tenable network security

Log Correlation Engine Backup Strategy

February 5, 2015

(Revision 2)

Table of Contents

Introduction

This document describes procedures for backing up Tenable Network Security's Log Correlation Engine™ (LCE) data. Installation, configuration, and management of LCE are covered by other documents. Please email any comments and suggestions to support@tenable.com.

A basic understanding of LCE functionality, SecurityCenter, and system administration is assumed.

Standards and Conventions

Throughout the documentation, filenames, daemons, and executables are indicated with a **courier bold** font such as **gunzip**, **httpd**, and **/etc/passwd**.

Command line options and keywords are also indicated with the **courier bold** font. Command line examples may or may not include the command line prompt and output text from the results of the command. Command line examples will display the command being run in **courier bold** to indicate what the user typed while the sample output generated by the system will be indicated in courier (not bold). Following is an example running of the Unix **pwd** command:

```
# pwd
/opt/sc4/daemons
#
```

Important notes and considerations are highlighted with this symbol and grey text boxes.

Tips, examples, and best practices are highlighted with this symbol and white on blue text.

Overview

Technologies and strategies to back up data can vary greatly from one organization to another. This document is intended to suggest some simple methods to back up LCE data for organizations that may not have a comprehensive backup policy or procedure.

The procedures described in this document are not intended to replace your current backup strategy, but are suggested methods to consider if you do not have an existing infrastructure to back up data.

Backing up LCE Data Using Syslog Forwarding

The "Data Forwarding" option is available as a part of the LCE GUI. While enabling this option does not back up the LCE server and configuration, the data can be forwarded to one or more dedicated syslog servers or other LCEs as the events arrive.

To configure "Syslog Forwarding", log in to the LCE GUI by navigating to the DNS name or the IP address of the LCE server over port 8836 (https://<DNS name or IP address>:8836) in your preferred web browser, and select "Configuration".

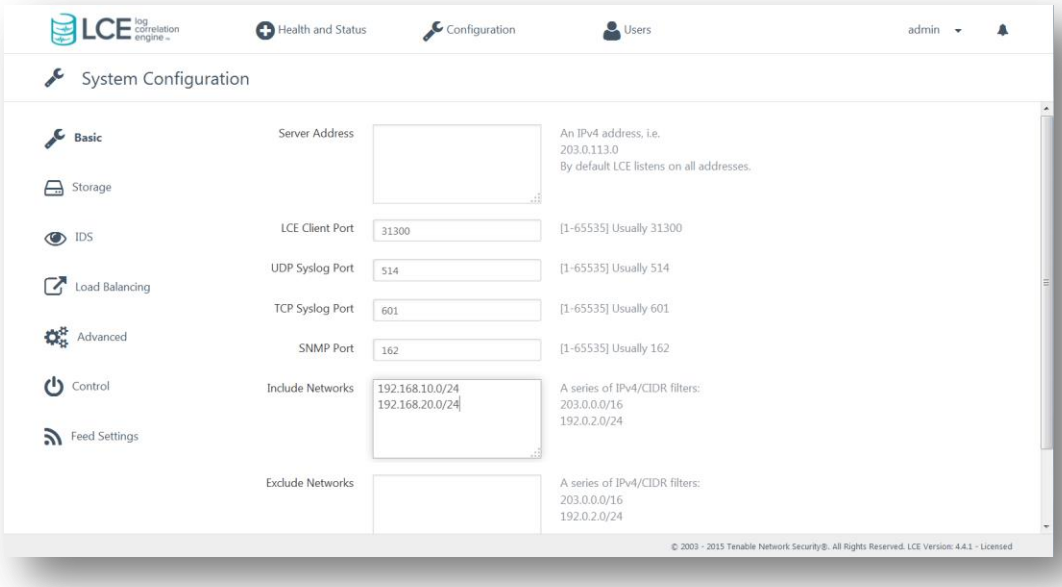

Chose the "Advanced" option and locate the "Data Forwarding" section as shown below:

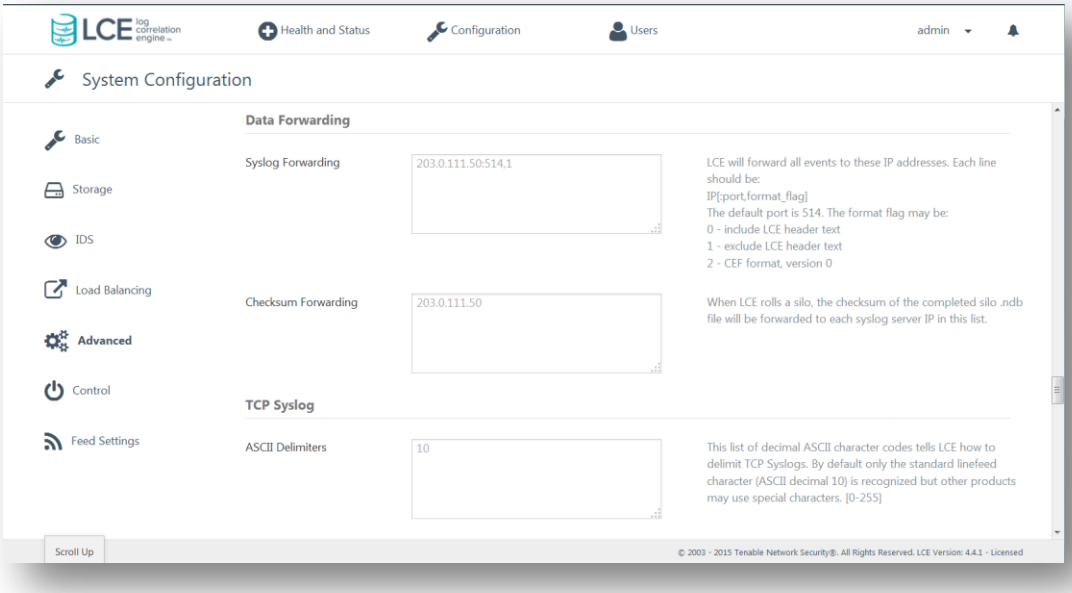

Add the desired IP address(es) of the host(s) that will receive the "Syslog Forwarding" data. The format of the entry is "IP address:port,format". If no port is specified, UDP 514 is assumed. When the "format" option is set to "1", only the original log is forwarded without any information that it was forwarded from the LCE server. If unset or set to "0", it will include the LCE header when forwarding logs. If "2 "is set, the log will be forwarded in CEF format.

If you have more than one LCE or syslog server to forward data to, specify each one on a separate line as shown in the following example:

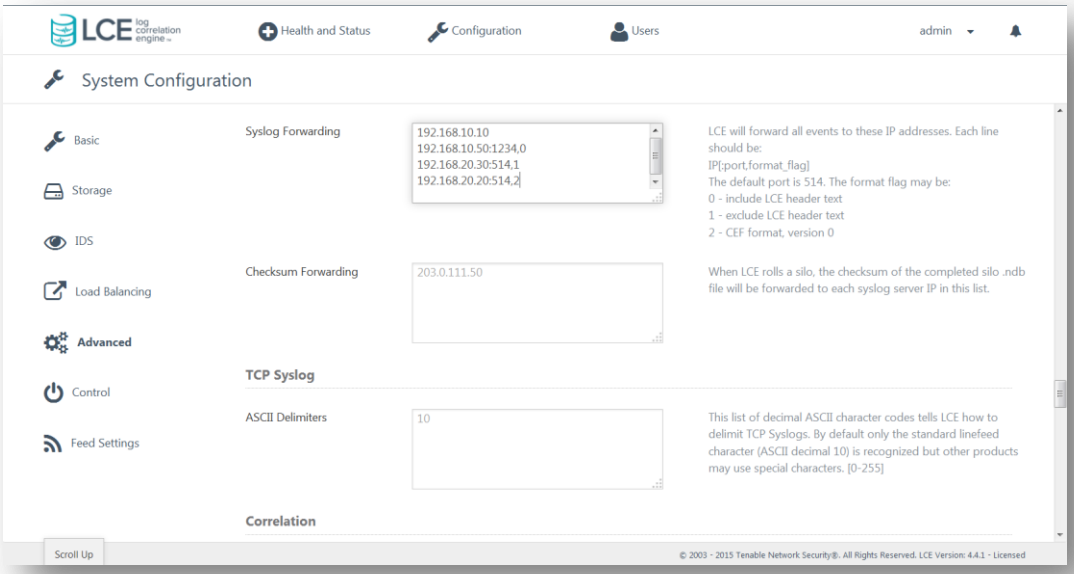

After the IP addresses have been added, scroll to the bottom of the page and select “Update” to apply the changes.

Backing up LCE Data Using rsync

The **`rsync`** utility enables you to back up only the LCE data that has changed since the last backup. There are a few issues to take into consideration when backing up an LCE server using **`rsync`**:

- Frequency and scope of backup
- Scheduling downtime

> Using **`rsync`** commands incorrectly can lead to data loss. It is very important to ensure the commands are entered correctly. If the source and target are reversed, the data from the backup location will overwrite the production data.

Frequency and Scope

It is recommended that the entire **`/opt/lce`** directory tree be backed up on a regular basis. The bulk of the space in this directory will consist of the LCE databases that contain the stored log data. Once the initial **`rsync`** is complete, the data copied will only consist of incremental changes to stored data, log files, PRMs, TASLs, and configuration data.

> Note: **`rsync`** is not intended for version management. This utility is designed to synchronize the data between two locations. The deltas that are backed up overwrite previous versions of files in the archive and do not maintain the older version. For example, if you performed a backup on a Tuesday and ran **`rsync`** every day for several days after, you could not roll back to the Tuesday version of the data.

Downtime

Downtime will need to be scheduled to gain the most complete backups possible. It is not recommended that **rsync** be performed on a live system for several reasons.

The first reason is that the current silo and associated files that are recording log data will be changing during the **rsync** process and will not match during verification attempts.

Another concern is that log files will likely change during the **rsync** process and will not match verification after the **rsync** is complete. These logs include those in **/opt/lce/admin/log** and **/opt/lce/log_archive** (if in use).

The last thing to note is that when performing the **rsync** while LCE is running, incoming data may be lost. On a busy LCE server, the incoming logs continue to be sent from their sources and the **rsync** process can send a lot of data across the same network interface(s). This may cause congestion that hinders both the collection of logs and the **rsync** process.

If the LCE server is shut down for the **rsync** process, data from LCE Clients will be cached until the server comes back online. However, data from syslog and other sources will be lost as they do not know when the LCE server is not available.

Running rsync

While specific environmental variables must be considered, testing has determined that the best generic method is the following command syntax as the root user:

```
# rsync -avzr -e ssh --delete /opt/lce/ root@<host>:/opt/lce/
```

The breakdown of the command is as follows:

Command Element	Description
-a	Sets the archive bit
-v	Displays verbose progress messages to the terminal screen
-z	Compresses the data before sending to the target machine
-r	Copies subdirectories
-e	Tells **rsync** which shell to use (in this case, "ssh" is specified to encrypt the data during transfer)
--delete	Compares the list of files available from the source host with the destination host. If the source no longer contains a file listed on the destination, the destination copy will be removed.
/opt/lce/	Represents the source directory to copy data from

root@<host>:/opt/lce/	Represents the destination directory. In this case, we are logging in as the root user. Replace <host> with the DNS name or IP address of the destination server, and specify the directory to copy to. If the directory does not exist, it will be created. Replace this with /path/to/local/directory if the destination is mounted on the same host as the LCE server.

Confirming Data Synchronization

Results of the **rsync** process can be confirmed by comparing MD5 checksum information of the source and copied data.

Collect the **md5sums** of the source data to a text file just prior to or just after **rsync** completion with a command such as:

```
# find /opt/lce/ -type f -print0| xargs -0 md5sum > lce_checksums.md5
```

This process may take quite some time depending on the number of files, their sizes, and the resources of the host server. The file will likely contain thousands of lines similar to:

```
0ef2b321b7d053f95fe4504c8fdd78ef  /opt/lce/lce.txt
```

The data at the start of the line is the **md5sum**, and the latter part is the associated file name. Once the process has completed, copy the **lce_checksums.md5** file to the target system.

On the target system, edit the **lce_checksums.md5** file to reflect the proper path to the **rsync** target if it was different than **/opt/lce/**. Next, run the following command:

```
# md5sum -c lce_checksums.md5 |grep -i failed > failed_check.txt
```

The command compares the **lce_checksums.md5** list of files and checksums against the targeted files on the local file system. Any of the files that return as failed in matching will be directed to the **failed_check.txt** file. This file may be reviewed for the files that show a different **md5sum** on the target machine than the source. If a file contains a different **md5sum**, that indicates a difference in the source and destination files. Leaving out the "`|grep -i failed`" portion of the command will save a list all of the compared files, both successful and failed. Leaving out the "`> failed_check.txt`" portion of the command will simply display the results in the terminal window without saving the output.

For More Information

Tenable has produced a variety of documents detailing the LCE's deployment, configuration, user operation, and overall testing. These documents are listed here:

- Log Correlation Engine 4.2 Architecture Guide – provides a high-level view of LCE architecture and supported platforms/environments.
- Log Correlation Engine 4.4 Administrator and User Guide – describes installation, configuration, and operation of the LCE.
- Log Correlation Engine 4.4 Quick Start Guide – provides basic instructions to quickly install and configure an LCE server. A more detailed description of configuration and management of an LCE server is provided in the "LCE Administration and User Guide" document.

- Log Correlation Engine 4.4 Client Guide – how to configure, operate, and manage the various Linux, Unix, Windows, NetFlow, and other clients.
- Log Correlation Engine 4.4 OPSEC Client Guide – how to configure, operate, and manage the OPSEC Client.
- LCE 4.4 High Availability Large Scale Deployment Guide – details various configuration methods, architecture examples, and hardware specifications for performance and high availability of large scale deployments of Tenable's Log Correlation Engine (LCE).
- LCE Best Practices – Learn how to best leverage the Log Correlation Engine in your enterprise.
- Tenable Event Correlation – outlines various methods of event correlation provided by Tenable products and describes the type of information leveraged by the correlation, and how this can be used to monitor security and compliance on enterprise networks.
- Tenable Products Plugin Families – provides a description and summary of the plugin families for Nessus, Log Correlation Engine, and the Passive Vulnerability Scanner.
- Log Correlation Engine Log Normalization Guide – explanation of the LCE's log parsing syntax with extensive examples of log parsing and manipulating the LCE's `.prm` libraries.
- Log Correlation Engine TASL Reference Guide – explanation of the Tenable Application Scripting Language with extensive examples of a variety of correlation rules.
- Log Correlation Engine 4.0 Statistics Daemon Guide – configuration, operation, and theory of the LCE's statistic daemon used to discover behavioral anomalies.
- Log Correlation Engine 3.6 Large Disk Array Install Guide – configuration, operation, and theory for using the LCE in large disk array environments.
- Example Custom LCE Log Parsing - Minecraft Server Logs – describes how to create a custom log parser using Minecraft as an example.

Documentation is also available for Nessus, the Passive Vulnerability Scanner, and SecurityCenter through the Tenable Support Portal located at https://support.tenable.com/.

There are also some relevant postings at Tenable's blog located at http://www.tenable.com/blog and at the Tenable Discussion Forums located at https://discussions.nessus.org/community/lce.

For further information, please contact Tenable at support@tenable.com, sales@tenable.com, or visit our web site at http://www.tenable.com/.

About Tenable Network Security

Tenable Network Security provides continuous network monitoring to identify vulnerabilities, reduce risk and ensure compliance. Our family of products includes SecurityCenter Continuous View™, which provides the most comprehensive and integrated view of network health, and Nessus®, the global standard in detecting and assessing network data. Tenable is relied upon by more than 20,000 organizations, including the entire U.S. Department of Defense and many of the world's largest companies and governments. For more information, visit tenable.com.