

Using vRealize Log Insight Importer

April 14, 2020

vRealize Log Insight 8.1



vmware®

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

<https://docs.vmware.com/>

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

Copyright © 2020 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

| | | |
|----------|--|----------|
| 1 | Using the vRealize Log Insight Importer | 4 |
| | Installing the vRealize Log Insight Importer | 4 |
| | Before You Install the vRealize Log Insight Importer | 5 |
| | Install the vRealize Log Insight Importer | 5 |
| | Running the vRealize Log Insight Importer | 6 |
| | About the vRealize Log Insight Importer Manifest File | 6 |
| | vRealize Log Insight Importer Manifest File Configuration Examples | 7 |
| | Run the vRealize Log Insight Importer | 8 |

Using the vRealize Log Insight Importer

1

The vRealize Log Insight Importer is a command-line utility used to import offline logs of historical data from local machines to the vRealize Log Insight server.

Use the Importer when you want to import logs that were collected in the past. You can import support bundles and archived logs and analyze logs from support bundles gathered from vRealize Log Insight or any VMware product.

vRealize Log Insight Importer includes the following features and capabilities.

- vRealize Log Insight Importer sends data over the ingestion API.
- It supports file log collection, including recursive directory collection.
- The Importer can read data from zip, tar, bzip, bzip2, or gz archive files. 7-Zip is not supported.
- You can stipulate that data be read recursively from a nested archive, such as a nested ZIP file, or from a directory within an archive.

This chapter includes the following topics:

- [Installing the vRealize Log Insight Importer](#)
- [Running the vRealize Log Insight Importer](#)

Installing the vRealize Log Insight Importer

You install the vRealize Log Insight Importer from an installation package you obtain from the VMware Download site. The installation packages include the MSI installer for Windows and POSIX installation packages (RPM, DEB and BIN) for Linux.

- [Before You Install the vRealize Log Insight Importer](#)

Check requirements and understand importer behavior before you install the importer.

- [Install the vRealize Log Insight Importer](#)

You can install vRealize Log Insight Importer on Windows and Linux. You can also install the vRealize Log Insight Importer on a vRealize Log Insight server and run it from the server.

Before You Install the vRealize Log Insight Importer

Check requirements and understand importer behavior before you install the importer.

Before you install, ensure that vRealize Log Insight has access to the NFS server on which archived data is stored. If the NFS server becomes inaccessible due to network failure or errors on the NFS server, importation of archived data might fail.

When logs are extracted from a bundle during ingestion, a log bundle name is automatically determined and added as a bundle tag to all extracted logs. The tag name is the filename of the log or the directory name in case of directory sources. Bundle tags differentiate bundles on a vRealize Log Insight server.

This tag overrides any tags with the same name that are specified in the manifest file. The tag can be overridden by command line tags that use the same name.

When you use the importer, be aware of the following behaviors:

- vRealize Log Insight Importer does not check for available disk space on the vRealize Log Insight virtual appliance. Therefore, the import of archived logs might fail if the virtual appliance runs out of disk space.
- vRealize Log Insight does not display progress information during log imports. As the import of archived data is in progress, you are unable to infer from the console output how much time is left before the import finishes or how much data is already imported.

Supported Operating Systems

The vRealize Log Insight Importer is supported on the following operating systems:

- Windows 32 bit and 64 bit
- Linux 32 bit and 64 bit

The Linux version does not run on an Apple Macintosh system.

Install the vRealize Log Insight Importer

You can install vRealize Log Insight Importer on Windows and Linux. You can also install the vRealize Log Insight Importer on a vRealize Log Insight server and run it from the server.

When you install vRealize Log Insight Importer, several VMware product manifest files are also installed. You can use these files or modify them for your needs when running vRealize Log Insight Importer. These manifest files are located in C:\Program Files (x86)\VMware\Log Insight Importer\Manifests for Windows, and /usr/lib/loginsight-importer/manifests for Linux.

If you uninstall the .bin package, also delete the /usr/bin/loginsight_importer symlink.

Prerequisites

- Verify that you can access the [VMware Download](#) site to download the vRealize Log Insight Importer.

Procedure

- 1 Download the vRealize Log Insight Importer installation package from the [VMware Download](#) site.

The installation packages include the MSI installer for Windows and POSIX installation packages (RPM, DEB, and BIN) for Linux.

- 2 Install the tool on your system.

After installation, the importer installation directory is added to the PATH environment variable on Windows, and a symlink to the executable file `loginsight-importer` is added to `/usr/bin/` on Linux. So the client can call `loginsight-importer` from the shell without specifying a path prefix.

The vRealize Log Insight Importer tool is installed in the following locations.

| Operating System | Filename | Installation Location |
|------------------|-------------------------|--|
| Windows | loginsight-importer.exe | C:\Program Files (x86)\VMware\Log Insight Importer |
| Linux | loginsight-importer | /usr/lib/loginsight-importer |

Running the vRealize Log Insight Importer

When you run the importer, you must include a manifest file. The manifest file provides information about log format, the location of the data to import, and source and destination information.

- [About the vRealize Log Insight Importer Manifest File](#)

vRealize Log Insight Importer uses a manifest configuration file to determine the log format and to specify the location of the data to import. The manifest file has the same format as the `liagent.ini` configuration file and is similar in structure.

- [vRealize Log Insight Importer Manifest File Configuration Examples](#)

The sample vRealize Log Insight Importer manifest files provide examples of parameter configurations.

- [Run the vRealize Log Insight Importer](#)

Run the vRealize Log Insight Importer to import offline logs of historical data to the vRealize Log Insight server.

About the vRealize Log Insight Importer Manifest File

vRealize Log Insight Importer uses a manifest configuration file to determine the log format and to specify the location of the data to import. The manifest file has the same format as the `liagent.ini` configuration file and is similar in structure.

Optionally, you can create your own manifest file to import arbitrary log files. One advantage of creating such a file is that you do not need to know the absolute path to the data files.

If you do not create a manifest file, vRealize Log Insight Importer uses the default manifest, which collects all .txt and .log files (include=*.log*;*.txt*) and applies the auto parser (extracts timestamp + kvp) on the extracted logs.

If the liagent.ini configuration file is used as a manifest file, vRealize Log Insight Importer extracts only the [filelog] sections as a manifest. All options for the [filelog] section are supported in vRealize Log Insight Importer.

For information about options supported in the [filelog] section and configuration examples, see the topic "Collect Events from a Log File" in the *Working with vRealize Log Insight Agents*.

To Create a Manifest File

You can copy and paste the contents of the agent configuration file into a new TXT file. To identify a dynamic path, remove the leading " / " before the directory path.

Specifying the Directory Path

The directory specified in the [filelog] section can be either relative to the source or absolute. To specify a relative path, do not include the leading slash under Linux, otherwise vRealize Log Insight Importer treats the path as absolute.

To indicate name patterns in the value of the directory key, you can use the * and ** characters.

- Use * as a placeholder for a single directory. Use it to indicate one level of nesting with an arbitrary folder name. For example, use directory = log_folder_* to indicate any folder that starts with the string log_folder_.
- Use ** to indicate an arbitrary level of nesting with any folder name. For example, you can use directory = **/log to indicate any folder with the name log at any level of nesting within the source directory.

vRealize Log Insight Importer Manifest File Configuration Examples

The sample vRealize Log Insight Importer manifest files provide examples of parameter configurations.

The value of the directory key must be either relative to the source or absolute. The following example shows how to collect logs from files with a .log extension which reside two levels lower than the source directory and name of the last folder ends with the _log string.

```
[filelog]importer_test]
directory=*\*_log
include=*.log
event_marker=^\\d{4}-\\d{2}-\\d{2} \\d{2}:\\d{2}:\\d{2} [A-Z]{4} LOG
```

The following example shows how to collect all files with the extension .log from all subfolders of the source directory including the source itself.

```
[filelog]sbimporter_test_channel]
directory = **
```

```
include = *.log
```

The following example shows how to collect logs from all files in the source directory (but not from subfolders) except files that have an `.ini` extension. We interpret files as UTF-16LE encoded.

```
[filelog|quotes_channel3]
directory=
charset=UTF-16LE
exclude=*.ini
tags={"Provider" : "Apache"}
```

The following example shows how to collect logs from all files with the extension `.log` in the source directory (but not from subfolders). The timestamp of events is parsed in the log file using the Common Log Format (CLF) parser and the extracted historical timestamp is applied. The log format parsed by the CLF parser is `2015-03-25 22:11:46,786 | DEBUG | pool-jetty-76 | AuthorizationMethodInterceptor | Authorizing method: public abstract.`

```
[filelog|vcd-container-debug]
directory=
include=*.log
parser=vcd

[parser|vcd]
base_parser=clf
format=%Y-%m-%d %H:%M:%S%f}t %M
```

Run the vRealize Log Insight Importer

Run the vRealize Log Insight Importer to import offline logs of historical data to the vRealize Log Insight server.

Prerequisites

- Review [About the vRealize Log Insight Importer Manifest File](#) and create a manifest file to use with the importer. For more information, see [vRealize Log Insight Importer Manifest File Configuration Examples](#).
- If you use the `honor_timestamp` parameter, verify that you have appropriate login credentials.
- If you import a support bundle, configure `honor_timestamp` and the user name and password.

Procedure

- 1 Start the vRealize Log Insight Importer tool by entering the following command at a command prompt.

```
/usr/bin/loginsight-importer.exe
```

- 2 Enter the manifest file name at the prompt.

3 Define the configuration parameters and press **Enter**.

The `--source` and `--server` parameters are required.

| Required Parameters | Description |
|---|--|
| <code>--source <path></code> | Specifies the path to a support bundle directory or path to a zip, gzip, bzip, bzip2, or tar archive. The value is added to all send messages as the value of the <code>bundle</code> tag. |
| <code>--server <hostname></code> | Destination server hostname or IP address. |
| Options | Description |
| <code>--port <port></code> | Port for connection. If not set then port 9000 is used for non-SSL connections and port 9543 is used for an SSL connection. |
| <code>--logdir <path></code> | Specifies the path to the logs directory. If this is not set, the path is: <code>\$(LOCALAPPDATA)\VMware\Log Insight Importer\log</code> on Windows and <code>~/loginsight-importer/log</code> on Linux. |
| <code>--manifest <file-path></code> | Specifies the path to the manifest file (.ini format). If this is not set, the <code>importer.ini</code> file in the source directory is used. If the <code>importer.ini</code> file does not exist or is not found in the source directory, vRealize Log Insight Importer applies the default (hardcoded) manifest and collects all .txt and .log files (<code>include=*.log*;*.txt*</code>), and also applies the auto parser (extracts timestamp + kvp). |
| <code>--no_ssl</code> | Do not use SSL for connections. This should not be set for authenticated connections (for example if <code>--honor_timestamp</code> is used). |
| <code>--ssl_ca_path <path></code> | Path to the trusted root certificates bundle file. |
| <code>--tags <tags></code> | Set tags for all sent events. For example <code>--tags "{ \"tag1\" : \"value1\", \"tag2\":\"value2\"}"</code> Note The tags option can accept hostname as a tag name. The value of the hostname tag from the command line is used instead of the FQDN of the sending machine as the value of the hostname field for all events extracted by vRealize Log Insight Importer. This is opposite of the tags parameter in the manifest file and extracted fields by parsers, which ignore the hostname field. A log bundle name, either a filename or a directory name in case of directory sources, is automatically determined and added as a <code>bundle</code> tag to all logs extracted from that specific bundle during the ingestion. This tag helps you to differentiate bundles on vRealize Log Insight Server. A <code>bundle</code> tag overrides tags with that same name from a manifest file. But it can be overridden by command line tags, if there is one with <code>bundle</code> name. |
| <code>--username <username ></code> | Username for authentication. Required if <code>--honor_timestamp</code> is set. |
| <code>--password <password></code> | Password for authentication. Required if <code>--honor_timestamp</code> is set. The username/password pair disables the allowed time-drift on vRealize Log Insight server so it is possible to import data with a historical timestamp. |

| Options | Description |
|--|---|
| <code>--honor_timestamp</code> | <p>Applies the extracted timestamp. The configured parsers extract the timestamp from the log entries and the <code>--honor_timestamp</code> applies the extracted timestamp.</p> <ul style="list-style-type: none"> ■ If the timestamp is extracted using configured parsers, then the events will have that timestamp applied. ■ If there is an event in the logs file, with no extracted timestamp, then the successfully extracted timestamp from the previous event in the same log file will be applied. ■ If no timestamp is found or parsed in the file then the MTIME of the log file will be applied as the timestamp. <p>Note If a manifest file was not provided, the default hardcoded manifest that the vRealize Log Insight Importer will use has the Automatic Log parser enabled. In this case, vRealize Log Insight Importer extracts the timestamp from the log entries if the <code>--honor_timestamp</code> parameter is used.</p> |
| <code>--debug_level <1 2></code> | Increases the verbosity level of the log file. This should only be changed when troubleshooting. Under normal operations this flag should not be used. |
| <code>--help</code> | Display help and exit. |

4 After the import is complete, press **Ctrl+C** on Windows or Linux to exit the tool.

Results

vRealize Log Insight Importer extracts the log entries from the directories specified in the parameters. The total number of processed files, extracted log messages, sent log messages, and the run time is displayed.

What to do next

From the vRealize Log Insight Interactive Analytics tab, you can refresh the view to list the imported log events. If you imported a support bundle and used the `honor_timestamp`, the dashboard should also display the events over time.