IBM Security QRadar
Version 7.1.0 (MR1)

Forwarding Logs Using Tail2Syslog Technical Note

Note: Before using this information and the product that it supports, read the information in Notices and Trademarks on page 9.

CONTENTS

1 FORWARDING LOGS USING TAIL2SYSLOG

The Tail2Syslog support script provides a method for monitoring and forwarding events to QRadar SIEM using syslog for real-time correlation. Tail2Syslog events forwarded to QRadar SIEM are intended for use with the universal DSM. Tail2Syslog operates by monitoring for a file matching a directory and file pattern (globbing pattern) from a configuration file you create. You can monitor a directory where a device is creating and appending to new log files by setting a date or size archive limit. The file monitored by Tail2Syslog is determined by the last modified date. The most recent log file is monitored until a new file with a more recent modified date is created in the directory matching the file pattern.

Unless otherwise noted, all references to QRadar SIEM refer to QRadar SIEM, IBM Security QRadar Log Manager, and IBM Security QRadar Network Anomaly Detection. References to flows do not apply to QRadar Log Manager.

This section includes the following topics:

Before You Begin

Before you install the Tail2Syslog support script, you must have the following:

- A Linux-based system to host Tail2Syslog with Perl 5.8.8 installed.
- Appropriate access to QRadar SIEM.

  Any firewall between the host system and QRadar SIEM must allow traffic on the syslog ports specified in the Tail2Syslog configuration parameters.

Installing Tail2Syslog

To install Tail2Syslog:

Step 1 Access the Qmmunity website:

https://qmmunity.q1labs.com/

Step 2 From the **Software** tab, select **Scripts**.

Step 3 Download the appropriate version of Tail2Syslog to your Linux-based host system:

```
63-tail2syslog.tar.gz
```

```
70-tail2syslog.tar.gz
```

Step 4 On the system hosting Tail2Syslog, create the following directory:

```
/opt/tail2syslog
```

Step 5 Extract the archive to /opt/tail2syslog/ on the host system:

```
tar -zxvf tail2syslog.tar.gz
```

The archive contains the following files:

- File/Tail.pm
- tail2syslog.pl

In order for the script to run properly, the extracted directory must contain a folder named **File**, which in turn must contain the Perl module Tail.pm.

Step 6 Type the following command to set the proper permissions for the script:

```
chmod +x tail2syslog.pl
```

You are now ready to configure tail2syslog on your host system. For more information, see **Configuring Tail2Syslog**.

Configuring Tail2Syslog

Before you run Tail2Syslog, you must create a configuration file. The configuration file contains the following information:

- The destination addresses of the QRadar SIEM or Event Collectors receiving syslog events.
- The log file directory path and file pattern (globbing pattern) of the file to monitor. Tail2Syslog monitors one file per directory and file pattern at a time.
- The archive directory path, which allows you to move old log files to a different directory.
- The threshold of the number of files to keep in the log file and archive directory.

For example,

```
Destinations=1.1.1.1
Globs=/root/Logs/file1*|/root/Audit/file2*
ArchiveDir=/store/complete/
DeletionThreshold=0
```

To configure Tail2Syslog:

Step 1 Create a configuration file for Tail2Syslog:

```
vi /opt/tail2syslog/example.cfg
```

The file **`example.cfg`** is created.

Step 2 Configure the following parameters in your configuration file:

a **Destinations** - Type an IP address for QRadar SIEM or Event Collectors to receive syslog events.

Multiple IP addresses can be used for load balancing syslog events, you must separate IP addresses using the pipe character (|). If you specify multiple IP addresses, events are distributed evenly between the IP addresses in the configuration file. This does not send all syslog events to multiple locations.

b **Globs** - Type the directory path and file pattern (glob) of the file to monitor.

You can configure the file to monitor multiple directories containing log files. This is in case your device creates folders for log file by date, allowing you to monitor files rolling across multiple directories. You must separate directory locations using the pipe character (|).

We recommend you use a wildcard (*) when specifying the log file to monitor. The wildcard allows you to read the latest file name if the device writing files is rolling log files due to size or day, such as files named **`audit1.log`** or **`audit2.log`**. You can use the configuration file to monitor multiple log files in different directories.

c **ArchiveDir** - Type a directory path for archive files.

The ArchiveDir parameter allows you to specify a path to move old log files to another directory for archiving. You can leave the path information for the ArchiveDir blank if DeletionThreshold=0.

d **DeletionThreshold** - Type a value for the DeletionThreshold.

The DeletionThreshold parameter allows you to specify the maximum number of files to keep in the archive directory by date. If you want to keep the 5 most recent files in the archive directory, then type **`DeletionThreshold=5`**. If you specify zero (0), the files are never deleted from the archive directory.

Step 3 Save the configuration file in VI.

You are now ready to use tail2syslog to monitor files and forward syslog events. For more information, see **Using Tail2Syslog**.

Using Tail2Syslog

To start Tail2Syslog monitoring files specified by your configuration file:

Step 1 Type the following command to run Tail2Syslog:

`./tail2syslog.pl -c <configuration file> <option parameters>`

Where:

`<configuration file>` is the configuration file you created. If your configuration file is not located in the directory containing **`tail2syslog.pl`**, you must type the full directory path to the configuration file.

`<option parameters>` is the list of any optional parameters required for running Tail2Syslog on your system.

The Tail2Syslog script supports several additional option parameters. For more information on optional parameters, see **Table 1-1**.

Table 1-1 Tail2Syslog Optional Parameters

Parameter	Description
-c	The -c parameter specifies the location of the configuration file for Tail2Syslog. Comment lines or blank lines are not allowed between parameters in the configuration file.
-p	The -p parameter specifies the port on the remote host where a syslog receiver is listening. If this parameter is not specified, Tail2Syslog uses TCP port 514 for sending events to QRadar SIEM.
-D	The -D parameter specifies that the script must run in the foreground. If the -D parameter is not specified, then Tail2Syslog runs as a background daemon and logs all internal messages to the local syslog service.
-a	The -a parameter adds a properly formatted syslog header to the message. Tail2Syslog typically sends files as they appear in the unmodified state from the file you are tailing. The -a parameter formats the syslog header of the form <PRI>Mmm dd hh:mm:ss tag. If you do not use the -a parameter, the options -t, -f, and -O have no effect.
-n	The -n parameter appends a new line to the end of the syslog message before forwarding the event.
-t	The -t parameter overrides the default tag name in the optional syslog header (see -a). By default, the tag name is the executable name of the script. The -t parameter overwrites the tag name with the filename from which the message was sent.
-u	The -u parameter overrides the default protocol and forces Tail2Syslog to send events using UDP. The default protocol for sending events is TCP as it ensures reliable delivery and prevents log messages being truncated, which can be the case when using UDP.
-s	The -s parameter sets the event per second (EPS) rate Tail2Syslog uses to forward events. The default rate is 200 EPS.
-f	The -f parameter allows you to add a syslog facility to the header in the syslog message. This parameter must be used in conjunction with the -a parameter. If a facility is not specified, then the default value is user.info.

Table 1-1 Tail2Syslog Optional Parameters (continued)

Parameter	Description
-O	The -O parameter overrides the default hostname in the optional syslog headers. This parameter must be used in conjunction with the -a parameter.
-l	The -l parameter allows you to define a logger for debug information. Your must specify a path and file if you use the -l parameter. For example, **`/bin/logger`**.
-v	The -v parameter displays the version information for the Tail2Syslog.

Step 2 Press the **Enter** key to start monitoring log files and forwarding events to QRadar SIEM.

Additional Examples

Additional examples are provided to assist with creating your configuration file:

- **Example 1: Common Usage**
- **Example 2: Load Balancing Events**
- **Example 3: Monitoring Multiple Folders**

Example 1: Common Usage

The most common usage of Tail2Syslog is to monitor a single log file rolling over by date or time.

```
Destinations=1.1.1.1
Globs=/root/logs/audit*
ArchiveDir=/store/complete/
DeletionThreshold=0
```

In the first example, Tail2Syslog is configured to monitor a single directory and forwards syslog events to the IP address `1.1.1.1`. Tail2Syslog monitors the most recent file matching the pattern **`audit*`**, until a more recent file appears. The newest file is then used to monitor and forward events. An archive directory is specified, but because **`DeletionThreshold=0`**, the files are moved for archiving, but never deleted to the archive directory.

Example 2: Load Balancing Events

The following example describes how to use the configuration file to load balance events provided to your event collectors when a device writes a large number of events to audit or log files.

```
Destinations=1.1.1.1|2.2.2.2
Globs=/root/logs/audit*
ArchiveDir=/store/complete/
DeletionThreshold=5
```

In example 2, Tail2Syslog is configured with two IP addresses, one for each event collector. Tail2Syslog monitors the most recent file matching the pattern **`audit*`**, until a more recent file appears and forward syslog events to **`1.1.1.1`** and **`2.2.2.2`**. Each destination receives an equal portion of the events from the file with the most recent modified date. The newest file used to monitor and forward events is kept in the globs directory, and old files are moved to the archive directory. Only the five most recent log files (by last modified date) are kept in the archive directory, the rest are purged to preserve disk space.

Example 3: Monitoring Multiple Folders

The following example describes how to use the configuration file to load balance events provided to your event collectors when a device writes a large number of events to audit or log files.

```
Destinations=1.1.1.1
Globs=/root/mon/audit*|/root/tue/audit*|/root/wed/audit*|/root/
thu/audit*|/root/fri/audit*|/root/sat/audit*|/root/sun/audit*
ArchiveDir=/store/complete/
DeletionThreshold=7
```

In example 3, Tail2Syslog monitors the most recent file matching the directory day of the week and the file pattern **`audit*`**, then forwards syslog events to the IP address **`1.1.1.1`**. Only the most recent log files (by last modified date) are kept in each globs directory, meaning mon, tue, wed, thu, fri, sat, and sun each have one **`audit*`** file, the rest are moved to the archive directory. Only the seven most recent log files (by last modified date) are kept in the archive directory, the rest are purged to preserve disk space.

A NOTICES AND TRADEMARKS

What's in this appendix:

This section describes some important notices, trademarks, and compliance information.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785 U.S.A.

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
170 Tracer Lane,
Waltham MA 02451, USA

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the

capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at *http:\\www.ibm.com/legal/copytrade.shtml*.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.