up default gateway...
Bringing down the eth1 interface...
Bringing up the eth1 interface...
done.
```

On the standby appliance:

```
> set failover configure
CHANGES HAVE NOT BEEN SAVED!
> save
Writing changes to disk...
Removing default gateway...
Bringing down the bond0 interface...
Bringing up the eth0 interface...
Setting up default gateway...
Bringing down the eth1 interface...
Bringing up the eth1 interface...
done.
```

To disable the failover configuration:

On the Standby system:

```
> set failover disable
> save
Writing changes to disk...
Removing default gateway...
Bringing down the eth0 interface...
Bringing down the eth1 interface...
Bringing up the bond0 interface...
Setting up default gateway...
done.
```

On the Active system:

```
> set failover disable

CHANGES HAVE NOT BEEN SAVED!
> save
Writing changes to disk...
hecking network configuration now...


[IPv4] Please select the network interface that will be the default
gateway.

0. 192.168.1.245  eth0
1. Do not save this configuration. Exit now.

> 0

The default gateway has been designated. Thank you.

Please select the IP address to use to generate the BlueCoat
certificate.

0. 192.168.1.245  eth0
1. Do not generate the BlueCoat certificate.

> 1

The BlueCoat certificate will not be generated. Thank you.

STOPPING MASTER TASK...
```

```
[writing new cluster configuration]

STARTING MASTER TASK...(ok)
done.
```

To set up network IP addresses for Ethernet interface 0:

```
> set ip 10.1.1.10 255.255.255.0 10.1.1.255 eth1
> show changes
Current changes that have not been saved:
ip address eth0 10.1.1.10 255.255.255.0 10.1.1.255
CHANGES HAVE NOT BE SAVED!
> save
```

To enable strong passwords and set each character minimum to 2, total minimum to 8, and expiration to 90:

```
> set strong_passwd enable
> set strong_passwd settings 2 2 2 2 8
> set strong_passwd expiration 90
```

show Command

The `show` command display the status of the current system interface information that is stored on the disk, history of changes made during this session, pending changes not yet saved, and current system date and time.

Type the following command from your command line.

```
show  [ current | digest | history | changes | date | regexsearches |
reverse_forward | san_ports | san_ devices | strong_passwd ]
```

show Syntax Parameters

Parameter	Description
current	Current interface information stored on disk.

Parameter	Description
digest	Shows the current SHA Digest being used by the Appliance.
history	History of saved changes for the current session.
changes	Pending changes that are not yet saved.
date	Current system date and time.
regexsearches	Shows the current simultaneous regular expression searches allowed, and the maximum number of simultaneous regular expression searches.
reverse_forward	Shows the status of tunnel_accept [on|off], tunnel_ init [on|off], partner IPs, and reverse forwarded status.
san_ports	Shows SAN interface information.
san_devices	Shows attached SAN devices.
strong_passwd	Shows the current settings for strong password requirements on the appliance for CLI users.

Examples

To show the current date and time on an Appliance:

```
> show date
Current Time:
       Wed Jul  7 22:09:44 CDT 2004
```

To show the current strong password settings on an Appliance:

```
> show strong_passwd
Strong password: disabled
Strong password settings:
    The minimum number of lower case letters: 1
    The minimum number of upper case letters: 1
    The minimum number of digits: 1
    The minimum number of non-alphanumeric characters: 1
    The minimum password length: 6
    Require password change after (days): never
```

swraid Command

The `swraid` command displays the status of software RAID devices.

Type the following command from your command line.

```
swraid
```

Examples

For software-raided model:

```
> swraid
Personalities : [raid1]
md2 : active raid1 sdb2[1] sda2[0]
      2097088 blocks [2/2] [UU]
md3 : active raid1 sdb3[1] sda3[0]
      972568192 blocks [2/2] [UU]
md1 : active raid1 sdb1[1] sda1[0]
      2096064 blocks [2/2] [UU]
unused devices: <none>
```

For non-software-raided model:

```
> swraid
swraid : No RAIDs defined
```

system Command

The `system` command implements system-wide changes.

Type the following command from your command line.

```
system [access | advanced_aggregation | data_client | data_vault | fips
| firewall | fsck | halt | iptables | ipv6_slaac| keycopy | logu |
monitoring_console | monthly_index | monthly_index_load_divisor | passwd
| reboot | secureuldp | sshkey_passphrase | storage_growth | update]
```

`system` Syntax Parameters

Parameter	Description	Options
access	Grants full access to the application. When Data Privacy mode is disabled, only one password is required to gain the access. The password can be changed using the `system passwd` command. When Data Privacy mode is enabled, the two Security Keys are required to gain access. You **cannot** change the Security Keys using the `system passwd` command. However, you can use the GUI (from **Administration > System Settings > General > Data Privacy Options**) to reset your Security Keys, see Data Privacy Settings.	None
advanced_ aggregation	Enable or disable the Advanced Aggregation features. After enabling the Advanced Aggregation option, the **Management > Rules > Aggregation** tab is visible to users and they can use the Advanced Aggregation features. By default, Advanced Aggregation is switched off. This feature can be enabled only if the Advanced Features option is enabled.	• `enable` - Enables the Advanced Aggregation feature. • `disable` - Disables the Advanced Aggregation feature. • `status` - Displays whether the Advanced

Parameter	Description	Options
	Important: Before disabling Advanced Aggregation, ensure that you delete or disable any advanced aggregation rules to avoid storing unnecessary aggregated data. To enable or disable this feature using the GUI, see the Advanced Aggregation setting.	Aggregation feature is enabled or disabled.
data_client	Creates or deletes a user account.	`add <username>` - Creates a new account, the following constraints apply to user names: • The first character of the username must be lower/upper case letter, or a number. • All characters, except the first character, must be lower/upper case letters, numbers, underscore character ('_') or period character ('.') `delete <username>` - Deletes the existing user account `list` - Displays all existing user accounts
data_vault	Manages encryption of all data volumes	`enable` - Enables the

Parameter	Description	Options
	including archives. By default, the data vault is disabled and the data volumes are in unlocked state. • If the Data Vault feature is enabled on LogLogic EVA, auto unlock is disabled, and you want to attach additional hard drives, run the following command before adding additional hard drives: `system data_vault enable_auto_unlock_once` When the appliance restarts, the saved encrypted password is used to automatically unlock the data vault. • If the Data Vault feature is enabled on LogLogic LMI and you are upgrading to 6.3.0, you must install the hotfix LMI-6.2.1_6.2.0-HF-LLCE-3207-3210 to first decrypt the data volumes. Then encrypt them again by using the `system data_vault` command. For more information, see Data Encryption.	encryption of data volumes. `status` - Displays whether the Data Vault feature is enabled or disabled. `unlock` - Unlocks the data vault after system reboot. `change_password` - Changes the password of the data vault. `enable_auto_unlock` - Saves the encrypted password to be used for automatically unlocking the data vault at boot time. `disable_auto_unlock` - Removes the saved password to be used for automatically unlocking the data vault at restart time. `enable_auto_unlock_once` - Automatically unlocks the data vault using the saved encrypted password, for only the next system restart. The password is deleted after one use.
fips	Enables or disables the Federal Information	To enable FIPS mode,

Parameter	Description	Options

	Processing Standard (FIPS) mode on the appliance. FIPS libraries are preinstalled in LogLogic LMI. Enabling the FIPS mode ensures that FIPS-compliant libraries are used during secure communication.	run the command:
		`> system fips enable`
		When prompted, type `yes` to reboot the appliance for the changes to take effect.
		To disable the FIPS mode, run the command:
		`> system fips disable`
		When prompted, type `yes` to reboot the appliance for the changes to take effect.
		Note: In an HA setup, disable the failover on both appliances, enable the FIPS mode, and then reenable the failover.
		`status` - Displays whether FIPS is enabled or disabled.

firewall	Configures the firewall setting. On the GUI, the firewall can be configured from **Administration > Firewall Settings**. See Adding an Input Rule.	`enable` - Enables the firewall. `disable` - Disables the firewall. `status` - Displays

Parameter	Description	Options
		whether the firewall is enabled or disabled. `list` - Displays a list of firewall rules in the system. `add <All/SingleIp/CIDR> <port> <TCP/UDP> <accept/deny>` - Adds a new set of IP address (All or Single IP/ CIDR), port number, protocol (TCP or UDP), and action (accept or deny). **Note:** The value is case-sensitive. `remove` - Removes a set of IP address, protocol, port number, and action. `port <add/remove> <TCP/UDP> <port> <desc> >` - Adds or removes a port for use in a firewall rule.
fsck	Performs a file system check.	`enable` - Enables fsck check on system reboot or startup. `disable` - Disables fsck check on system reboot or startup. `status` - Displays

Parameter	Description	Options
		whether fsck is enabled or disabled.
halt	Halts the appliance.	None
iptables	Enables or disables the appliance iptables. This can be used for Firewall Settings.	`on` - Enables the appliance iptables. `off` - Disables the appliance iptables.
ipv6_slaac	Manages the Stateless Autoconfiguration (SLAAC) feature of IPv6. By default, the feature is turned off.	`enable` - Enables SLAAC. `disable` - Disables SLAAC. `status` - Displays whether SLAAC is on or off.
keycopy	By default, copies the RSA public key of the LogLogic product family to establish secure file transfer access with another server. The public key is used for user authentication when transferring files using the secure protocols SCP or SFTP.	`dsa` - Copies the Digital Signature Algorithm (DSA) public key to the target server. This parameter is available for backward compatibility.
logu	Enables or disables the Advanced Features. The default is No. After running `logu enable`, you must exit from the root shell for `mtask` to restart and the changes to take effect. To enable or disable this feature from the GUI, see Advanced Features .	`enable` - Enables the Advanced Features. `disable` - Disables the Advanced Features. `status` - Displays whether Advanced Features are enabled or disabled.
monitoring_	Enables or disables the Monitoring Console and	`enable` - Enables the

Parameter	Description	Options
console	displays the **Monitoring > Console** menu. This feature can be enabled only if the Advanced Features option is enabled. To enable or disable this feature from the GUI, see Monitoring Console.	Monitoring Console. `disable` - Disables the Monitoring Console.
monthly_index	Enables or disables the Monthly Index feature. The default is No. This feature can be enabled only if the Advanced Features option is enabled. To disable archiving of indexes while the raw data is archived, see Monthly Index.	`enable` - Enables the Monthly Index feature. `disable` - Disables the Monthly Index feature. `status` - Displays whether the Monthly Index feature is enabled or disabled.
monthly_ index_load_ divisor	Controls what fraction of the monthly index terms are loaded into memory during an Advanced Search. To enable or disable this feature from the GUI, see Monthly Index Load Divisor.	`show` - Displays the value of monthly index load divisor. `set [1-5]` - Sets the value of the monthly index load divisor.
passwd	Changes the password for the CLI or system account. If an old password is present, the system prompts you for the old password and compares it against the stored password. After the system authenticates the user, password aging information is checked to see if the user is permitted to change their password. If the user is authenticated, the system prompts for a replacement password. If the password is accepted, `passwd` prompts again and compares the second entry against the first. Both entries must match to successfully change the	This command with no option indicates to change the password for CLI or shell access. `cli` - Change password for the CLI account. `shell` - Change password for the shell account.

Parameter	Description	Options
	password.	
reboot	Reboots the appliance.	None
secureuldp	If secureuldp is On, you must manually restart `engine_uldpcollector` after installing or deleting the rootCA or LogLogic LMI certificate: `mtask -s engine_uldpcollector restart`	`create csr` - Creates a certificate signing request. `install rootCA` - Parses and installs the rootCA certificate. `install certificate` - Parses and installs the certificate. `delete rootCA` - Deletes the rootCA certificate. `delete certificate` - Deletes the certificate from the appliance. `show csr` - Displays the certificate signing request.
sshkey_ passphrase	This command controls the sshkey_passphrase feature. Once this feature is enabled, the SSH private key is stored in an encrypted format. The private key can only be used after being unlocked with assigned passphrase every time the system boots up. If the passphrase is not unlocked, any file collection or backup configurations using an SSH-based communication channel. HA is affected and stopped until the passphrase is unlocked.	`enable` - Enables the SSH private key encryption feature. `disable` - Disables the SSH private key encryption feature. The private key is stored in plain text format. `unlock` - Decrypts the encrypted SSH private key and stores the key

Parameter	Description	Options
	Note: The following constraints apply to this feature to work in HA (failover) mode: • The feature cannot be enabled or disabled when HA is configured. • To use the feature in HA mode, the feature must be enabled separately on both nodes in the HA pair. • In HA pair, the unlocked private key is not passed from the active node to standby node. This means that, if one node in the pair is rebooted, it requires manual step to log in to the node and unlock the private key, for HA to work properly.	in the key management daemon. `change_pass` - Assigns a new passphrase to the current SSH private key. `status` - Displays whether sshkey_ passphrase feature is enabled or disabled.
storage_ growth	By default, the feature is disabled. To attach additional storage to LogLogic EVA at the time of system boot, you must enable the feature. The feature remains enabled through every system restart until it is disabled again. For information about how to attach additional storage, see *TIBCO LogLogic® Enterprise Virtual Appliance Quick Start*.	`enable` - Enables attaching more storage volume when the system boots. Once enabled, additional storage is checked at every system boot, and if found, it is attached to the local storage. `disable` - Disables attaching more storage volume when the system boots. `status` - Displays whether the feature is enabled or disabled.
update	Checks and updates files from one version to another version. You can use this command to update files on a smaller scale.	None

The system access command differs from the `system passwd` command. For example, currently the application is password protected. The `system access` command lets you access the application and use the `system passwd` command to change the password for the CLI or system account.

To enable IP tables:

```
> system iptables on
```

To reboot the system:

```
> system reboot
```

To change the console password:

```
> system passwd cli
Enter password:
Re-enter new password:
```

To apply file updates:

```
> system update
Choose an upgrade file from the list:
0: update.tar.bz2
1: exit
>> 0
```

Copying the Public Key to Another Server

Before you begin

For LogLogic LMI 6.2.0 or earlier: Set the permissions of the `~/.ssh/authorized_keys` file to 600 by running the following command:

```
$ chmod 600 ~/.ssh/authorized_keys
```

Procedure

1. In the appliance CLI, copy the public SSH key of the appliance to the server:
    a. Run the system `keycopy` command.

    ```
    > system keycopy
    ```

    The appliance asks whether to test or copy the key.

    b. Enter `C` to copy the key.

    The appliance copies the key to the server and displays its pathname.

    c. Note down the displayed server path where the key is copied.

    You later need to append this file to `~/.ssh/authorized_keys` on the server.
    The appliance asks for the server IP address.

    d. Enter the server IP address (provided by your Administrator).

    The appliance asks for the server user name.

    e. Enter the user name (provided by your Administrator).

    The appliance asks for confirmation of the displayed host IP address and RSA key fingerprint.

    f. Enter `yes`.

    The appliance reports that it permanently added the appliance as a known host, and then asks for the password.

    g. Enter the password.

    The appliance prompts you to configure the server with the appliance's key, appending it to `~/.ssh/authorized_keys` on the server. For example:

    ```
    SCP Server: IP-address
    login as: scpdata
    ============================================================
    Machine Name:  sqalinux
    Owner: SQA Administrator
    Groups: RE/SQA/Documentation
    ```

```
Last Update: Mar 25, 2009
============================================================
SCP_server:~> ls -l /tmp/LOGLOGICPUBKEY
-rw-r--r--    1 scpdata  users         611 2009-12-03 18:07
LOGLOGICPUBKEY
SCP_server:~> cat /tmp/LOGLOGICPUBKEY >> ~/.ssh/authorized_
keys
```

The server setup is complete.

2. Verify the server setup.
   a. Run the system `keycopy` command.

      ```
      > system keycopy
      ```

      The appliance asks whether to test or copy the key.
   b. Enter `T` to test the key.

      The appliance tasks for the server IP address.
   c. Enter the server IP address (provided by your Administrator).

      The appliance asks for the server user name.
   d. Enter the user name (provided by your Administrator).

      The appliance copies a test file (`scptestfile`) to the server and then copies it back to the LogLogic appliance.

      The appliance displays when the test copy is complete successfully.

Applying the File Updates

```
> system update
Choose an upgrade file from the list:
0: update.tar.bz2
1: exit
>> 0
```

unset Command

The `unset` command removes certain configurations from the appliance.

Type the following command from your command line:

```
unset [ data migration | net ]
```

`unset` Syntax Parameters

Parameter	Description
data migration	Removes the data migration configuration for the appliance.
net	Removes a configured network interface from the appliance.

These commands have no arguments or keywords.

Note: Network configuration changes can conflict with cluster configurations, so a cluster must be disabled before you can change network configurations using the `unset net` command.

Example

On an appliance configured with two NICs (eth0 and eth1), to remove eth1 and leave eth0 as the only configured NIC:

```
> unset net
Please select the network interface to unconfigure.
0. 10.1.35.5 eth0
1. 10.1.35.6 eth1
2. Do not unconfigure a network interface. Exit now.
> 1
Network interface eth1 has been designated for unconfiguration.
You must save the change for it to take effect.
> save
```

watch Command

The `watch` command displays the current state of the appliance in real-time.

Example

```
> watch
```

The command shows the following submenu at the bottom of the screen. You can navigate across the different screens using it.

```
1)Overview 2)Queues 3)Forward 4)Cluster 5)HTTP Streams 6)Alerts
7)LLTunnels 8)TCP Dest 9)Upstream b)Backup j)Sched Jobs l,m)Devs
n)Syslog-NG s)Sys Alerts t)Trapsu)Users v)VPN +/-)Change refresh speed
[^C to exit]
```

System Shell Commands

This section describes system shell commands.

llversion Command

The `llversion` command displays versions of all components of LogLogic LMI.

Log in as `toor` user and type the following command:

```
$ llversion [-a]
```

where the option `-a` displays more information such as revision number and the time stamp of the build.

The versions of components on the appliance are displayed. An example output if the option `-a` is not used:

```
LMI Version: 6.3.0
Hotfix Version: N/A
LSP Version: 34.1
LSPC Version: 2.11
Logu Version: 6.3.0.0
llshred Version: 1.0.2
MCAgent Version: 2.4.8
EDC Version: 2.1.0
```

llshred Command

To comply with the GDPR regulations, LogLogic LMI provides a CLI utility - the `llshred` command - to shred selected data.

Log in as `toor` user and type the following command from the command line:

```
$ llshred [ -dryRun | -f | -v | -vv ]
```

`llshred` Syntax Parameters

Parameter	Description	Default
-dryRun	If set to false, log data is prefixed with the string `LLSHREDDER` and each character is replaced with an `X`. If set to true, processes the input file but does not alter events. It is highly recommended to run with `dryRun=true` before running the command with `dryRun = false`.	true
-f	Specifies the path to the file containing the log events to be erased.	n/a
-v	Reports verbose console messages. Provides detailed information about the files that include shredded events.	false
-vv	Reports more verbose console messages. Provides more information about events that are shredded.	false

Command examples

An example with `dryRun` as false

```
$ llshred -f llshred_key.csv -dryRun false

----- Shred Report -----
Command Line        :llshred -f llshred_key.csv -dryRun false
Status              :Success Shredded Events saved.
Version             :version=1.0.1-
SNAPSHOT,timestamp=20180516092641,revision=132636
Appliance IP        :10.114.76.60
Start               :Thursday, June 28, 2018 10:31:48 AM PDT
End                 :Thursday, June 28, 2018 10:31:49 AM PDT
```

```
Duration            :0:00:00.015
Events Shredded     :38
Files Modified      :5
/loglogic/data/vol1/2018/06/28/1700/rawdata_10013_1530205920_60-
3650.txt.gz,25
/loglogic/data/vol1/2018/06/28/1600/rawdata_10030_1530203340_60-
3650.txt.gz,1
/loglogic/data/vol1/2018/06/28/1600/rawdata_10048_1530204420_60-
3650.txt.gz,1
/loglogic/data/vol1/2018/06/28/1600/rawdata_10054_1530204780_60-
3650.txt.gz,10
/loglogic/data/vol1/2018/06/28/1700/rawdata_10002_1530205260_60-
3650.txt.gz,1
```

An example with `dryRun` as true

```
$ llshred -f llshred_key.csv -dryRun true

----- Shred Report -----
Command Line        :llshred -f llshred_key.csv -dryRun true
Status              :Success Dry Run - Shredded Events not saved.
Version             :version=1.0.1-
SNAPSHOT,timestamp=20180516092641,revision=132636
Appliance IP        :10.114.76.60
Start               :Thursday, June 28, 2018 10:29:27 AM PDT
End                 :Thursday, June 28, 2018 10:29:28 AM PDT
Duration            :0:00:00.963
Events Shredded     :38
Files to be Modified:5
/loglogic/data/vol1/2018/06/28/1700/rawdata_10013_1530205920_60-
3650.txt.gz,25
/loglogic/data/vol1/2018/06/28/1600/rawdata_10030_1530203340_60-
3650.txt.gz,1
/loglogic/data/vol1/2018/06/28/1600/rawdata_10048_1530204420_60-
3650.txt.gz,1
/loglogic/data/vol1/2018/06/28/1600/rawdata_10054_1530204780_60-
3650.txt.gz,10
/loglogic/data/vol1/2018/06/28/1700/rawdata_10002_1530205260_60-
3650.txt.gz,1
```

An example with error in the .csv file

```
$ llshred -f llshred_key.csv

----- Shred Report -----
```

```
Command Line        :llshred -f llshred_key.csv
Status              :Invalid EventKey:8Dda28272E0, No stdf record with
gmtHourPath,seqno = /2018/07/28/0800/,10279 found.
Version             :version=1.0.1-
SNAPSHOT,timestamp=20180516092641,revision=132636
Appliance IP        :10.114.76.60
Start               :Thursday, June 28, 2018 10:41:25 AM PDT
End                 :Thursday, June 28, 2018 10:41:25 AM PDT
Duration            :0:00:00.360
Events Shredded     :0
Files to be Modified:0
------------------------
```

Related Topics

For more information, see Shredding LogLogic LMI Event Data.

Simple Network Management Protocol (SNMP)

Simple Network Management Protocol (SNMP) is an application layer protocol that facilitates the exchange of management information between network devices.

SNMP uses UDP/IP protocol stack. The current implementation supports SNMP version v2c.

A typical Network Management Station (NMS) runs one or more management applications in order to control and monitor managed Appliances. A node/Appliance can also be monitored and controlled by more than one NMS. Since SNMP is platform independent, an agent running on an Appliance does not need any extra functionality to support a specific type of operating system or hardware.

An SNMP agent running on the Appliance responds to the SNMP queries. Applications can query both an operating system and LogLogic product details. The following example shows one NMS managing all the LogLogic Appliances.

The SNMP receiver can be any entity capable of receiving SNMP traps v1 or v2c. The following diagram illustrates a single NMS managing all connected LogLogic appliances.

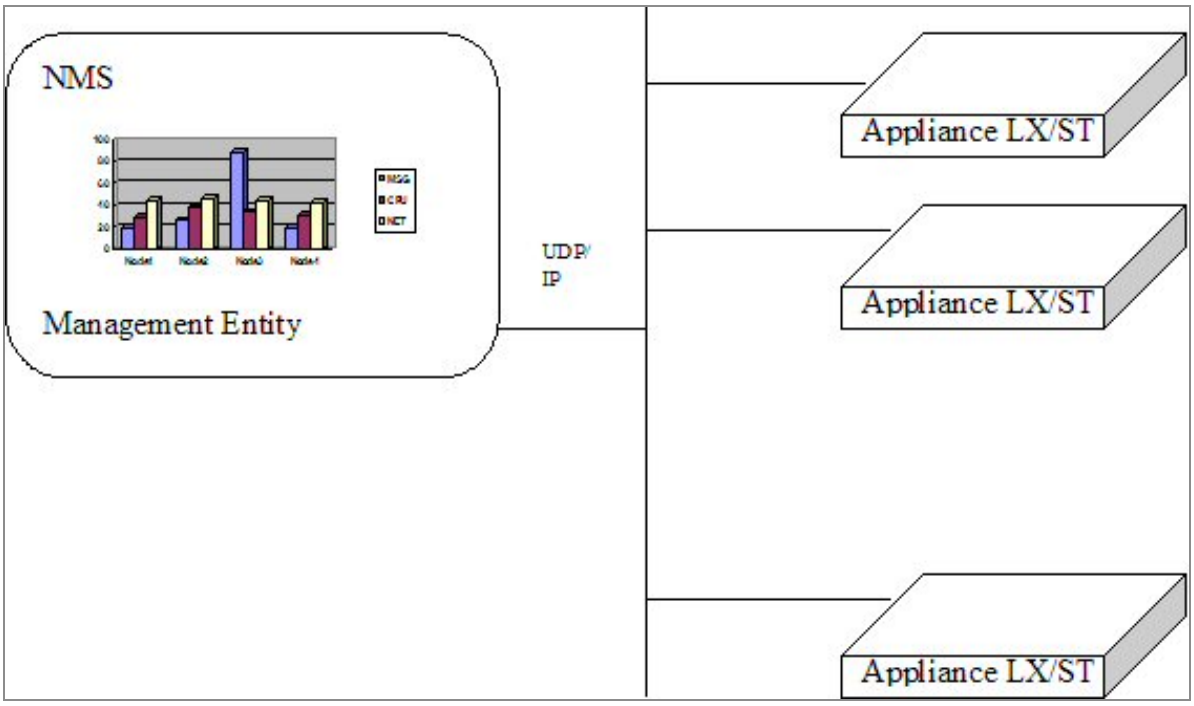

Enabling SNMP

You can enable Simple Network Management Protocol (SNMP) to facilitate the exchange of management information between network devices.

Procedure

1. Log in to your appliance.
2. Select the **Administration > System Settings > General** tab to modify the default SNMP community string.
3. Select the **Enable SNMP Daemon** check box.
4. Click **Update**.
5. Select **Administration > Firewall Settings**.

   The **Firewall Settings** tab is displayed.
6. Select the **Enable IP Firewall** checkbox to activate the fields.
7. Select the **IP Address All** radio button from the **Input Rule** box.
8. Select the UDP protocol from the **Protocol** drop down box.
9. Select SNMP: 161 from the **Port** drop down box.
10. Select Accept from the **Action** radio buttons.
11. Click **Add**.
12. Verify that port 161 is added to the list.
13. Click **Apply** to save your changes.

> **Note:** In IPv6 environments when entering IPv6 addresses ensure that the address is in square brackets and is preceded by udp6: followed by 161 (where 161 is the snmp port number).

Management Information Base

A Management Information Base (MIB) is a collection of information that is organized hierarchically. MIBs are accessed using a network-management protocol such as SNMP.

They are comprised of managed objects and are identified by object identifiers.

The LogLogic MIB is an integrated MIB for your appliances. The LogLogic SNMP agent responds to requests from SNMP managers.

NMS, Agent, and MIB Relationship

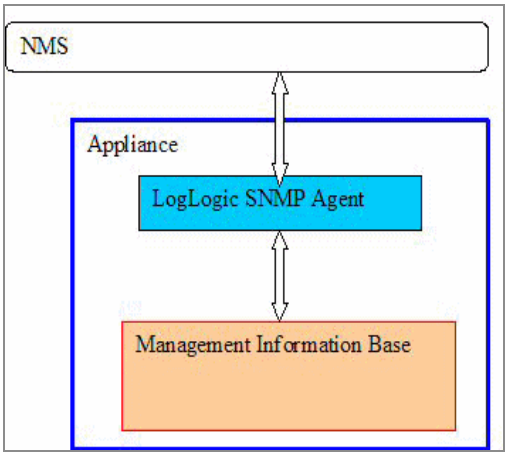

A copy of LOGLOGIC-SNMP-MIB.txt is located under `/usr/share/snmp/mibs` folder or from the TIBCO Support Website, you can download `LOGLOGIC-SNMP-MIB.txt` to load into your NMS. After it is loaded and configured, NMS is able to query your LogLogic appliance. In addition, you should enable SNMP the appliance you want to monitor and control. For details about enabling SNMP, see Enabling SNMP.

SNMP is a request-response protocol used to transfer management information between entities acting in a manager role and entities acting in an agent role. The NMS requests data from the appliance to display it in a user defined form. Snmpwalk (or snmpget) is a command line tool to get data from your appliance.

Sample Object IDs

IPv4 Examples

The following IPv4 examples show a few query and response details on an appliance with a sample IP address of 10.1.1.226 and community string public. The sample OIDs display with their sample output in the code.

Note: If you are using IPv6 addresses, ensure that you follow the correct address format. For more information, see IPv6 Address Formats.

```
$ snmpwalk -v2c -c public 10.1.1.226  SNMPv2-
SMI::enterprises.18552.1.2.1
Prints all the LX MIB.
$ snmpwalk -v2c -c public 10.1.1.226  SNMPv2-
SMI::enterprises.18552.2.2.1
Prints all the ST MIB.
Note that  SNMPv2-SMI::enterprises can be replaced with dotted number
format .1.3.6.1.4.1
$ snmpwalk -v2c -c public 10.1.1.226 1.3.6.1.2
To poll system, interfaces, etc
$ snmpwalk -v2c -c public 10.1.1.226 1.3.6.1.2.1.1.3.0
SNMPv2-MIB::sysUpTime.0 = Timeticks: (25555392) 2 days, 22:59:13.92
To poll for system uptime.
Corresponding system command
$ uptime
 15:32:54 up 2 days, 23:01,  4 users,  load average: 2.43, 2.59, 2.49
$ snmpwalk -v2c 10.1.1.226 -c public 1.3.6.1.2.1.2.2.1
To walk the network interfaces table. bond0, eth0 and eth1 are listed in
this table along with their corresponding network stats
IF-MIB::ifIndex.1 = INTEGER: 1
IF-MIB::ifIndex.2 = INTEGER: 2
IF-MIB::ifIndex.3 = INTEGER: 3
IF-MIB::ifIndex.4 = INTEGER: 4
IF-MIB::ifIndex.5 = INTEGER: 5
IF-MIB::ifIndex.6 = INTEGER: 6
IF-MIB::ifDescr.1 = STRING: lo
IF-MIB::ifDescr.2 = STRING: eth0
IF-MIB::ifDescr.3 = STRING: eth1
IF-MIB::ifDescr.4 = STRING: dummy0
IF-MIB::ifDescr.5 = STRING: eql
IF-MIB::ifDescr.6 = STRING: bond0
IF-MIB::ifType.1 = INTEGER: softwareLoopback(24)
IF-MIB::ifType.2 = INTEGER: ethernetCsmacd(6)
IF-MIB::ifType.3 = INTEGER: ethernetCsmacd(6)
IF-MIB::ifType.4 = INTEGER: ethernetCsmacd(6)
IF-MIB::ifType.5 = INTEGER: slip(28)
IF-MIB::ifType.6 = INTEGER: ethernetCsmacd(6)
IF-MIB::ifMtu.1 = INTEGER: 16436
IF-MIB::ifMtu.2 = INTEGER: 1500
IF-MIB::ifMtu.3 = INTEGER: 1500
IF-MIB::ifMtu.4 = INTEGER: 1500
IF-MIB::ifMtu.5 = INTEGER: 576
```

```
IF-MIB::ifMtu.6 = INTEGER: 1500
IF-MIB::ifSpeed.1 = Gauge32: 10000000
IF-MIB::ifSpeed.2 = Gauge32: 100000000
IF-MIB::ifSpeed.3 = Gauge32: 0
IF-MIB::ifSpeed.4 = Gauge32: 10000000
IF-MIB::ifSpeed.5 = Gauge32: 0
IF-MIB::ifSpeed.6 = Gauge32: 10000000
IF-MIB::ifPhysAddress.1 = STRING:
IF-MIB::ifPhysAddress.2 = STRING: 0:2:b3:e9:33:80
IF-MIB::ifPhysAddress.3 = STRING: 0:2:b3:e9:33:80
IF-MIB::ifPhysAddress.4 = STRING:
IF-MIB::ifPhysAddress.5 = STRING:
IF-MIB::ifPhysAddress.6 = STRING: 0:2:b3:e9:33:80
IF-MIB::ifAdminStatus.1 = INTEGER: up(1)
IF-MIB::ifAdminStatus.2 = INTEGER: up(1)
IF-MIB::ifAdminStatus.3 = INTEGER: up(1)
IF-MIB::ifAdminStatus.4 = INTEGER: down(2)
IF-MIB::ifAdminStatus.5 = INTEGER: down(2)
IF-MIB::ifAdminStatus.6 = INTEGER: up(1)
IF-MIB::ifOperStatus.1 = INTEGER: up(1)
IF-MIB::ifOperStatus.2 = INTEGER: up(1)
IF-MIB::ifOperStatus.3 = INTEGER: down(2)
IF-MIB::ifOperStatus.4 = INTEGER: down(2)
IF-MIB::ifOperStatus.5 = INTEGER: down(2)
IF-MIB::ifOperStatus.6 = INTEGER: up(1)
IF-MIB::ifInOctets.1 = Counter32: 179847582
IF-MIB::ifInOctets.2 = Counter32: 3672236919
IF-MIB::ifInOctets.3 = Counter32: 0
IF-MIB::ifInOctets.4 = Counter32: 0
IF-MIB::ifInOctets.5 = Counter32: 0
IF-MIB::ifInOctets.6 = Counter32: 3672414769
IF-MIB::ifInUcastPkts.1 = Counter32: 1928357
IF-MIB::ifInUcastPkts.2 = Counter32: 1353515244
IF-MIB::ifInUcastPkts.3 = Counter32: 0
IF-MIB::ifInUcastPkts.4 = Counter32: 0
IF-MIB::ifInUcastPkts.5 = Counter32: 0
IF-MIB::ifInUcastPkts.6 = Counter32: 1353515828
IF-MIB::ifInDiscards.1 = Counter32: 0
IF-MIB::ifInDiscards.2 = Counter32: 44
IF-MIB::ifInDiscards.3 = Counter32: 0
IF-MIB::ifInDiscards.4 = Counter32: 0
IF-MIB::ifInDiscards.5 = Counter32: 0
IF-MIB::ifInDiscards.6 = Counter32: 44
IF-MIB::ifInErrors.1 = Counter32: 0
IF-MIB::ifInErrors.2 = Counter32: 44
IF-MIB::ifInErrors.3 = Counter32: 0
```

```
IF-MIB::ifInErrors.4 = Counter32: 0
IF-MIB::ifInErrors.5 = Counter32: 0
IF-MIB::ifInErrors.6 = Counter32: 44
IF-MIB::ifOutOctets.1 = Counter32: 179847582
IF-MIB::ifOutOctets.2 = Counter32: 547984552
IF-MIB::ifOutOctets.3 = Counter32: 0
IF-MIB::ifOutOctets.4 = Counter32: 0
IF-MIB::ifOutOctets.5 = Counter32: 0
IF-MIB::ifOutOctets.6 = Counter32: 547984923
IF-MIB::ifOutUcastPkts.1 = Counter32: 1928357
IF-MIB::ifOutUcastPkts.2 = Counter32: 947178
IF-MIB::ifOutUcastPkts.3 = Counter32: 0
IF-MIB::ifOutUcastPkts.4 = Counter32: 0
IF-MIB::ifOutUcastPkts.5 = Counter32: 0
IF-MIB::ifOutUcastPkts.6 = Counter32: 947182
IF-MIB::ifOutDiscards.1 = Counter32: 0
IF-MIB::ifOutDiscards.2 = Counter32: 0
IF-MIB::ifOutDiscards.3 = Counter32: 0
IF-MIB::ifOutDiscards.4 = Counter32: 0
IF-MIB::ifOutDiscards.5 = Counter32: 0
IF-MIB::ifOutDiscards.6 = Counter32: 0
IF-MIB::ifOutErrors.1 = Counter32: 0
IF-MIB::ifOutErrors.2 = Counter32: 0
IF-MIB::ifOutErrors.3 = Counter32: 0
IF-MIB::ifOutErrors.4 = Counter32: 0
IF-MIB::ifOutErrors.5 = Counter32: 0
IF-MIB::ifOutErrors.6 = Counter32: 0
IF-MIB::ifOutQLen.1 = Gauge32: 0
IF-MIB::ifOutQLen.2 = Gauge32: 0
IF-MIB::ifOutQLen.3 = Gauge32: 0
IF-MIB::ifOutQLen.4 = Gauge32: 0
IF-MIB::ifOutQLen.5 = Gauge32: 0
IF-MIB::ifOutQLen.6 = Gauge32: 0
IF-MIB::ifSpecific.1 = OID: SNMPv2-SMI::zeroDotZero
IF-MIB::ifSpecific.2 = OID: SNMPv2-SMI::zeroDotZero
IF-MIB::ifSpecific.3 = OID: SNMPv2-SMI::zeroDotZero
IF-MIB::ifSpecific.4 = OID: SNMPv2-SMI::zeroDotZero
IF-MIB::ifSpecific.5 = OID: SNMPv2-SMI::zeroDotZero
IF-MIB::ifSpecific.6 = OID: SNMPv2-SMI::zeroDotZero
System command output:
$ ifconfig
bond0     Link encap:Ethernet  HWaddr 00:02:B3:E9:33:80
          inet addr:10.1.1.226 Bcast:10.1.1.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MASTER MULTICAST  MTU:1500  Metric:1
          RX packets:1404595735 errors:44 dropped:44 overruns:0 frame:0
          TX packets:1064539 errors:0 dropped:0 overruns:0 carrier:0
```

```
          collisions:0 txqueuelen:0
          RX bytes:1629444858 (1553.9 Mb)  TX bytes:570661638 (544.2 Mb)
eth0      Link encap:Ethernet  HWaddr 00:02:B3:E9:33:80
          UP BROADCAST RUNNING SLAVE MULTICAST  MTU:1500  Metric:1
          RX packets:1404595720 errors:44 dropped:44 overruns:0 frame:0
          TX packets:1064539 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:1629441929 (1553.9 Mb)  TX bytes:570661638 (544.2 Mb)
          Base address:0x8440 Memory:fe020000-fe040000
eth1      Link encap:Ethernet  HWaddr 00:02:B3:E9:33:80
          UP BROADCAST NOARP SLAVE MULTICAST  MTU:1500  Metric:1
          RX packets:0 errors:0 dropped:0 overruns:0 frame:0
          TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)
lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          UP LOOPBACK RUNNING  MTU:16436  Metric:1
          RX packets:1947784 errors:0 dropped:0 overruns:0 frame:0
          TX packets:1947784 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:181891708 (173.4 Mb)  TX bytes:181891708 (173.4 Mb)
$ snmpwalk -v2c 10.1.1.226 -c public 1.3.6.1.4.1.2021.10.1
System load averages.
UCD-SNMP-MIB::laIndex.1 = INTEGER: 1
UCD-SNMP-MIB::laIndex.2 = INTEGER: 2
UCD-SNMP-MIB::laIndex.3 = INTEGER: 3
UCD-SNMP-MIB::laNames.1 = STRING: Load-1
UCD-SNMP-MIB::laNames.2 = STRING: Load-5
UCD-SNMP-MIB::laNames.3 = STRING: Load-15
UCD-SNMP-MIB::laLoad.1 = STRING: 2.69
UCD-SNMP-MIB::laLoad.2 = STRING: 2.61
UCD-SNMP-MIB::laLoad.3 = STRING: 2.44
UCD-SNMP-MIB::laConfig.1 = STRING: 12.00
UCD-SNMP-MIB::laConfig.2 = STRING: 12.00
UCD-SNMP-MIB::laConfig.3 = STRING: 12.00
UCD-SNMP-MIB::laLoadInt.1 = INTEGER: 268
UCD-SNMP-MIB::laLoadInt.2 = INTEGER: 260
UCD-SNMP-MIB::laLoadInt.3 = INTEGER: 243
UCD-SNMP-MIB::laLoadFloat.1 = Opaque: Float: 2.690000
UCD-SNMP-MIB::laLoadFloat.2 = Opaque: Float: 2.610000
UCD-SNMP-MIB::laLoadFloat.3 = Opaque: Float: 2.440000
UCD-SNMP-MIB::laErrorFlag.1 = INTEGER: 0
UCD-SNMP-MIB::laErrorFlag.2 = INTEGER: 0
UCD-SNMP-MIB::laErrorFlag.3 = INTEGER: 0
UCD-SNMP-MIB::laErrMessage.1 = STRING:
```

```
UCD-SNMP-MIB::laErrMessage.2 = STRING:
UCD-SNMP-MIB::laErrMessage.3 = STRING:
$ snmpwalk -v2c 10.1.1.226 -c public 1.3.6.1.4.1.2021.11
CPU usage.
snmpwalk -v2c 10.1.1.226 -c public 1.3.6.1.4.1.2021.11
UCD-SNMP-MIB::ssIndex.0 = INTEGER: 1
UCD-SNMP-MIB::ssErrorName.0 = STRING: systemStats
UCD-SNMP-MIB::ssSwapIn.0 = INTEGER: 2
UCD-SNMP-MIB::ssSwapOut.0 = INTEGER: 1
UCD-SNMP-MIB::ssIOSent.0 = INTEGER: 34
UCD-SNMP-MIB::ssIOReceive.0 = INTEGER: 2
UCD-SNMP-MIB::ssSysInterrupts.0 = INTEGER: 2
UCD-SNMP-MIB::ssSysContext.0 = INTEGER: 19
UCD-SNMP-MIB::ssCpuUser.0 = INTEGER: 13
UCD-SNMP-MIB::ssCpuSystem.0 = INTEGER: 22
UCD-SNMP-MIB::ssCpuIdle.0 = INTEGER: 63
UCD-SNMP-MIB::ssCpuRawUser.0 = Counter32: 7283712
UCD-SNMP-MIB::ssCpuRawNice.0 = Counter32: 6917114
UCD-SNMP-MIB::ssCpuRawSystem.0 = Counter32: 22873470
UCD-SNMP-MIB::ssCpuRawIdle.0 = Counter32: 65327096
UCD-SNMP-MIB::ssCpuRawKernel.0 = Counter32: 22873470
UCD-SNMP-MIB::ssIORawSent.0 = Counter32: 1529174198
UCD-SNMP-MIB::ssIORawReceived.0 = Counter32: 2322564824
UCD-SNMP-MIB::ssRawInterrupts.0 = Counter32: 1805767855
UCD-SNMP-MIB::ssRawContexts.0 = Counter32: 3197292578
UCD-SNMP-MIB::systemStats.62.0 = Counter32: 567530
UCD-SNMP-MIB::systemStats.63.0 = Counter32: 349201
System command output:
$ vmstat 1 10
procs -----------memory---------- ---swap-- -----io---- --system-- ----
cpu----
 r  b   swpd   free   buff  cache   si   so    bi    bo   in    cs us sy
id wa
 1  1 204464  49616  77736 1763564    2    1    34     2    2    19 14
22 64  0
 4  2 204464  46388  77744 1764116    0    0     0 30240 7020 11827  7
30 63  0
 2  0 204464  49304  77760 1762952 1216    0  1216 10880 6806 10088  8
48 44  0
 1  0 204464  49264  77772 1763136    0    0     0     0 6544 11681  6
21 74  0
 2  0 204464  45792  77592 1763304    0    0     0     0 6441 11931  8
22 71  0
 1  0 204464  48384  77736 1763328    0    0     0   284 6478  8893  8
18 75  0
 3  0 204464  48448  77748 1763288    0    0     0     0 6502 11732  6
```

```
21 73  0
 1  0 204464  48864  77760 1763460    0    0     0     0 6478 11506  7
18 75  0
 2  1 204464  49416  77772 1761916    0    0     0 16916 6580 11303  8
19 73  0
 1  1 204464  49468  77664 1763008    0    0     0 32784 6675 10961  9
18 73  0
$ snmpwalk -v2c 10.1.1.226 -c public 1.3.6.1.4.1.2021.9
Disk stats.
UCD-SNMP-MIB::dskIndex.1 = INTEGER: 1
UCD-SNMP-MIB::dskIndex.2 = INTEGER: 2
UCD-SNMP-MIB::dskIndex.3 = INTEGER: 3
UCD-SNMP-MIB::dskIndex.4 = INTEGER: 4
UCD-SNMP-MIB::dskPath.1 = STRING: /
UCD-SNMP-MIB::dskPath.2 = STRING: /failsafe
UCD-SNMP-MIB::dskPath.3 = STRING: /tmp
UCD-SNMP-MIB::dskPath.4 = STRING: /loglogic
UCD-SNMP-MIB::dskDevice.1 = STRING:
/dev/scsi/host0/bus0/target0/lun0/part1
UCD-SNMP-MIB::dskDevice.2 = STRING:
/dev/scsi/host0/bus0/target0/lun0/part2
UCD-SNMP-MIB::dskDevice.3 = STRING:
/dev/scsi/host0/bus0/target0/lun0/part5
UCD-SNMP-MIB::dskDevice.4 = STRING:
/dev/scsi/host0/bus0/target2/lun0/part1
UCD-SNMP-MIB::dskMinimum.1 = INTEGER: 10000
UCD-SNMP-MIB::dskMinimum.2 = INTEGER: 10000
UCD-SNMP-MIB::dskMinimum.3 = INTEGER: 100000
UCD-SNMP-MIB::dskMinimum.4 = INTEGER: 100000
UCD-SNMP-MIB::dskMinPercent.1 = INTEGER: -1
UCD-SNMP-MIB::dskMinPercent.2 = INTEGER: -1
UCD-SNMP-MIB::dskMinPercent.3 = INTEGER: -1
UCD-SNMP-MIB::dskMinPercent.4 = INTEGER: -1
UCD-SNMP-MIB::dskTotal.1 = INTEGER: 1494236
UCD-SNMP-MIB::dskTotal.2 = INTEGER: 1035692
UCD-SNMP-MIB::dskTotal.3 = INTEGER: 1035692
UCD-SNMP-MIB::dskTotal.4 = INTEGER: 1960866168
UCD-SNMP-MIB::dskAvail.1 = INTEGER: 1188804
UCD-SNMP-MIB::dskAvail.2 = INTEGER: 554520
UCD-SNMP-MIB::dskAvail.3 = INTEGER: 948528
UCD-SNMP-MIB::dskAvail.4 = INTEGER: 1552997244
UCD-SNMP-MIB::dskUsed.1 = INTEGER: 229528
UCD-SNMP-MIB::dskUsed.2 = INTEGER: 428560
UCD-SNMP-MIB::dskUsed.3 = INTEGER: 34552
UCD-SNMP-MIB::dskUsed.4 = INTEGER: 407868924
UCD-SNMP-MIB::dskPercent.1 = INTEGER: 16
```

```
UCD-SNMP-MIB::dskPercent.2 = INTEGER: 44
UCD-SNMP-MIB::dskPercent.3 = INTEGER: 4
UCD-SNMP-MIB::dskPercent.4 = INTEGER: 21
UCD-SNMP-MIB::dskPercentNode.1 = INTEGER: 4
UCD-SNMP-MIB::dskPercentNode.2 = INTEGER: 8
UCD-SNMP-MIB::dskPercentNode.3 = INTEGER: 0
UCD-SNMP-MIB::dskPercentNode.4 = INTEGER: 0
UCD-SNMP-MIB::dskErrorFlag.1 = INTEGER: 0
UCD-SNMP-MIB::dskErrorFlag.2 = INTEGER: 0
UCD-SNMP-MIB::dskErrorFlag.3 = INTEGER: 0
UCD-SNMP-MIB::dskErrorFlag.4 = INTEGER: 0
UCD-SNMP-MIB::dskErrorMsg.1 = STRING:
UCD-SNMP-MIB::dskErrorMsg.2 = STRING:
UCD-SNMP-MIB::dskErrorMsg.3 = STRING:
UCD-SNMP-MIB::dskErrorMsg.4 = STRING:
System command output:
$ df
Filesystem           1K-blocks      Used Available Use% Mounted on
/dev/scsi/host0/bus0/target0/lun0/part1
                       1494236    229536   1188796  17% /
/dev/scsi/host0/bus0/target0/lun0/part2
                       1035692    428560    554520  44% /failsafe
/dev/scsi/host0/bus0/target0/lun0/part5
                       1035692     34552    948528   4% /tmp
/dev/scsi/host0/bus0/target2/lun0/part1
                     1960866168 407891296 1552974872  21% /loglogic
$ snmpwalk -v2c 10.1.1.226 -c public 1.3.6.1.4.1.2021.4
Memory stats.
UCD-SNMP-MIB::memIndex.0 = INTEGER: 0
UCD-SNMP-MIB::memErrorName.0 = STRING: swap
UCD-SNMP-MIB::memTotalSwap.0 = INTEGER: 1052248
UCD-SNMP-MIB::memAvailSwap.0 = INTEGER: 847824
UCD-SNMP-MIB::memTotalReal.0 = INTEGER: 2067508
UCD-SNMP-MIB::memAvailReal.0 = INTEGER: 49424
UCD-SNMP-MIB::memTotalFree.0 = INTEGER: 897252
UCD-SNMP-MIB::memMinimumSwap.0 = INTEGER: 16000
UCD-SNMP-MIB::memShared.0 = INTEGER: 0
UCD-SNMP-MIB::memBuffer.0 = INTEGER: 79240
UCD-SNMP-MIB::memCached.0 = INTEGER: 1765480
UCD-SNMP-MIB::memSwapError.0 = INTEGER: 0
UCD-SNMP-MIB::memSwapErrorMsg.0 = STRING:
System command output:
$ free
             total       used       free     shared    buffers
cached
Mem:       2067508    1937788     129720          0      79360
```

```
1680940
-/+ buffers/cache:     177488    1890020
Swap:      1052248     204424     847824
```

IPv6 Examples

The following IPv6 examples show a few query and response details on an appliance with a sample IP address of fd00::aaaa:a72:4a2b and community string public.

```
$snmpwalk -v2c udp6:[fd00::aaaa:a72:4a2b]:161 -c public .1.3.6.1.2.1.1.3.0
```

Supported Object IDs

There are several other supported Object IDs (OID) on LogLogic appliances.

To view them, access them from top node 1 or provide the OID explicitly to access data. You can also contact support@tibco.com if you cannot find a particular OID.

For every symbolic name, there are two OIDs: one for LogLogic LX Appliance or LogLogic MX Appliance, and another for LogLogic ST Appliance.

Object ID (OID)	Symbolic Name	Definitions
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.10 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.10	macAddr	System etho interface media access control address
For LX or MX OID:SNMPv2-SMI::enterprises.18552.1.2.1.11 For ST OID:SNMPv2-SMI::enterprises.18552.2.2.1.11	productSelected	Name of the LogLogic appliance product family (LogLogic MX Appliance, LogLogic LX Appliance, or LogLogic MX Appliance).

Object ID (OID)	Symbolic Name	Definitions
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.12 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.12	modelSelected	Model of the appliance.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.13 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.13	messageQueueInserts	Slot number that the most recent message was placed. This number indicate that the appliance received this many number of messages so far from boot time.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.14 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.14	messageQueueReads	Slot number that the parser refers to read next message from the received message queue.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.15 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.15	messageDrops	Number of messages that were dropped. This number indicate that the appliance is losing messages because of the lack of resources.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.16 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.16	msgRatePerMin	Number of messages received during the last minute.
For LX or MX OID: SNMPv2-	msgRatePer5Min	Number of

Object ID (OID)	Symbolic Name	Definitions
SMI::enterprises.18552.1.2.1.17 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.17		messages received during the last five minutes.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.18 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.18	msgRatePer15Min	Number of messages received during the last fifteen minutes.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.19 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.19	totalSyslogSources	Number of syslog log sources that the appliance is receiving messages from.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.20 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.20	totalSyslogReceivers	Total number of active syslog receivers in the list.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.21 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.21	totalTrapReceivers	Appliance is able to receive and forward SNMP traps. This is a counter that represents the total number of trap senders to the appliance.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.22 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.22	totalLEAServers	Total number of Check Point LEA servers.
For LX or MX OID: SNMPv2-	startLEAAgent	ID of the LEA agent

Object ID (OID)	Symbolic Name	Definitions
SMI::enterprises.18552.1.2.1.23 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.23		to start.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.24 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.24	totalUsers	Total number of users on the appliance.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.25 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.25	totalSyslogMessages	Number of log messages received by the appliance from the last boot.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.26 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.26	totalLEAMessages	Number of log messages received through LEA server but deprecated in 32.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.27 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.27	totalProcessedMessages	Total of all the processed syslog message counters.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.28 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.28	totalUnapprovedLEAMessages	Number of log messages received through LEA server but not approved bur deprecated in 3.2 version
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.29	totalUnknownSyslogMessages	Total number of messages that are not recognized by

Object ID (OID)	Symbolic Name	Definitions
For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.29		the appliance.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.30 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.30	totalOtherSyslogMessages	Total messages difference between total messages received by the appliance minus all known and unknown messages.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.31 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.31	totalSkippedSyslogMessages	Total number of skipped syslog messages. Some messages needs to be skipped not to count twice such as SNMP trap messages or device may be disabled.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.32 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.32	totalApprovedSyslogMessages	Total number of approved messages. The message is approved if the used enable the device or auto discover is turned on. Deprecated in LogLogic Release 3.2.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.33 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.33	totalUnapprovedSyslogMessages	Messages come into the appliance and counted under unapproved as long as the

Object ID (OID)	Symbolic Name	Definitions
		corresponding device is in the approved list.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.34 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.34	totalAcceptedSyslogMessages	Total number of messages accepted by the appliance in firewall category.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.35 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.35	totalDeniedSyslogMessages	Total number of messages denied by the appliance in firewall category.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.36 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.36	totalSecuritySyslogMessages	Total number of messages categorized as firewall security messages received.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.37 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.37	totalSystemSyslogMessages	Total number of syslog messages received by the appliance in firewall category.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.38 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.38	totalFTPSyslogMessages	Number of messages received through FTP protocol in firewall category.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.39 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.39	totalURLSyslogMessages	Total number of URL messages received in firewall category.

Object ID (OID)	Symbolic Name	Definitions
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.40 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.40	totalNortelVPNAuthMessages	Total number of Nortel Authentication messages received. They are login success and failed messages only.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.41 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.41	totalVPNMessages	Total number of VPN messages received.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.42 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.42	dbQueueInserts	Slot number where the most recent SQL query added to the queue.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.43 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.43	dbQueueReads	Slot number where the most recent SQL query to be executed.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.44 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.44	seq	Sequence number to assign to each message.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.45 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.45	rsenderQueueInUse	Spin lock semaphore flag for rsender queue.
For LX or MX OID: SNMPv2-	rsenderQueueInserts	Slot number for the

Object ID (OID)	Symbolic Name	Definitions
SMI::enterprises.18552.1.2.1.46 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.46		next message to be inserted for rsender.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.47 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.47	rsenderQueueReads	Slot number from which the next message to be read by rsender.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.48 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.48	rsenderDrops	Number of message that were dropped by rsender.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.49 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.49	startTime	Time when the engines shared memory segment was created.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.50 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.50	messageTooLong	Number of messages that were unable to fit into buffer slot.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.51 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.51	msgRatePerSec	Number of messages received during the last second.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.60 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.60	CPUFanSpeed	Variable set for a group of data. Each group is in sensor number, fan speed RPM, high

Object ID (OID)	Symbolic Name	Definitions
		threshold, and low threshold form.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.61 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.61	SysFanSpeed	Variable set for a group of data. Each group is in sensor number, fan speed RPM, high threshold, and low threshold form.
For LX or MX OID: SNMPv2-SMI::enterprises.18552.1.2.1.62 For ST OID: SNMPv2-SMI::enterprises.18552.2.2.1.62	CPUTemperature	Variable set for a group of data. Each group is in sensor number, Celsius degrees, high threshold, and low threshold form.
For LX: SNMPv2-SMI::enterprises.18552.1.2.1.63 For ST: SNMPv2-SMI::enterprises.18552.2.2.1.63	CpuRatePerMin	CPU rate percentage, 1 minute average
For LX: SNMPv2-SMI::enterprises.18552.1.2.1.64 For ST: SNMPv2-SMI::enterprises.18552.2.2.1.64	CpuRatePer5Min	CPU rate percentage, 5 minutes average
For LX: SNMPv2-SMI::enterprises.18552.1.2.1.65 For ST: SNMPv2-SMI::enterprises.18552.2.2.1.65	CpuRatePer15Min	CPU rate percentage, 15 minutes average:
For LX: SNMPv2-SMI::enterprises.18552.1.2.1.66	DiskUsageMB	Disk used space (megabytes)

Object ID (OID)	Symbolic Name	Definitions
For ST: SNMPv2-SMI::enterprises.18552.2.2.1.66		
For LX: SNMPv2-SMI::enterprises.18552.1.2.1.67 For ST: SNMPv2-SMI::enterprises.18552.2.2.1.67	Version	LogLogic LMI version
For LX: SNMPv2-SMI::enterprises.18552.1.2.1.68 For ST: SNMPv2-SMI::enterprises.18552.2.2.1.68	HotfixVersion	Maintenance Hotfix version, for example, HF1 or HF2
For LX: SNMPv2-SMI::enterprises.18552.1.2.1.69 For ST: SNMPv2-SMI::enterprises.18552.2.2.1.69	LSPVersion	LogLogic LSP version
For LX: SNMPv2-SMI::enterprises.18552.1.2.1.70 For ST: SNMPv2-SMI::enterprises.18552.2.2.1.70	DiskFreeMB	Disk free space (megabytes)
For LX: SNMPv2-SMI::enterprises.18552.1.2.1.71 For ST: SNMPv2-SMI::enterprises.18552.2.2.1.71	MemoryFreeMB	Memory available (megabytes)
For LX: SNMPv2-SMI::enterprises.18552.1.2.1.72 For ST: SNMPv2-SMI::enterprises.18552.2.2.1.72	MemoryUsedMB	Memory in use (megabytes)

List of Available Traps

An SNMP trap is an asynchronous event-generated message that an appliance sends to its client.

The client is a trap receiver which is normally a network monitoring station. The LogLogic appliance supports a set of SNMP traps sent to alert a user.

For SNMP version 1, use the ESTN in the following table.

Categories of Alerts Available as SNMP Traps running version 1

Alert Category	Alert Name	Enterprise-OID	ESTN	Description
Cisco PIX Messages Alert	Cisco PIX Messages Alert	ent.18552.1.3 ent.18552.2.3 (ST)	1	The messages per second rate for a specific PIX message code is outside (later or earlier) specified rates
Network Policy Alert	Network Policy Alert	ent.18552.1.3 ent.18552.2.3 (ST)	92	A network policy messages was received with an Accept or Deny Policy action
Adaptive Baseline Alert	Adaptive Baseline Alert	ent.18552.1.3 ent.18552.2.3 (ST)	93	The messages per second rate is outside (later or earlier) the nominal traffic rate
Ratio Based Alert	Ratio Based Alert	ent.18552.1.3 ent.18552.2.3 (ST)	94	The specified message count is outside a specified percentage of total messages
VPN Statistics Alert	VPN Statistics Alert	ent.18552.1.3 ent.18552.2.3 (ST)	95	Recorded statistics on VPN or RADIUS messages match relative or absolute criteria

Alert Category	Alert Name	Enterprise-OID	ESTN	Description
VPN Connections Alert	VPN Connections Alert	ent.18552.1.3 ent.18552.2.3 (ST)	96	A VPN Disconnects or Deny event occurred
VPN Messages Alert	VPN Messages Alert	ent.18552.1.3 ent.18552.2.3 (ST)	97	A VPN Message Alert triggered on combinations of specific VPN message area, severity, and code
Message Volume Alert	Message Volume Alert	ent.18552.1.3 ent.18552.2.3 (ST)	98	The messages per second rate is outside (later or earlier) specified limits. If the user sets the “Zero Message Alert” check box, an alert is triggered only if zero messages are received within the timespan set.
Pre-defined Search Filter	Pre-defined Search Filter Alert	ent.18552.1.3 ent.18552.2.3 (ST)	99	A text search filter matched message fields
System Alert	Dropped- message	ent.18552.1.3 ent.18552.2.3 (ST)	189	Dropped messages exceeded the user-specified limit
	Failover	ent.18552.1.3 ent.18552.2.3 (ST)	191	A failover occurred
	Disk Usage	ent.18552.1.3	192	Disk usage exceeded the specified threshold

Alert Category	Alert Name	Enterprise-OID	ESTN	Description
		ent.18552.2.3 (ST)		
	Network Connection Speed	ent.18552.1.3 ent.18552.2.3 (ST)	193	Network connection throughput fell under the specified threshold
	Network Interface	ent.18552.1.3 ent.18552.2.3 (ST)	194	The configured network interface failed
	Data Migration	ent.18552.1.3 ent.18552.2.3 (ST)	195	A data migration completed
	CPU temperature	ent.18552.1.3 ent.18552.2.3 (ST)	197	The CPU temperature exceeded the specified limit
	Synchronization Failure	ent.18552.1.3 ent.18552.2.3 (ST)	198	Data synchronization failed after a failover
	Secure Tunnel connection status	ent.18552.1.3 ent.18552.2.3 (ST)	199	The configured TCP forward connection failed

If you are using SNMP version 2, the TrapOID differs for each trap. The following table lists the TrapOID for each alert category.

Note: The TIBCO LogLogic appliance uses SNMP v1 by default. In order to use v2c the admin needs to edit the following file:

```
Logapp root:~$ cat /loglogic/conf/snmpd_alerts_trap_version
  version: 1 type: snmptrap  Logapp root:~$
```

Update the version from 1 to 2 in the file. If this file does not exist, create the file with:

```
version: 1 type: snmptrap
```

The file can then be edited to version 2 if required.

Categories of Alerts Available as SNMP Traps running version 2

Alert Category	Alert Name	TrapOID	Description
Cisco PIX Messages Alert	Cisco PIX Messages Alert	ent.18552.1.3.3 ent.18552.2.3.3 (ST)	The messages per second rate for a specific PIX message code is outside (later or earlier) specified rates
Network Policy Alert	Network Policy Alert	ent.18552.1.3.92 ent.18552.2.3.92 (ST)	A network policy messages was received with an Accept or Deny Policy action
Adaptive Baseline Alert	Adaptive Baseline Alert	ent.18552.1.3.93 ent.18552.2.3.93 (ST)	The messages per second rate is outside (later or earlier) the nominal traffic rate
Ratio Based Alert	Ratio Based Alert	ent.18552.1.3.94 ent.18552.2.3.94 (ST)	The specified message count is outside a specified percentage of total messages
VPN Statistics Alert	VPN Statistics Alert	ent.18552.1.3.95 ent.18552.2.3.95	Recorded statistics on VPN or RADIUS messages match relative or absolute criteria

Alert Category	Alert Name	TrapOID	Description
		(ST)	
VPN Connections Alert	VPN Connections Alert	ent.18552.1.3.96 ent.18552.2.3.96 (ST)	A VPN Disconnects or Deny event occurred
VPN Messages Alert	VPN Messages Alert	ent.18552.1.3.97 ent.18552.2.3.97 (ST)	A VPN Message Alert triggered on combinations of specific VPN message area, severity, and code
Message Volume Alert	Message Volume Alert	ent.18552.1.3.98 ent.18552.2.3.98 (ST)	The messages per second rate is outside (later or earlier) specified limits. If the user sets the “Zero Message Alert” check box, an alert is triggered only if zero messages are received within the timespan set.
Pre-defined Search Filter	Pre-defined Search Filter Alert	ent.18552.1.3.99 ent.18552.2.3.99 (ST)	A text search filter matched message fields
System Alert	Dropped-message	ent.18552.1.3.189 ent.18552.2.3.189 (ST)	Dropped messages exceeded the user-specified limit
	Failover	ent.18552.1.3.191 ent.18552.2.3.191 (ST)	A failover occurred
	Disk Usage	ent.18552.1.3.192 ent.18552.2.3.192 (ST)	Disk usage exceeded the specified threshold

Alert Category	Alert Name	TrapOID	Description
	Network Connection Speed	ent.18552.1.3.193 ent.18552.2.3.193 (ST)	Network connection throughput fell under the specified threshold
	Network Interface	ent.18552.1.3.194 ent.18552.2.3.194 (ST)	The configured network interface failed
	Data Migration	ent.18552.1.3.195 ent.18552.2.3.195 (ST)	A data migration completed
	CPU temperature	ent.18552.1.3.197 ent.18552.2.3.197 (ST)	The CPU temperature exceeded the specified limit
	Synchronization Failure	ent.18552.1.3.198 ent.18552.2.3.198 (ST)	Data synchronization failed after a failover
	Secure Tunnel connection status	ent.18552.1.3.199 ent.18552.2.3.199 (ST)	The configured TCP forward connection failed

Other Traps

The SNMP agent can send other application-specific traps.

Besides the LogLogic application-specific traps, the SNMP agent running on your appliance also sends the following traps:

- `SNMPv2-MIB::coldStart`
- `SNMPv2-MIB::authenticationFailure`

- NET-SNMP-AGENT-MIB::nsNotifyShutdown
- NET-SNMP-AGENT-MIB::nsNotifyRestart

These traps, and the following sample messages, are normal messages received in the specified situations.

Sample Messages Received at a Receiver

When a system where SNMP is enabled starts up:

```
2005-08-12 00:57:24 10.1.1.226 [10.1.1.226]:
SNMPv2-MIB::sysUpTime.0 = Timeticks: (9) 0:00:00.09     SNMPv2-
MIB::snmpTrapOID.0 = OID: SNMPv2-MIB::coldStart SNMPv2-
MIB::snmpTrapEnterprise.0 = OID: NET-SNMP-MIB::netSnmpAgentOIDs.10
```

When a system where SNMP is enabled stops:

```
2005-08-12 00:57:21 10.1.1.226 [10.1.1.226]:
NMPv2-MIB::snmpTrapEnterprise.0 = OID: NET-SNMP-MIB::netSnmpAgentOIDs.10
2005-08-12 00:57:21 10.1.1.226 [10.1.1.226]:SNMPv2-MIB::sysUpTime.0 =
Timeticks: (2938) 0:00:29.38  SNMPv2-MIB::snmpTrapOID.0 = OID: NET-SNMP-
AGENT-MIB::nsNotifyShutdown
```

When the agent receives a SIGHUP signal:

```
2005-08-12 00:59:49 10.1.1.226 [10.1.1.226]:
SNMPv2-MIB::sysUpTime.0 = Timeticks: (14462) 0:02:24.62 SNMPv2-
MIB::snmpTrapOID.0 = OID: NET-SNMP-AGENT-MIB::nsNotifyRestart
```

When the agent receives a request using an unknown community name:

```
SNMPv2-MIB::sysUpTime.0 = Timeticks: (147115816) 17     SNMPv2-
MIB::snmpTrapOID.0 = OID: SNMPv2-MIB::authenticationFailureSNMPv2-
MIB::snmpTrapEnterprise.0 = OID: NET-SNMP-MIB::netSnmpAgentOIDs.10
```

Trap Attributes for LogLogic LX Appliances or LogLogic MX Appliances

The alerts on LogLogic LX Appliances or LogLogic MX Appliances generate SNMP traps with these VarBind parameters:

```
{lxTrapDesc, lxLocalAddress, lxTime, lxTrapSourceIP, lxUserOid,
lxTrapMsg}
```

The community string for these trap attributes is public.

Parameter	OID Description
lxTrapDesc	OID: SNMPv2-SMI::enterprises.18552.1.2.1.3
lxLocalAddress	OID: SNMPv2-SMI::enterprises.18552.1.2.1.4
lxTime	OID: SNMPv2-SMI::enterprises.18552.1.2.1.5
lxTrapSourceIP	OID: SNMPv2-SMI::enterprises.18552.1.2.1.6
lxUserOid	OID: SNMPv2-SMI::enterprises.18552.1.2.1.7
lxTrapMsg	OID: SNMPv2-SMI::enterprises.18552.1.2.1.8

Trap Attributes for a LogLogic ST Appliance

The alerts on LogLogic ST Appliances generate SNMP traps with the following VarBind parameters:

```
{stTrapDesc, stLocalAddress, stTime, stTrapSourceIP, stUserOid,
stTrapMsg}
```

The community string for these trap attributes is public.

Parameter	OID Description
stTrapDesc	OID: SNMPv2-SMI::enterprises.18552.2.2.1.3
stLocalAddress	OID: SNMPv2-SMI::enterprises.18552.2.2.1.4
stTime	OID: SNMPv2-SMI::enterprises.18552.2.2.1.5
stTrapSourceIP	OID: SNMPv2-SMI::enterprises.18552.2.2.1.6
stUserOid	OID: SNMPv2-SMI::enterprises.18552.2.2.1.7
stTrapMsg	OID: SNMPv2-SMI::enterprises.18552.2.2.1.8

Alert Message Samples

The following are alert message samples. The log message print format is specific to the trap receiver. You can change these according to your production environment.

SNMP version: 1

```
2008-06-19 19:13:01 10.0.30.33(via UDP: [172.16.0.1]:50326) TRAP, SNMP
v1, community public
       SNMPv2-SMI::enterprises.18552.1.3 Enterprise Specific Trap (98) Uptime:
21
       SNMPv2-SMI::enterprises.18552.1.2.1.3 = STRING: "(Description) "
SNMPv2-SMI::enterprises.18552.1.2.1.4 = IpAddress: 10.0.30.33   SNMPv2-
SMI::enterprises.18552.1.2.1.5 = Timeticks: (1213927981) 140    SNMPv2-
SMI::enterprises.18552.1.2.1.6 = STRING:
"10.0.30.61,10.2.1.13,10.0.30.33,10.2.1.14,10.0.0.11,10.0.30.33,10.0.30.
19,10.0.30.34,10.0.30.61,10.2.1.10,10.2.1.9,10.2.1.12,20.11.2.211,20.11.
1.211,20.11.4.211,20.11.8.211,20.11.10.211,20.11.7.211,20.11.11.211,20.1
1.5.211,20.11.9.211,20.11.14.211,20.11.15.211,20.11.16.211,20.11.17.211,
20.11.18.211,20.11.13.211,20.11.12.211,20.11.20.211,20.11.19.211,20.11.2
1.211,20.11.24.211,20.11.25.211,20.11.22.211,20.11.23.211,20.11.26.211,2
0.11.27.211,20.11.28.211,20.11.29.211,20.11.30.211"      SNMPv2-
SMI::enterprises.18552.1.2.1.7 = OID: SNMPv2-SMI::enterprises.18552
SNMPv2-SMI::enterprises.18552.1.2.1.8 = STRING: "Optional Text;High-ABS-
MESSAGE-VOLUME;volume alert1;msgrate:2833;highthresh:110;source:
0.0.0.0"
```

SNMP version: 2

```
2008-06-19 18:37:25 NET-SNMP version 5.2.1 Started.
2008-06-19 18:37:56 172.16.0.1 [UDP: [172.16.0.1]:50276]:
SNMPv2-MIB::sysUpTime.0 = Timeticks: (1818885) 5:0     SNMPv2-
MIB::snmpTrapOID.0 = OID: SNMPv2-SMI::enterprises.18552.1.3.0.98        SNMPv2-
SMI::enterprises.18552.1.2.1.3 = STRING: "(Description) "       SNMPv2-
SMI::enterprises.18552.1.2.1.4 = IpAddress: 10.0.30.33  SNMPv2-
SMI::enterprises.18552.1.2.1.5 = Timeticks: (1213925876) 140    SNMPv2-
SMI::enterprises.18552.1.2.1.6 = STRING:
"10.0.30.61,10.2.1.13,10.0.30.33,10.2.1.14,10.0.0.11,10.0.30.33,10.0.30.
19,10.0.30.34,10.0.30.61,10.2.1.10,10.2.1.9,10.2.1.12,20.11.2.211,20.11.
1.211,20.11.4.211,20.11.8.211,20.11.10.211,20.11.7.211,20.11.11.211,20.1
1.5.211,20.11.9.211,20.11.14.211,20.11.15.211,20.11.16.211,20.11.17.211,
20.11.18.211,20.11.13.211,20.11.12.211,20.11.20.211,20.11.19.211,20.11.2
1.211,20.11.24.211,20.11.25.211,20.11.22.211,20.11.23.211,20.11.26.211,2
0.11.27.211,20.11.28.211,20.11.29.211,20.11.30.211"      SNMPv2-
SMI::enterprises.18552.1.2.1.7 = OID: SNMPv2-SMI::enterprises.18552
SNMPv2-SMI::enterprises.18552.1.2.1.8 = STRING: "Optional Text;High-ABS-
MESSAGE-VOLUME;volume alert1;msgrate:3254;highthresh:110;source:
0.0.0.0"
```

Definition of Configuration Rule Files

You must define a configuration file containing a rule to format log messages before they are forwarded.

The formatted forwarding feature is used to format log messages prior to forwarding. The Forwarding function uses the configuration file that defines the formatting rules. All log messages that match the forwarding rule are formatted.

Rules consist of Regular expressions that are applied to the log messages. If log messages are matched with the Regular expression, then the extracted strings are substituted into the forwarded message before forwarding it to the defined destination.

You can upload only one configuration file for each Message Routing rule. Each configuration rule file can be used across multiple destinations or rules, and can have a maximum file size of 1 GB.

The configuration rule file is used for:

- filtering logs
- transforming/formatting log messages
- filtering character patterns from a log (shredding) and, optionally, replacing them with literal strings

A configuration rule file can consist of multiple rules. A rule consist of either regexp/template pattern or shred/replace pattern.

A regexp/template rule consists of two lines; a Regular expression used to match and extract patterns and the forwarded message template. A shred/replace pattern rule defines shred= option. You can also have both types of rules in the same configuration rule file.

First, regexp/template rules are applied in the order of their appearance in the configuration file. The system stops applying rules as soon as the first match is detected. The shred/replace rules are applied after a single scan across all regexp/template rules. Only those log messages that match any regexp rule are forwarded. However, if you use forward all option, any log message are forwarded even without matching any regexp rule.

The configuration rule file must satisfy the following criteria:

- rule file cannot be empty

- must consist of one or more rules; either regexp/template rule or shred/replace rule or both types can be in the same rule file
- regexp <regular expression> must be valid; cannot be empty

  regexp is a keyword followed by a PCRE-compatible regular expression. A `regexp` has to succeed in order for a log message to be modified.

  For example,

  ```
  regexp(.+)\s(.+)\slx_scheduler:\s(.+)\s(.+):\s\((.+)\)\slx_
  scheduler:\sending\sconfig\smgmt\sjob
  ```

  Note: It is good practice not to end the regular expression with a capturing pattern (.+), because it affects performance when capturing a large portion of the message.

- template must be valid; cannot be empty

  The rule has to specify a template for the contents of the forwarded log. The template can contain literal strings as well as portions of the original message. The template is applied to create a new message. After template-driven formatting is completed, the new log message is forwarded to the destination of the forwarding rule. $1, $2, and so on, in the template refer to the patterns extracted by `regexp`. Extracted patterns are defined by a matching pair of parentheses, for example (.+) in the previous regular expression example. $1 refers to the first pair of parentheses.

  It is good practice to use up to $10 matching patterns in the template. You can use the matching patterns repeatedly, however, the digits after the $ sign should be 10 or less than 10.

  For example,

  ```
  template LOGS|LogLogic MODEL|CC01|Application configuration
  change|5|deviceExternalId=62968-1 msg=ending config mgmt job
  shost=$2 dhost=$2 suser=system suid=system spriv=User
  ```

- shred/replace option (Optional)

  Shred option removes the portion of the transformed message based on the template and replaces with the new string, and then forwards the message to the destination. For example,

  ```
  shred=\d{16} replace=XYZ-XYZ-XYZ-XYZ
  ```

`shred=` regular expression is applied to every message repeatedly in order to identify a sensitive pattern. The sensitive pattern can have more than one instance in the message, all instances are detected. The `shred` in the config file must be all in lower-case letters.

`replace=` optionally defines a substitution string that replaces sensitive string.

Additional Options in the Configuration Rule File

It is good practice to add additional options (such as `source_type` or `match`) in the configuration rule file, which help filter log messages faster and speed up the performance.

If the rule file includes these options, the following criteria must be satisfied:

- source-type must have a valid log type; cannot be empty

  `source_type=` is used to limit regexp filters to logs of specific source type.

  For example,

  ```
  source_type=LogLogic Appliance
  ```

- Match pattern must be valid; cannot be empty

  `match=` <unique string matching part of the log message>

  The “match” can be used and validated only when written as “`match=`”. If you use “match” only (without =), then it is not validated.

  This applies the matching string to each log. Only those logs containing the string are tested against the regular expression. The match option is very helpful when the regexp gets complicated.

  For example,

  ```
  match=action:login;status:success;
  ```

- forwardall

  To enable all logs to be forwarded to the destination, use the `forwardall` option. This option is required with shred/replace, when there is no regexp/template in the rule.

Warning: Multiple rules can refer to the same configuration rule file. If you overwrite the configuration file, all rules referencing to that configuration file are affected. Therefore, when creating a configuration rule file, always use a different file name to preserve an existing rule file and keep a copy of all configuration rule files that have been uploaded.

Examples of Configuration Rules

The following sample rule file shows both types of rules. In the following examples, $ pattern in the template is replaced in the formatted message. Each matching $ pattern is color coded in the following examples.

```
source_type=LogLogic Appliance
######################
# LOGLOGIC EVENTS #
######################
Scenario#1: Transform messages before forwarding to downstream
match=ending config mgmt job
regexp(.+)\s(.+)\slx_scheduler:\s(.+)\s(.+):\s\((.+)\)\slx_
scheduler:\sending\sconfig\smgmt\sjob
template LOGS|LogLogic MODEL|CC01|Application configuration
change|5|deviceExternalId=62968-1 msg=ending config mgmt job shost=$2
dhost=$2 log_type=$3 process_id=$4 session_start=$5 suser=system
suid=system spriv=User
```

Note: When writing a regular expression, care should be taken to define as few capturing patterns as possible. Define only those capturing patterns that must be used in the template to format the forwarded message. Adding unused capturing patterns to the regular expression can quickly degrade the forwarding performance.

If the original message is as follows:

```
<11>Sep 12 20:49:41 localhost lx_scheduler: %LOGLOGIC-PRI-6 8329:
(1315860581) lx_scheduler: ending config mgmt job
The formatted log message will appear as follows:
LOGS|LogLogic MODEL|CC01|Application configuration
change|5|deviceExternalId=62968-1 msg=ending config mgmt job
```

```
shost=localhost dhost=localhost log_type=%LOGLOGIC-PRI-6 process_
id=8329: session_start=1315860581 suser=system suid=system spriv=User
In the previous example, $2 in the template is replaced with localhost;
$3 is replaced with %LOGLOGIC-PRI-6; $4 is replaced with 8329:; $5 is
replaced with 1315860581.
Scenario#2: Extract fields only
match=action:logoff;  status:success;
regexp(.+)\s+?(%L.+)\s+?user:(.+);\s+?module:(.+);\s+?action:
(.+);\s+?status:(.+);\s+?session_id:(.+);\s+?client_ip:(.+);\s+?target_
ip:(.+);\s+?session_start:(.+);\s+?session_duration:(.+);\s+?disconnect_
reason:(.+);\s+?info:
template LOGS|LogLogic MODEL|$3 $4 $5 $6 $7 $8 $9 $10
If the original message is as follows:
<182> Sep 13 18:50:24 20.20.20.20 %LOGLOGIC-6-3102: user:admin;
module:user_intfc;  action:logoff; status:success; session_
id:4203070123;  client_ip:10.10.0.1;  target_ip:20.20.20.20; session_
start:1315965001;  session_duration:23;  disconnect_reason:user_logoff;
info:sign out, orig_session_id,FA85C2AB28037AC810F8A8BCB71B4A33,
Then, after running the rule, the formatted log message will appear as
follows:
LOGS|LogLogic MODEL|admin user_intfc logoff success 4203070123 10.10.0.1
20.20.20.20 1315965001
# Scenario#3: forward all RAW messages and shred 4203070123 to replace
it with XXXXXXXXX
source_type=LogLogic Appliance
########################
# LOGLOGIC EVENTS #
########################
forwardall
shred=4203070123 replace=XXXXXXXXX
If the original log message is as follows:
<182> Sep 13 18:50:24 20.20.20.20 %LOGLOGIC-6-3102: user:admin;
module:user_intfc;  action:logoff;  status:success;  session_
id:4203070123;  client_ip:10.10.0.1;  target_ip:20.20.20.20;  session_
start:1315965001;  session_duration:23;  disconnect_reason:user_logoff;
info:sign out, orig_session_id,FA85C2AB28037AC810F8A8BCB71B4A33,
Then, after running the rule, the formatted log message will appear as
follows; where 4203070123 will be replaced with XXXXXXXXX:
<182> Sep 13 18:50:24 20.20.20.20 %LOGLOGIC-6-3102: user:admin;
module:user_intfc;  action:logoff;  status:success;  session_
id:XXXXXXXXXX;  client_ip:10.10.0.1;  target_ip:20.20.20.20;  session_
start:1315965001;  session_duration:23;  disconnect_reason:user_logoff;
info:sign out, orig_session_id,FA85C2AB28037AC810F8A8BCB71B4A33,
```

Defining a Configuration Rule File

You can define a configuration rule file.

Procedure

1. Navigate to **Administration > Message Routing** navigation menu.
2. Select a rule and click **Add Destination**.
3. Specify the fields as described in Adding Destinations to the All Sources Rule or Creating a New Outbound Routing Rule.
4. To specify the **Format Rule Definition**, click the **Browse** button to select the saved configuration rule file.
5. Click the **Upload** button to upload the configuration file.

   The file is validated and the **Test** button gets enabled if the file is validated successfully.
6. To verify if the formatting rules in the configuration file work as expected, click the **Test** button.

   Note:
   - For SNMP protocol, the format forwarding rule feature is not supported.
   - The uploaded configuration rule file name is automatically converted into lower case.
7. Enter a sample log message in the **Message** field and click the **Test Now** button. If the message matches the regexp, the information is extracted from the file and the **Regex**, **Template**, and **Output** fields are displayed.

   Extracted Configuration Rule File

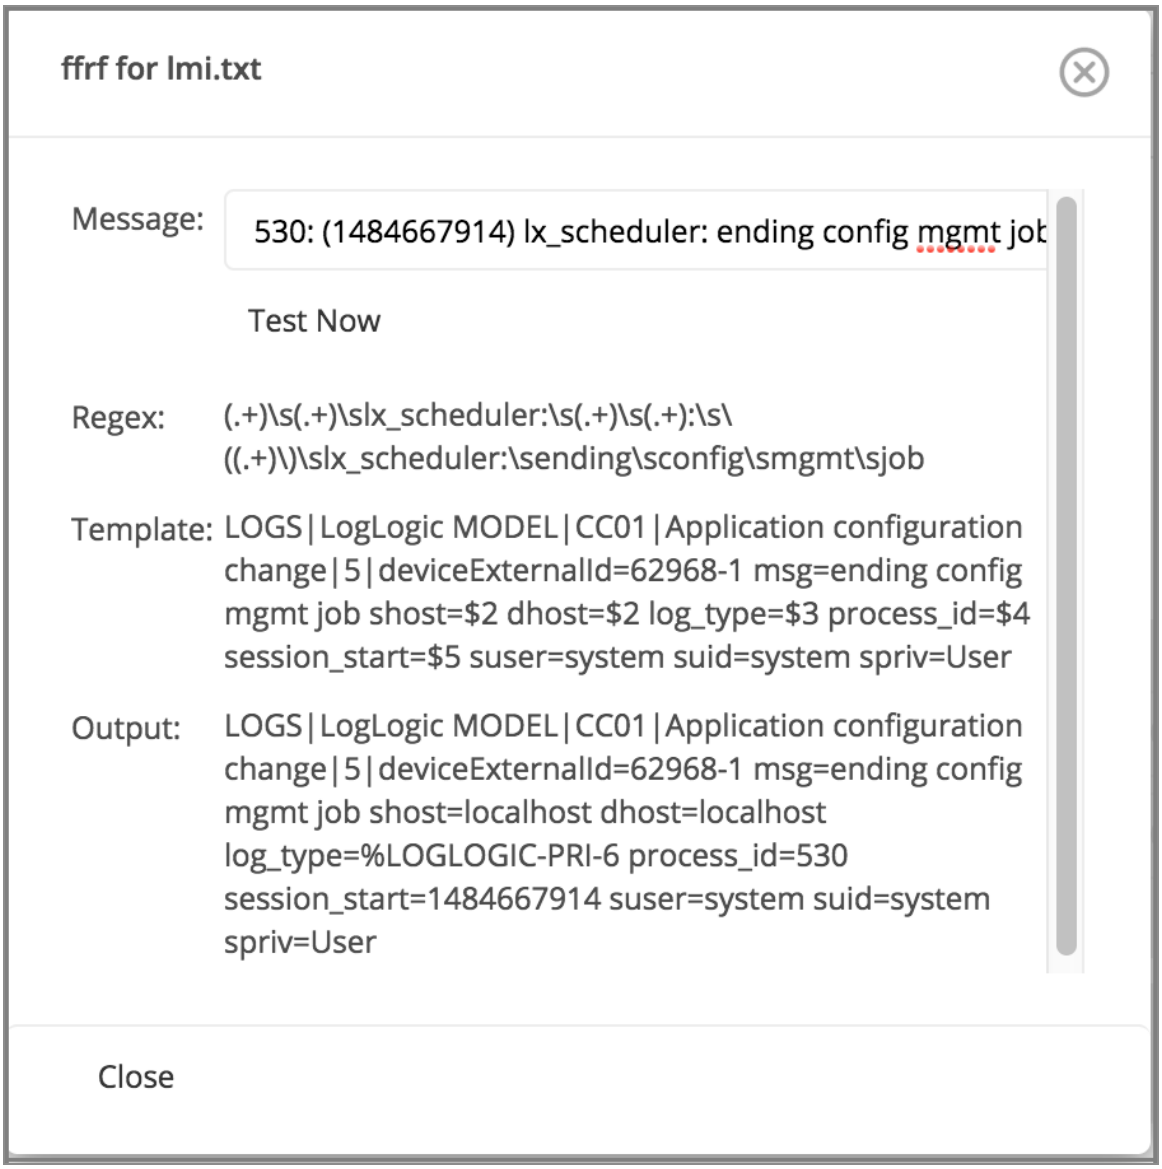

Note: Based on the template, the transformed message in the **Output** field could be impacted if `shred/replace` option is currently used in the configuration rule file. **Regex** and **Template** fields on the Test window can be empty if `forwardall` option is used in the configuration rule file.

8. Click **Close** to close the Test window.

Configuration of LogLogic iDRAC

You must configure the iDRAC network connectivity.

Beginning with H4 appliances, TIBCO LogLogic® appliances include the Dell iDRAC utility for a more convenient low-level TIBCO LogLogic® appliance administration.

The iDRAC interface is available as a local console and a web interface. The web interface is enabled by default on all TIBCO LogLogic® appliances, and relies on the iDRAC-designated interface being connected to the network infrastructure. If this interface is left disconnected, the iDRAC interface is not accessible remotely, but is still accessible in the local console.

By default on TIBCO LogLogic® appliances, the labeled iDRAC network interface has an assigned static IPv4 address of `192.168.0.120/24`. By connecting the iDRAC network interface to a network infrastructure, the iDRAC web interface becomes available through HTTPS, at `https://192.168.0.120` as well as telnet and SSH to the same default IP.

For instructions to change the network connectivity from the local console, see *TIBCO LogLogic® Log Management Intelligence Hardware Installation Guide*.

LogLogic LMI Ports

This section contains all ports that are open on a LogLogic LMI host, including LogLogic LSP ports.

TIBCO LogLogic® Universal Collector ports are not documented in this table.

A port serves as a destination on either the appliance or a foreign host depending on direction. In the following tables, inbound direction refers to the port number on the local appliance and outbound direction indicates the port number on the remote host.

Socket interface: The LogLogic LMI appliances use a dual stack IPv4/v6 configuration by mapping the IPv4 addresses of the appliance into the IPv6 address space.

Note: Some ports might not be used at all times such as when HA mode is enabled.

References to IPv6 in the following table do not indicate being able to collect data using IPv6, and hence, the reason the appliance has a dual stack configuration. All external communication must still occur over IPv4.

Port Assignments

A list of ports, directions, and description.

This section includes a list of ports:

- LogLogic LMI Port Assignments - inbound
- Bidirectional ports
- Internal ports
- Outbound ports

LogLogic LMI processes that do not require ports are listed in the LogLogic LMI Processes that do not Require Ports section.

LogLogic LMI Port Assignments - inbound

Inbound port assignments are listed in the following table.

Port	Socket Interface	Transport	Process Name	Description	LogLogic LMI or OS?
22	all (IPv4)	tcp	sshd	CLI access for root/toor using Secure Shell (SSH) / TCP syslog and LLTCP with encryption.	OS
80	all (IPv6)	tcp	java(Tomcat)	HTTP access to the web GUI. Redirects to 443.	LogLogic LMI
123	all (IPv4) IPv6 local link	udp	ntpd	Network Time Protocol (NTP) service for using the appliance as a time source.	OS
161	all (IPv4)	udp	snmpd	Listens for poll requests by SNMP monitoring applications gathering SNMP-related info about appliance.	OS
162	all (IPv4)	udp	engine_ trapcollector	To receive SNMP traps from log sources.	OS
443	all (IPv6)	tcp	java(Tomcat)	HTTPS access to	LogLogic

Port	Socket Interface	Transport	Process Name	Description	LogLogic LMI or OS?
				the web GUI.	LMI
514	all (IPv4)	udp	engine_ collector	Receives syslog (UDP syslog) messages.	LogLogic LMI
514, 6514	all (IPv4)	tcp	engine_ tcpcollector	Receives syslog (TCP syslog) messages and TLS syslog messages.	LogLogic LMI
2055 9555 9995	all	tcp	LogLogic LSP Collector	LogLogic LSP Collector for Netflow	LogLogic LMI
3306	all (IPv4)	tcp	mysqld	MySQL database.	LogLogic LMI
4400	all (IPv4)	tcp	engine_cluster_ membership	Rsync replication failover service (receives connection from peer node) (HA mode only).	LogLogic LMI
4433	all (IPv4)	tcp	engine_http_ collector	http-based log collection (Blue Coat, NetApp, and so on)	LogLogic LMI
4433	all (IPv6)	tcp	java (Tomcat)	Management station: Used to receive updates from a remote appliance	LogLogic LMI

Port	Socket Interface	Transport	Process Name	Description	LogLogic LMI or OS?
5514	all (IPv4)	tcp	engine_ rcollector	ULDP prior to LogLogic LMI 5.2	LogLogic LMI
5514	all (IPv4)	tcp	engine_ rcollector	LogLogic TCP-based message routing.	LogLogic LMI
5515	all (IPv4)	tcp	stunnel	Secure ULDP collection.	LogLogic LMI
5516	all (IPv4)	tcp	engine_ uldpcollector	ULDP for LogLogic LMI 5.2 and later.	LogLogic LMI
7000 - 8000	localhost (IPv4 & v6)	tcp	ssh	Used as the tunnel mechanism by engine_stunnel for forwarding to downstream appliances when authentication and encryption are enabled. Four ports are used at a time. The specific 4 ports used increment each time when a particular tunnel is started so that there are no conflicts. The first port of the set is for forwarding syslog traffic, the second port is for http data, the	LogLogic LMI

Port	Socket Interface	Transport	Process Name	Description	LogLogic LMI or OS?
				third is for file data using rcollector and the fourth is for Check Point data.	
8080	all (IPv6)	tcp	java (Tomcat)	Provides a destination for web browser redirects during LogLogic LMI upgrade.	LogLogic LMI
9013	all (IPv6)	tcp	java	Used for listening by TIBCO eventdistributor client.	LogLogic LMI
9680	all	tcp	logu-web	Webapp service	LogLogic LMI
9681	all	tcp	logu-querynode	Query node REST service	LogLogic LMI
9683	all	tcp	logu-datanode	Data node REST service	LogLogic LMI
11965	default gw	tcp	ll_tunnel	Message forwarding when using LogLogic TCP with encryption. Note: This is deprecated for 5514/tcp w/o encryption and	LogLogic LMI

Port	Socket Interface	Transport	Process Name	Description	LogLogic LMI or OS?
				22/tcp with encryption.	

LogLogic LMI Port Assignments - bidirectional

Bidirectional port assignments are listed in the following table.

Port	Socket Interface	Transport	Process Name	Description	LogLogic LMI or OS?
9611	all	tcp	logu-datanode	Data node ingest service	LogLogic LMI
9620	all	tcp	logu-querynode	Query node query service	LogLogic LMI
9622	all	tcp	logu-datanode	Data node streaming service	LogLogic LMI
9626	all	tcp	logu-aggregationnode	Aggregation node query service	LogLogic LMI
9682	all	tcp	logu-correlationnode	Correlation node REST service	LogLogic LMI
9683	all	tcp	logu-datanode	Distributed Advanced Search	LogLogic LMI
9685	all	tcp	logu-aggregationnode	Aggregation node REST service	LogLogic LMI

Port	Socket Interface	Transport	Process Name	Description	LogLogic LMI or OS?
9687	all	tcp	logu-monitoringconsole	Monitoring console REST service	LogLogic LMI
9688	all	tcp	logu-monitoringconsole	Monitoring Console cluster service for the domain type `LMI Domain`	LogLogic LMI
9690 - 9700	all	tcp	logu-monitoringconsole	Recommended for use for additional domains in Monitoring Console	LogLogic LMI
9880	all	tcp	logu-web	WebApp HTTP: Redirect to HTTPS	LogLogic LMI

LogLogic LMI Port Assignments - internal

Port	Socket Interface	Transport	Process Name	Description	LogLogic LMI or OS?
199	localhost (IPv4)	tcp	snmpd	SNMP Unix Multiplexer.	OS
768	all (IPv4)	raw	engine_ collector	Used for internal logging	LogLogic LMI
768	all (IPv4)	raw	engine_	Used for internal	LogLogic

Port	Socket Interface	Transport	Process Name	Description	LogLogic LMI or OS?
			highpri_reader	logging	LMI
768	all (IPv4)	raw	engine_lx_ scheduler	Used for internal logging	LogLogic LMI
768	all (IPv4)	raw	engine_lx_ parser	Used for internal logging	LogLogic LMI
768	all (IPv4)	raw	engine_ tcpcollector	Used for internal logging	LogLogic LMI
768	all (IPv4)	raw	engine_ tcpforwarder	Used for internal logging	LogLogic LMI
768	all (IPv4)	raw	engine_ trapcollector	Used for internal logging	LogLogic LMI
768	all (IPv4)	raw	engine_ uldpcollector	Used for internal logging	LogLogic LMI
1099	all (IPv6)	tcp	java (LogLogic LSP)	Used for LogLogic LSP core communication to Java RMI registry.	LogLogic LMI
1514	all (IPv6)	udp	engine_ collector	Used for logs with Domain ID	LogLogic LMI
2098	all (IPv6)	tcp	java (MC Agent)	Java RMI Registry service for Tomcat (only when MC Agent installed).	LogLogic LMI
2099	all (IPv6)	tcp	java (MC Agent)	Java instance listening for Shutdown/Reboot	LogLogic LMI

Port	Socket Interface	Transport	Process Name	Description	LogLogic LMI or OS?
				command (only when MC Agent installed).	
2508	all (IPv6)	tcp	java (MC Agent)	MCAgent	LogLogic LMI
4401	all (IPv4)	tcp	engine_ cluster_ membership	Cluster membership monitor (receives connection from cluster_ membership and mysqld engines) (HA mode only).	LogLogic LMI
8005	localhost (IPv6)	tcp	java (Tomcat)	Tomcat administration port.	LogLogic LMI
8180	localhost (IPv6)	tcp	java (MC Agent)	SSH port for Karaf - (only when MC agent is installed).	LogLogic LMI
9443	all	tcp	java (Tomcat)	HTTPS Remote Control	LogLogic LMI
9600	all (IPv4)	tcp	llzk	Used by zookeper for configuration of Advanced Features	LogLogic LMI
9621	all	tcp	logu-datanode	Data node query service	LogLogic LMI
31000	localhost (IPv6)	tcp	java (LogLogic LSP)	LogLogic LSP Core.	LogLogic LMI
32000	localhost	tcp	java (LogLogic	Wrapper binary for	LogLogic

Port	Socket Interface	Transport	Process Name	Description	LogLogic LMI or OS?
			LSP)	LogLogic LSP.	LMI
32001	localhost	tcp	java (MC Agent)	Wrapper binary for MC Agent (only when MC Agent installed).	LogLogic LMI
32768-61000	all (IPv4)	udp	engine_archive	Performs archiving on LogLogic ST Appliances.	LogLogic LMI
32768-61000	all (IPv4)	udp	engine_ collector	Manages real-time syslog collection	LogLogic LMI
32768-61000	all (IPv4)	udp	engine_ filecollector	Manages file Xfer rules, deep parses file-based log data, assists with forwarding of file-data.	LogLogic LMI
32768-61000	all (IPv4)	udp	engine_ highpri_reader	Handles message forwarding, search filter alerts (LogLogic LX Appliance only), real-time view feeds.	LogLogic LMI
32768-61000	all (IPv4)	udp	engine_lx_ scheduler	Handles periodic tasks such as aggregation, cleanup, alerts.	LogLogic LMI
32768-61000	all (IPv4)	udp	engine_rsender	Handles forwarding when LogLogic TCP	LogLogic LMI

Port	Socket Interface	Transport	Process Name	Description	LogLogic LMI or OS?
				is used as the protocol.	
32768-61000	all (IPv4)	udp	engine_st_reporter	Handles regex searches.	LogLogic LMI
32768-61000	all (IPv4)	udp	engine_syslog	Replays /var/log/sys.log file back into UDP collector so we can parse our own syslog messages.	LogLogic LMI
32768-61000	all (IPv4)	udp	engine_sysmon	Monitors system and issues system alerts. Monitors memory, system load avg, # of zombie processes and logs to sys.log file every 5 minutes.	LogLogic LMI
32768-61000	all (IPv4)	udp	engine_tcpcollector	Involved in collection when using syslog-ng (TCP syslog).	LogLogic LMI
32768-61000	all (IPv4)	udp	engine_tcpforwarder	Used for internal logging	LogLogic LMI
32768-61000	all (IPv4)	udp	engine_trapcollector	Used for internal logging	LogLogic LMI
32768-61000	all (IPv4)	udp	engine_uldpcollector	Process and forward SNMP traps to remote hosts.	LogLogic LMI

LogLogic LMI Outbound Port Assignments

Outbound port assignments are listed in the following table.

Dest Port	Socket Interface	Transport	Process Name	Description	LogLogic LMI or OS?
22	default gateway	tcp	ssh	SSH-based backups	OS
25	default gateway	tcp	llmail, msmtp, or Tomcat	Sends emails to an SMTP server. The process used is dictated by what is being sent (alerts, reports, and so on).	LogLogic LMI
49	default gateway	tcp	java (Tomcat)	TACACS authentication (but no authorization) for users.	LogLogic LMI
68	all (IPv4)	udp	dhclient	Manages DHCP client IP settings.	LogLogic LMI
88	default gateway	udp	java (Tomcat)	Kerberos feature when using LDAP.	LogLogic LMI
111	default gateway	tcp	Sun RPC portmapper	LogLogic LMI NFS backups and archiving: mount command communicates to Sun RPC Port mapper to get	OS

Dest Port	Socket Interface	Transport	Process Name	Description	LogLogic LMI or OS?
				port # for mountd (NFS v3 only)	
111	default gateway	udp	Sun RPC portmapper	LogLogic LMI NFS backups and archiving: mount command communicates to Sun RPC Port mapper to get port # for mountd (NFS v3 only)	OS
123	default gateway	udp	ntpd	Network Time Protocol (NTP) service for using the appliance as a time source.	OS
389	default gateway	tcp	java (Tomcat)	LDAP to Active Directory.	LogLogic LMI
636	default gateway	tcp	java (Tomcat)	LDAP to Active Directory.	LogLogic LMI
>1023	default gateway	tcp	various	Interact with multiple server daemons (statd, lockd, rquotad, mountd) for using NFS.	OS
1433	default gateway	tcp	java (LogLogic LSP)	Microsoft SQL Server JDBC or GDBC collection (with and without	LogLogic LMI

Dest Port	Socket Interface	Transport	Process Name	Description	LogLogic LMI or OS?
				TLS enabled)	
1521	default gateway	tcp	java (LogLogic LSP)	Oracle Database JDBC or GDBC collection (when TLS is disabled)	LogLogic LMI
1812	default gateway	tcp	java (Tomcat)	RADIUS	LogLogic LMI
2049	default gateway	tcp	nfs	LogLogic LMI NFS (v3) backups and archiving: data transfer occurs using this port. v4 supported from version 6.2.0. v4 uses this port for mounting, locking, and data transfer.	OS
2484	default gateway	tcp	nfs	Oracle database GDBC collection (only when TLS is enabled)	OS
2561	default gateway	tcp	java (Hawk console node)	Hawk console to Hawk TCP daemon on agents	LogLogic LMI
2581	default gateway	tcp	java (Hawk console node)	Hawk console self host	LogLogic LMI

Dest Port	Socket Interface	Transport	Process Name	Description	LogLogic LMI or OS?
3306	default gateway	tcp	java (LogLogic LSP)	MySQL Database GDBC collection.	LogLogic LMI
4433	all (IPv4)	tcp	engine_http_ collector	File-based message routing	LogLogic LMI
4433	all (IPv6)	tcp	java (Tomcat)	Management station: Used to send requests to a remote appliance.	LogLogic LMI
7222	default gateway	tcp	LLCollectors	TIBCO Enterprise Message Service™ collection (TLS disabled)	LogLogic LMI
7243	default gateway	tcp	LLCollectors	TIBCO Enterprise Message Service™ collection (TLS enabled)	LogLogic LMI
9000	all	tcp	engine_ filecollector	Used by HDFS client to connect to HDFS cluster. See how to change the port number.	LogLogic LMI
9092	default gateway	tcp	LLCollectors	Apache Kafka (TLS disabled)	LogLogic LMI
9093	default gateway	tcp	LLCollectors	Apache Kafka (TLS enabled)	LogLogic LMI
9600	all (IPv4)	tcp	llzk	Used by	LogLogic

Dest Port	Socket Interface	Transport	Process Name	Description	LogLogic LMI or OS?
				ZooKeeper	LMI
18184	default gateway	tcp	chkpt_agent	Used by LEA for log export from LEA server.	LogLogic LMI
18190	default gateway	tcp	chkpt_agent	Used by CheckPoint Mgmt Interface (CPMI) for communication between LogLogic LMI and Mgmt Module.	LogLogic LMI
18210	default gateway	tcp	chkpt_agent	Used by Secure Internal Communication (SIC) for pulling certificates from Mgmt Module.	LogLogic LMI
32768-61000	all (IPv4)	tcp	engine_ tcpforwarder	Perform message routing when using syslog-ng (TCP syslog).	LogLogic LMI
dynamic port	default gateway	tcp	rpc.mountd	NFS sharing: port used by the mount command over TCP outbound to an NFS server	OS
dynamic port	default gateway	tcp	NFS client	NFS file sharing: used for file	OS

Dest Port	Socket Interface	Transport	Process Name	Description	LogLogic LMI or OS?
				locking	
dynamic port	default gateway	udp	NFS client	Used by NFS v3 to access the rpc.mountd service on the NFS server for performing the actual mount operation	OS
dynamic port	default gateway	udp	NFS client	Used by NFS v3 to access the rpc.lockd service on the NFS server to acquire a file lock when accessing archived data	OS

LogLogic LMI Processes that do not Require Ports

The following LogLogic LMI processes do not need to bind to any port for accepting data from other components.

Process	Description
engine_alerting	Manages some types of alerts such as baseline ratio-based, message rate alerts, and so on
engine_backup	Mirrors the existing data stores (MySQL database, raw logs in /loglogic/data/vol1, system configuration files) to a remote host.
engine_cluster_monitor	Monitors the replication of data and the replication

Process	Description
	configuration, and restarts it if it does not respond.
engine_mysqld	Monitors mysqld and restarts it if it does not respond
engine_ntp	Monitors ntp and restarts it if it does not respond.
engine_tcp_scheduler	Monitors the data files created by engine_rsender in /loglogic/data/rsender/ready so they can be transmitted to their destination.
ll_opsec_manager	Manages OPSEC suite of protocols for Check Point log sources. Uses chkpt_agent for the actual work and manages the startup and shutdown of those agent processes.

Examples of LogLogic Port Assignments

Log Message Push

Description	Protocol	Port #	Comments
Syslog	UDP	514	Used for incoming syslog data. You can change this port number from 514 in the **System Settings > General** tab **Syslog UDP Port** field. If you change this port number, you must add the other port number here.
Blue Coat/Netcache	HTTP/ HTTPS	4433	Used for incoming HTTPS streams from log sources such as Blue Coat ProxySG and NetApp Netcache.

Check Point

Description	Protocol	Port #	Comments
lea_server	LEA/TCP	18184	Used to transfer log messages.
cpmi_server	TCP	18190	Default port. Used for rule listing and firewall/interface auto-discover. Note: Must match Check Point Manager Server.
SIC	TCP	18210	Used to establish connection with the Check Point Management Interface (CPMI). SIC - Secure Internal Communication
CMPI Forwarding	UDP	5514	Used for collecting LogLogic streams from the Check Point Management Interface through the rtchpk utility.

GUI

Description	Protocol	Port #	Comments
Browser	HTTP	80	Used for internal web browser access requests to the LogLogic Appliance. The requests are redirected to port 443 (HTTPS).
Browser	HTTPS	443	Used for incoming HTTPS requests to the GUI. The requests are redirected from port 80 (HTTP).
Browser	HTTP	8080	Browser redirects during upgrade.

Note: If you are running java 1.8.0_x, you must perform the following steps:

1. As administrator, update your `C:\Program Files (x86)\Java\jre1.8.0_x\lib\security\java.policy` file and grant the following permission to non-abbreviated IPv6 address:

   ```
   grant { permission java.net.SocketPermission
   "fd00:0:0:0:0:aaaa:a73:1a3d", "connect,resolve"; };
   ```

   You can also add permissions to both abbreviated and non-abbreviated addresses:

   ```
   grant { permission java.net.SocketPermission
   "fd00:0:0:0:0:aaaa:a73:1a3d", "connect,resolve"; };
   grant { permission java.net.SocketPermission "fd00::aaaa:a73:1a3d",
   "connect,resolve"; };
   ```

   The IP address should be replaced with the IP address of your appliance.

2. In **Control Panel > java > Security** add the following to the exception list:

   ```
   https://[fd00::aaaa:a73:1a3d]:443, where "fd00::aaaa:a73:1a3d” is
   your appliance IP
   https://[fd00:0:0:0:0:aaaa:a73:1a3d]:443, where
   "fd00:0:0:0:0:aaaa:a73:1a3d” id the non-abbreviated version for
   your appliance IP
   ```

Note: Appliance IP address can be either IPv4 or IPv6. Both are supported.

Miscellaneous

Description	Protocol	Port #	Comments
CLI Access	SSH	22	Used for SSH client access. Configured on/off.
NTP	NTP	123	Used by the Network Time Protocol Daemon (NTPD).
Browser	HTTPS	443	Used for SSL two-way handshake.

Failover

Description	Protocol	Port #	Comments
High Availability Failover	Rsync	4400	Used by the replication sync failover service.
High Availability Failover	MySQL	3306	Used by the MySQL failover service.

Outbound Traffic

Description	Protocol	Port #	Comments
LogLogic TCP	TCP	5514	Used for collecting LogLogic streams from the Check Point Management Interface via the `rtchpk` utility.
LogLogic TCP	TCP	9443	Used by Management Station to send requests from the Management Station to a Remote Appliance.
LogLogic TCP	TCP	9443	Used for sending updates from a Remote Appliance to the Management Station.
Syslog Alert	UDP	514	Used for incoming syslog data. You can change this port number on the **System Settings > General** tab > the **Syslog UDP Port** field. If you change this port number, you must add the other port number here.
SNMP Alerts	UDP	161	Used for incoming SNMP client requests.
SNMP Notification	UDP	162	Used for incoming and outgoing SNMP trap messages. (Internal alerts from LogLogic LX Appliance or LogLogic ST Appliance, and log collection)

TIBCO Documentation and Support Services

For information about this product, you can read the documentation, contact TIBCO Support, and join TIBCO Community.

How to Access TIBCO Documentation

Documentation for TIBCO products is available on the TIBCO Product Documentation website, mainly in HTML and PDF formats.

The TIBCO Product Documentation website is updated frequently and is more current than any other documentation included with the product.

Product-Specific Documentation

The following documentation for this product is available on the TIBCO LogLogic® Log Management Intelligence Product Documentation page.

- *TIBCO LogLogic® Log Management Intelligence Release Notes*
- *TIBCO LogLogic® Log Management Intelligence Administration*
- *TIBCO LogLogic® Log Management Intelligence Configuration and Upgrade*
- *TIBCO LogLogic® Log Management Intelligence Hardware Installation Guide*
- *TIBCO LogLogic® Log Management Intelligence Log Source Report Mapping*
- *TIBCO LogLogic® Log Management Intelligence Security Guidelines*
- *TIBCO LogLogic® Log Management Intelligence SSD Hardware Field Installation*
- *TIBCO LogLogic® Log Management Intelligence Syslog Alert Message Format Quick Reference*
- *TIBCO LogLogic® Log Management Intelligence User Guide*
- *TIBCO LogLogic® Log Management Intelligence XML Import/Export Entities Reference*
- *TIBCO LogLogic® Enterprise Virtual Appliance Quick Start*

Other TIBCO Product Documentation

The following documents for TIBCO LogLogic® Log Source Packages are available on the TIBCO eDelivery website or TIBCO Support website after logging in.

- *TIBCO LogLogic® Log Source Packages Release Notes*
- *TIBCO LogLogic® Log Source Packages Installation and Upgrade*
- *TIBCO LogLogic® Log Source Packages Log Configuration Guides*
- *TIBCO LogLogic® Log Source Packages Log Collector Guides*

How to Contact TIBCO Support

Get an overview of TIBCO Support. You can contact TIBCO Support in the following ways:

- For accessing the Support Knowledge Base and getting personalized content about products you are interested in, visit the TIBCO Support website.
- For creating a Support case, you must have a valid maintenance or support contract with TIBCO. You also need a user name and password to log in to TIBCO Support website. If you do not have a user name, you can request one by clicking **Register** on the website.

How to Join TIBCO Community

TIBCO Community is the official channel for TIBCO customers, partners, and employee subject matter experts to share and access their collective experience. TIBCO Community offers access to Q&A forums, product wikis, and best practices. It also offers access to extensions, adapters, solution accelerators, and tools that extend and enable customers to gain full value from TIBCO products. In addition, users can submit and vote on feature requests from within the TIBCO Ideas Portal. For a free registration, go to TIBCO Community.

Legal and Third-Party Notices

SOME TIBCO SOFTWARE EMBEDS OR BUNDLES OTHER TIBCO SOFTWARE. USE OF SUCH EMBEDDED OR BUNDLED TIBCO SOFTWARE IS SOLELY TO ENABLE THE FUNCTIONALITY (OR PROVIDE LIMITED ADD-ON FUNCTIONALITY) OF THE LICENSED TIBCO SOFTWARE. THE EMBEDDED OR BUNDLED SOFTWARE IS NOT LICENSED TO BE USED OR ACCESSED BY ANY OTHER TIBCO SOFTWARE OR FOR ANY OTHER PURPOSE.

USE OF TIBCO SOFTWARE AND THIS DOCUMENT IS SUBJECT TO THE TERMS AND CONDITIONS OF A LICENSE AGREEMENT FOUND IN EITHER A SEPARATELY EXECUTED SOFTWARE LICENSE AGREEMENT, OR, IF THERE IS NO SUCH SEPARATE AGREEMENT, THE CLICKWRAP END USER LICENSE AGREEMENT WHICH IS DISPLAYED DURING DOWNLOAD OR INSTALLATION OF THE SOFTWARE (AND WHICH IS DUPLICATED IN THE LICENSE FILE) OR IF THERE IS NO SUCH SOFTWARE LICENSE AGREEMENT OR CLICKWRAP END USER LICENSE AGREEMENT, THE LICENSE(S) LOCATED IN THE “LICENSE” FILE(S) OF THE SOFTWARE. USE OF THIS DOCUMENT IS SUBJECT TO THOSE TERMS AND CONDITIONS, AND YOUR USE HEREOF SHALL CONSTITUTE ACCEPTANCE OF AND AN AGREEMENT TO BE BOUND BY THE SAME.

TIBCO, the TIBCO logo, the TIBCO O logo, and LogLogic are either registered trademarks or trademarks of TIBCO Software Inc. in the United States and/or other countries.

All other product and company names and marks mentioned in this document are the property of their respective owners and are mentioned for identification purposes only.

This software may be available on multiple operating systems. However, not all operating system platforms for a specific software version are released at the same time. See the readme file for the availability of this software version on a specific operating system platform.

THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

THIS DOCUMENT COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION HEREIN; THESE CHANGES WILL BE INCORPORATED IN NEW EDITIONS OF THIS DOCUMENT. TIBCO SOFTWARE INC. MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S) AND/OR THE PROGRAM(S) DESCRIBED IN THIS DOCUMENT AT ANY TIME.

THE CONTENTS OF THIS DOCUMENT MAY BE MODIFIED AND/OR QUALIFIED, DIRECTLY OR INDIRECTLY, BY OTHER DOCUMENTATION WHICH ACCOMPANIES THIS SOFTWARE, INCLUDING BUT NOT LIMITED TO ANY RELEASE NOTES AND "READ ME" FILES.

This and other products of TIBCO Software Inc. may be covered by registered patents. Please refer to TIBCO's Virtual Patent Marking document (https://www.tibco.com/patents) for details.