VMware vRealize Log Insight Getting Started Guide

vRealize Log Insight 2.5

This document supports the version of each product listed and supports all subsequent versions until the document is replaced by a new edition. To check for more recent editions of this document, see http://www.vmware.com/support/pubs.

EN-001658-00

vmware®

You can find the most up-to-date technical documentation on the VMware Web site at:

http://www.vmware.com/support/

The VMware Web site also provides the latest product updates.

If you have comments about this documentation, submit your feedback to:

docfeedback@vmware.com

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Contents

VMware vRealize Log Insight Getting Started Guide

The *VMware vRealize Log Insight Getting Started Guide* provides information about deploying and configuring VMware® vRealize™ Log Insight™, including how to size the Log Insight virtual appliance to receive log messages from your environment.

Intended Audience

This information is intended for anyone who wants to install, configure, or maintain Log Insight. The information is written for experienced Linux system administrators who are familiar with virtual machine technology and datacenter operations.

1 Before You Install Log Insight

To start using Log Insight in your environment, you must deploy the Log Insight virtual appliance and apply several basic configurations.

This chapter includes the following topics:

Supported Log Files and Archive Formats in Log Insight

You can use Log Insight to analyse any unstructured or structured log files.

If the source host has the capability to send log events over either the syslog protocol or HTTP, Log Insight can accept the data. To analyze historic data you can import log files that were archived by Log Insight.

See the topic Import a Log Insight Archive into Log Insight in the *VMware vCenter Log Insight Administration Guide*.

NOTE Although Log Insight can handle historic data and real-time data simultaneously, you are advised to deploy a separate instance of Log Insight to process imported log files.

Security Requirements

To ensure that your virtual environment is protected from external attacks, you must observe certain rules.

- Always install Log Insight in a trusted network.
- Always save Log Insight support bundles in a secure location.

For information about securing your virtual environment, see the *VMware vSphere Security Guide* and the Security Center on the VMware Web site.

Product Compatibility

Log Insight collects data over the syslog protocol and HTTP, can connect to vCenter Server to collect events, tasks, and alarms data, and can integrate with vRealize Operations Manager to send notification events and enable launch in context . Check the *VMware vRealize Log Insight Release Notes* for latest updates on supported product versions.

Virtual Appliance Deployment

You must and deploy the Log Insight virtual appliance using vSphere. Always use a vSphere Client to connect to a vCenter Server. The Log Insight virtual appliance should be deployed on an ESX/ESXi host version 4.1 or later that is managed by vCenter Server version 4.1 or later.

Syslog Feeds

Log Insight collects and analyses syslog data over the following ports and protocols.

- 514/UDP
- 514/TCP
- 1514/TCP (SSL)

You must configure environment components such as operating systems, applications, storage, firewalls, and network devices to push their syslog feeds to Log Insight.

API Feeds

The Log Insight Ingestion API collects data over the following port and protocol.

- 9000/TCP

vSphere Integration

You can configure Log Insight to pull data for tasks, events, and alarms that occurred in one or more vCenter Server instances. Log Insight uses the vSphere API to connect to vCenter Server systems and collect data.

You can configure ESXi hosts to forward syslog data to Log Insight.

Table 1-1. Supported vSphere Product Versions

Type of Integration	Supported Product Versions
Tasks, events, and alarms data collection	vCenter Server 5.1 and later
Syslog feeds	ESXi 5.x and later

See topic *Connect Log Insight to vCenter Server 5.1.x Systems* in the Log Insight Administration Guide.

vRealize Operations Manager Integration

Log Insight and vRealize Operations Manager vApp or Installable can be integrated in two independent ways.

- Log Insight can send notification events to vRealize Operations Manager.

  See the topic *Configure Log Insight to Send Notification Events to vRealize Operations Manager* in the Log Insight Administration Guide.

- The launch in context menu of vRealize Operations Manager can display actions related to Log Insight.

  See topic *Enable Launch in Context for Log Insight in vRealize Operations Manager* in the Log Insight Administration Guide.

The following table contains the versions of vRealize Operations Manager that support notifications and launch in context.

Product Deliverable	Notification Events	Launch in Context
vRealize Operations Manager vApp	■ vSphere UI 5.7.1 and later ■ Custom UI 5.6 and later	5.7.1 and later
vRealize Operations Manager Installable	5.7.0 Hot Fix 1 and later	5.7.1 and later

Minimum Requirements

VMware distributes Log Insight as a virtual appliance in OVA file format. Various resources and applications must be available for the virtual appliance to run successfully.

Virtual Hardware

During deployment of the Log Insight virtual appliance you can select different sizes according to the ingestion requirements for the environment. An extra small configuration is the smallest supported configuration and can support log volumes of 3GB a day for about 10 users. The extra small configuration requires the following virtual resources.

- 2 vCPUs, 2GHz each
- 4GB RAM
- Approximately 144GB storage space

For complete resources requirements based on ingestion requirements, see "Sizing the Log Insight Virtual Appliance," on page 10

Supported Browsers

You can use one of the following browsers to connect to the Log Insight Web user interface.

IMPORTANT Cookies must be enabled in your browser.

- Mozilla Firefox 10.x, 19.x, 20.x, 21.0, 23 and 29.0.1
- Safari 6.0 , 7.0.2
- Google Chrome 25.x, 26.x, 27.x, 29 and 34
- Internet Explorer 10.x and 11.x

NOTE Internet Explorer Document mode must be set to **Standards Mode**. Other modes are not supported. **Browser Mode:** Compatibility View is not supported.

Required Network Ports

The following network ports must be externally accessible.

Port	Protocol
80/TCP	HTTP
443/TCP	HTTPS

Port	Protocol
514/UDP, 514/TCP	Syslog
1514/TCP	Syslog
9000/TCP	Log Insight Ingestion API

Sizing the Log Insight Virtual Appliance

By default, the Log Insight virtual appliance has 2 vCPUs, 4GB of virtual memory, and 144GB of disk space provisioned. Log Insight uses 100GB of the disk space to store raw data, index, metadata, and so on.

Standalone Deployment

You can change the settings according to the environment for which you intend to collect logs.

During the virtual appliance deployment, you can select the size of the appliance as follows.

Option	Log Ingest Rate	vCPUs	Memory	IOPS	Syslog Connections	Events per Second
Extra Small	3GB/day	2	4GB	75	20	200
Small	15GB/day	4	8GB	500	100	1000
Medium	37.5GB/day	8	16GB	1000	250	2500
Large	112.5GB/day	16	32GB	1500	750	7500

NOTE You can use a syslog aggregator to increase the number of syslog connections that send events to Log Insight. However, the maximum number of events per second is fixed and does not depend on the use of a syslog aggregator. A Log Insight instance cannot be used as a syslog aggregator.

The sizing is based on the following assumptions.

- Each vCPU is at least 2GHz.
- Each ESXi host sends up to 10 messages per second with an average message size of 170 bytes/message. This is roughly equivalent to 150MB/day/host.

NOTE For large installations, you must upgrade the virtual hardware version of the Log Insight virtual machine. Log Insight supports virtual hardware version 7 or later. Virtual hardware version 7 can support up to 8 vCPUs. Therefore, you must upgrade to virtual hardware version 8 or later for ESXi 5.x if you plan to provision 16 vCPUs. You use the vSphere Client to upgrade the virtual hardware. If you want to upgrade virtual hardware to the latest version, read and understand the information in the VMware knowledge base article Upgrading a virtual machine to the latest hardware version (1010675) .

Cluster Deployment

Use a large configuration for the master and worker nodes in a Log Insight cluster. The number of events per second increases linearly with the number of nodes.

Reducing the Memory Size

If you want to use the **Extra Small** version of the appliance on your laptop, but the laptop does not have enough memory, you can reduce the memory size to 2GB.

Installing Log Insight 2

Log Insight is delivered as a virtual appliance that you must deploy in your environment.

To deploy the Log Insight virtual appliance, follow the standard OVF deployment procedure.

This chapter includes the following topics:

Deploy the Log Insight Virtual Appliance

Download the Log Insight virtual appliance. VMware distributes the Log Insight virtual appliance as an `.ova` file. Deploy the Log Insight virtual appliance by using the vSphere Client.

Prerequisites

- Verify that you have a copy of the Log Insight virtual appliance `.ova` file.
- Verify that you have permissions to deploy OVF templates to the inventory.
- Verify that your environment has enough resources to accommodate the minimum requirements of the Log Insight virtual appliance. See the topic *Minimum Requirements* in the *VMware vCenter Log Insight Getting Started Guide*.
- Verify that you read and understand the virtual appliance sizing recommendations. See the topic *Sizing the Log Insight Virtual Appliance* in *VMware vCenter Log Insight Getting Started Guide*.

Procedure

1. In the vSphere Client, select **File > Deploy OVF Template**.
2. Follow the prompts in the Deploy OVF Template wizard.
3. On the Deployment Configuration page, select the size of the Log Insight virtual appliance based on the size of the environment for which you intend to collect logs.

   Small is the minimum requirement for production environments.

Option	Log Ingest Rate	vCPUs	Memory	IOPS	Syslog Connections	Events per Second
Extra Small	3GB/day	2	4GB	75	20	200
Small	15GB/day	4	8GB	500	100	1000

Option	Log Ingest Rate	vCPUs	Memory	IOPS	Syslog Connections	Events per Second
Medium	37.5GB/day	8	16GB	1000	250	2500
Large	112.5GB/day	16	32GB	1500	750	7500

NOTE If you select **Large**, you must upgrade the virtual hardware on the Log Insight virtual machine after the deployment.

4 On the Disk Format page, select a disk format.

- **Thick Provision Lazy Zeroed** creates a virtual disk in a default thick format. Space required for the virtual disk is allocated when the virtual disk is created. The data remaining on the physical device is not erased during creation, but is zeroed out on demand at a later time, on first write from the virtual appliance.
- **Thick Provision Eager Zeroed** creates a type of thick virtual disk that supports clustering features such as Fault Tolerance. Space required for the virtual disk is allocated at creation time. In contrast to the flat format, the data remaining on the physical device is zeroed out when the virtual disk is created. it might take much longer to create disks in this format than to create other types of disks.

  IMPORTANT Deploy the Log Insight virtual appliance with thick provisioned eager zeroed disks whenever possible for better performance and operation of the virtual appliance.

- **Thin Provision** creates a disk in thin format. The disk grows as the data saved on it grows. If your storage device does not support thick provisioning disks or you want to conserve unused disk space on the Log Insight virtual appliance, deploy the virtual appliance with thin provisioned disks.

NOTE Shrinking disks on the Log Insight virtual appliance is not supported and might result in data corruption or data loss.

5 (Optional) On the Properties page, set the networking parameters for the Log Insight virtual appliance.

If you do not provide network settings, such as IP address, DNS servers, and gateway, Log Insight utilizes DHCP to set those settings.

CAUTION Do not specify more than two domain name servers. If you specify more than two domain name servers, all configured domain name servers are ignored in the Log Insight virtual appliance.

Use comma to separate domain name servers.

6 (Optional) On the Properties page, set the root password for the Log Insight virtual appliance.

7 Follow the prompts to complete the deployment.

For information on deploying virtual appliances, see the *User's Guide to Deploying vApps and Virtual Appliances*.

After you power on the virtual appliance, an initialization process begins. The initialization process takes several minutes to complete. At the end of the process, the virtual appliance restarts.

8 Navigate to the **Console** tab and check the IP address of the Log Insight virtual appliance.

IP Address Prefix	Description
https://	The DHCP configuration on the virtual appliance is correct.
http://	The DHCP configuration on the virtual appliance failed. a Power off the Log Insight virtual appliance. b Right-click the virtual appliance and select **Edit Settings**. c Set a static IP address for the virtual appliance.

What to do next

- To enable SSH connections to the Log Insight virtual appliance, configure the root password in the virtual appliance console. See topic *Configure the Root SSH Password for the Log Insight Virtual Appliance* in the *Log Insight Administration Guide*.
- If you want to configure a standalone Log Insight deployment, see topic *Configure New Log Insight Deployment* in the *Log Insight Getting started Guide*.

  The Log Insight Web interface is available at https://*log-insight-host*/ where *log-insight-host* is the IP address or host name of the Log Insight virtual appliance.

Start New Standalone Deployment

When you access the Log Insight Web interface for the first time after the virtual appliance deployment or after removing a worker node from a cluster, you must complete the initial configuration steps.

All settings that you modify during the initial configuration are also available in the Administration Web user interface.

For information about the trace data that Log Insight might collect and send to VMware if you choose to participate in the Customer Experience Improvement Program, see Chapter 3, "The Customer Experience Improvement Program," on page 17.

Prerequisites

- In the vSphere Client, note the IP address of the Log Insight virtual appliance. For more information, see "Start New Standalone Deployment," on page 13
- Verify that you are using a supported browser, see the VMware vSphere Documentation.
- Verify that you have a valid license key. You can request an evaluation or permanent license key by using your account to My VMware™ https://my.vmware.com/ .
- If you want to use local, vCenter Server, or Active Directory credentials to integrate Log Insight with vRealize Operations Manager, verify that these users are imported in vRealize Operations Manager Custom user interface. For instructions about configuring LDAP, see the vRealize Operations Manager Administration Guide.

Procedure

1. Use a supported browser to navigate to the Web user interface of Log Insight.

   The URL format is https://*log_insight-host*/, where *log_insight-host* is the IP address or host name of the Log Insight virtual appliance.

   The initial configuration wizard opens.
2. Click **Start New Deployment**.
3. Set the password for the Admin user and click **Save and Continue**.

   Optionally, you can provide an email address for the admin user.
4. Enter the license key, click **Set Key**, and click **Continue**.
5. On the General Configuration page, type the email address to receive system notifications from Log Insight.
6. To opt out of the Customer Experience Improvement Program, unselect **Send weekly Trace Data to VMware as part of the Customer Experience Improvement Program**. Click **Save and Continue**.

7 On the Time Configuration page, set how time is synchronized on the Log Insight virtual appliance and click **Test**.

Option	Description
NTP server (recommended)	By default, Log Insight is configured to synchronize time with public NTP servers. If an external NTP server is not accessible due to firewall settings, you can use the internal NTP server of your organization. Use commas to separate multiple NTP servers.
ESX/ESXi host	If no NTP servers are available, you can sync the time with the ESXi host where you deployed the Log Insight virtual appliance.

8 Click **Save and Continue**.

9 (Optional) To allow active directory users access in Log Insight, select **Enable Active Directory support** and click **Save and Continue**.

NOTE Active Directory is enabled for Log Insight only and does not affect the integration of Log Insight with other VMware products.

10 Specify the properties of an SMTP server to enable outgoing alert and system notification emails.

To verify that the SMTP configuration is correct, type a valid email address and click **Test**. Log Insight sends a test email to the address that you provided.

11 Click **Save and Continue**.

12 (Optional) Configure the integration between Log Insight and vRealize Operations Manager.

a Type the host name and user credentials for the UI VM of the vRealize Operations Manager vApp, and click **Test Connection** to verify the connection.

NOTE You must provide the user credentials of a vRealize Operations Manager administrator user.

b To enable Log Insight to send alert notifications triggered by Log Insight alarms, select **Enable alerts integration**.

c To allow vRealize Operations Manager to launch Log Insight with an object-specific query, click **Enable Launch in Context**.

This operation might take a few minutes.

13 Click **Save and Continue**.

14 (Optional) To archive log data to an NFS location, select **Enable Data Archiving**, enter the path to the storage location, and click **Test** to verify that Log Insight can connect to that storage.

15 Click **Save and Continue**.

16 Click **Restart** to complete the initial setup of Log Insight.

After the Log Insight process restarts, you are redirected to the **Dashboards** tab of Log Insight.

What to do next

- Go to the **Administration** page by selecting the drop-down menu icon  in the navigation bar and use the **vSphere Integration** page to configure Log Insight to pull tasks, events, and alerts from vCenter Server instances, and to configure ESXi hosts to send syslog feeds to Log Insight.
- See the topic *Assign a Permanent License to Log Insight* in the *Log Insight Administration Guide*.
- See the topic *Install the Log Insight Adapter in vRealize Operations Manager Standalone* in the *Log Insight Administration Guide*.

- Install the Log Insight Windows Agent to collect events from Windows event channels, Windows directories, and flat text log files. See the topic *Installing the Log Insight Windows Agent as a Windows Service* in the *Log Insight Administration Guide*.

Join Existing Deployment

After you deploy and set up a standalone Log Insight node, you can deploy a new Log Insight instance and add it to the existing node to form a Log Insight cluster.

Log Insight can scale out by using multiple virtual appliance instances. This enables linear scaling of the ingestion throughput, increases query performance and allows for ingestion high availability. In cluster mode, Log Insight provides master and worker nodes. Both master and worker nodes are responsible for a subset of data. Master nodes can query all subsets of data and aggregate the results.

IMPORTANT It is highly recommended that you configure a minimum of three nodes in a Log Insight cluster to provide ingestion, configuration, and user space High Availability.

Prerequisites

- In the vSphere Client, note the IP address of the worker Log Insight virtual appliance.
- Verify that you have the IP address or host name of the master Log Insight virtual appliance.
- Verify that you have an administrator account on the master Log Insight virtual appliance.
- Verify that the versions of the Log Insight master and worker nodes are in sync. Do not add an older version Log Insight worker to a newer version Log Insight master node.
- You must synchronize the time on the Log Insight Linux Agent virtual appliance with an NTP server. See the topic *Synchronize the Time on the Log Insight Virtual Appliance* in the *Log Insight Administration Guide*.
- For information on supported browser versions, see the topic *Minimum Requirements* in the *VMware vCenter Log Insight Getting Started Guide*.

Procedure

1 Use a supported browser to navigate to the Web user interface of the Log Insight worker.

  The URL format is https://*log_insight-host*/, where *log_insight-host* is the IP address or host name of the Log Insight worker virtual appliance.

  The initial configuration wizard opens.

2 Click **Join Existing Deployment**.

3 Enter the IP address or host name of the Log Insight master and click **Go**.

  The worker sends a request to the Log Insight master to join the existing deployment.

4 Click the **Click here to access the Cluster Management page** link.

5 Log in as an administrator.

  The Cluster page loads.

6 Click **Allow**.

  The worker joins the existing deployment and Log Insight begins to operate in a cluster.

What to do next

- To add another worker, deploy a new Log Insight instance and add it to the cluster using the startup wizard.

- Repeat the procedure to add a minimum of two Log Insight worker nodes.

The Customer Experience Improvement Program 3

You can configure Log Insight to collect data to help improve your user experience with VMware products. The following section contains important information about the Customer Experience Improvement Program.

The goal of the Customer Experience Improvement Program is to quickly identify and address problems that might be affecting your experience. If you choose to participate in the VMware Customer Experience Improvement Program, Log Insight will regularly send encrypted trace data to VMware. You can use trace data for product development and troubleshooting purposes. Log Insight anonymizes and encrypts any personal identification information from your systems or servers before transferring any trace data to VMware.

If you have any questions or concerns regarding the Customer Experience Improvement Program for Log Insight, contact li-info@vmware.com.

Trace Data that Log Insight Collects

To provide the benefits of the Customer Experience Improvement Program, Log Insight collects trace data directly from log files stored on your Log Insight virtual appliance and transfers the data to VMware on a weekly basis.

Categories of Information in Trace Data

Trace data contains the following categories of information.

runtime.log Contains information about low-level system trace activities conducted by Log Insight, including indexing, garbage collection, and monitoring activities. If an error occurs while Log Insight is processing data or a query, information about the error appears in the `runtime.log` file.

ui.log Contains information regarding interactions with user interface components and parameters, such as which buttons were pressed or which options were selected from a drop-down menu.

usage.log Contains information regarding the queries that the query engine runs. Each line has the exact query that the search engine runs, including the time it was started, the length of time it ran, and if an error occurred during its execution.

watchdog.log Contains information from the watchdog process that monitors Log Insight and restarts the application if it fails or becomes unresponsive. The `watchdog.log` file contains information documenting when such failures are detected.

vcenter-operations.log	Contains infromation regarding the integration between Log Insight and vRealize Operations Manager, sending query alerts to vRealize Operations Manager and registering Log Insight with vRealize Operations Manager for launch in context.
li-vsphere.log	Contains information regarding the integration between Log Insight and vSphere.
alert.log	Contains information about user defined alerts that have been triggered.
upgrade.log	Contains information about events that occur during Log Insight upgrade.
loginsight_daemon_stdout.log	Contains information about the standard output of Log Insight daemon.
systemalert.log	Contains information about system alerts that Log Insight sends.
watchdog_log*	Contains information about the run time events of the watch dog process, which is responsible for restarting Log Insight if it is shutdown for some reason.

Personal Information in Trace Data

Trace data can also contain personal information, including:

- Email addresses
- MAC addresses
- Internet protocol addresses
- User names
- Host names
- Query content
- Search word content

Personal information found inside trace data files is anonymized and encrypted inside your Log Insight virtual appliance before being transferred to VMware. Trace data is encrypted using public key cryptography and sent through email using your SMTP server. Trace data is stored in the VMware internal secured network and is not shared with third parties.

You can view the files at any time by remotely logging in to your Log Insight virtual appliance using SSH, and navigating to `/storage/var/loginsight`.

You can stop the transfer of trace data to VMware at any time. See the topic *Stop Sending Trace Data to VMware* in the *Log Insight Administration Guide*.

If you have any questions or concerns regarding the Customer Experience Improvement Program for Log Insight, contact li-info@vmware.com.

Index