EXTENDED SERVICES

# Splunk

Log Management and Correlation

Version 04
12/12/2013

# TABLE OF CONTENTS

# Splunk
# Log Management and Correlation

## 1 SCOPE

CERT-EU offers a starter kit licenses for Splunk with access to relevant and valuable detection rules (IOCs) to facilitate log correlation and detection of anomalous events in the network of the constituent.

It uses the same IOCs as the extended services IDS system for searching past events in the log files using a multi-tier index or detecting events in real-time using rules applied to log files from multiple sources.

For this to work is necessary that CERT-EU Splunk console has access to constituent local indexer server.

Splunk licence cost is based on GB of log data been indexed per day. CERT-EU can advise on possible ways to reduce the amount of data before uploading them into Splunk ensuring maximum optimisation of data volume usage.

Splunk can be deployed by using a single software component and configuration. It can coexist with existing infrastructure or be deployed as a universal platform for accessing IT data.

The simplest deployment is indexing and searching on the same server; but Splunk components should be deployed on different servers to address load and availability requirements.

## 2 SPLUNK COMPONENTS

### *2.1 FORWARDER*

Forwarder is a Splunk component that forward originated data to remote indexers servers for indexing and storage. In most cases, they do not index data themselves. The forwarders can easily co-exist on the machines generating the data, because the data-consuming function has minimal impact on machine performance. Forwarders consume data locally and then forward the data across the network to another Splunk component, called the indexer.

## 2.2 INDEXER

The indexer indexes the data from forwarders and runs searches. It provides indexing capability for local and remote data and host the primary Splunk datastore, as well as Splunk Web. It should reside on a machine by itself.

The following diagram shows several forwarders sending data to a single indexer:

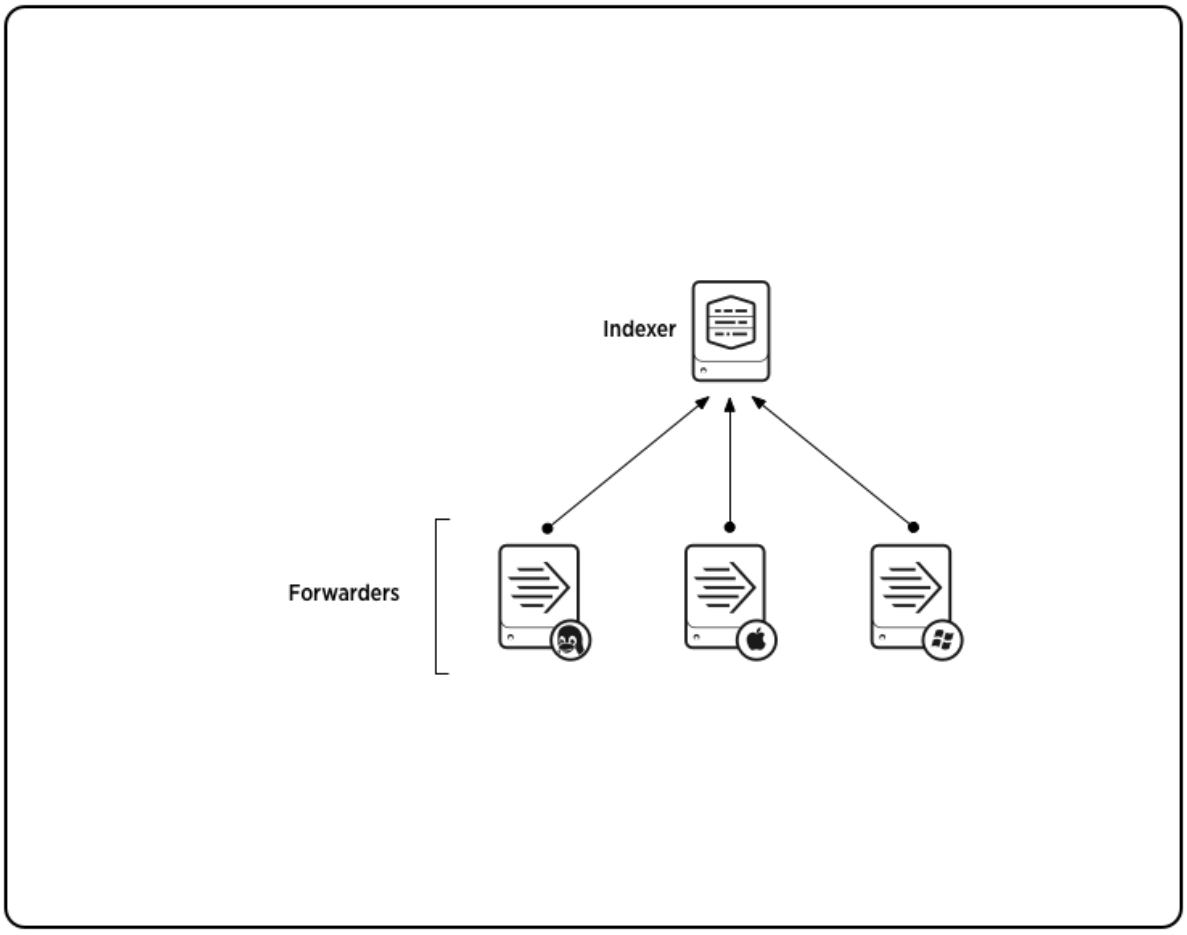

For a larger deployment, we might have hundreds of forwarders sending data to a number of indexers. We can configure load balancing on the forwarders, so that they distribute their data across some or all of the indexers. Not only does load balancing help with scaling, but it also provides a fail-over capability if one of the indexers goes down. The forwarders automatically switch to sending their data to any indexers that remain alive.

In following diagram, each forwarder load-balances its data across two indexers:

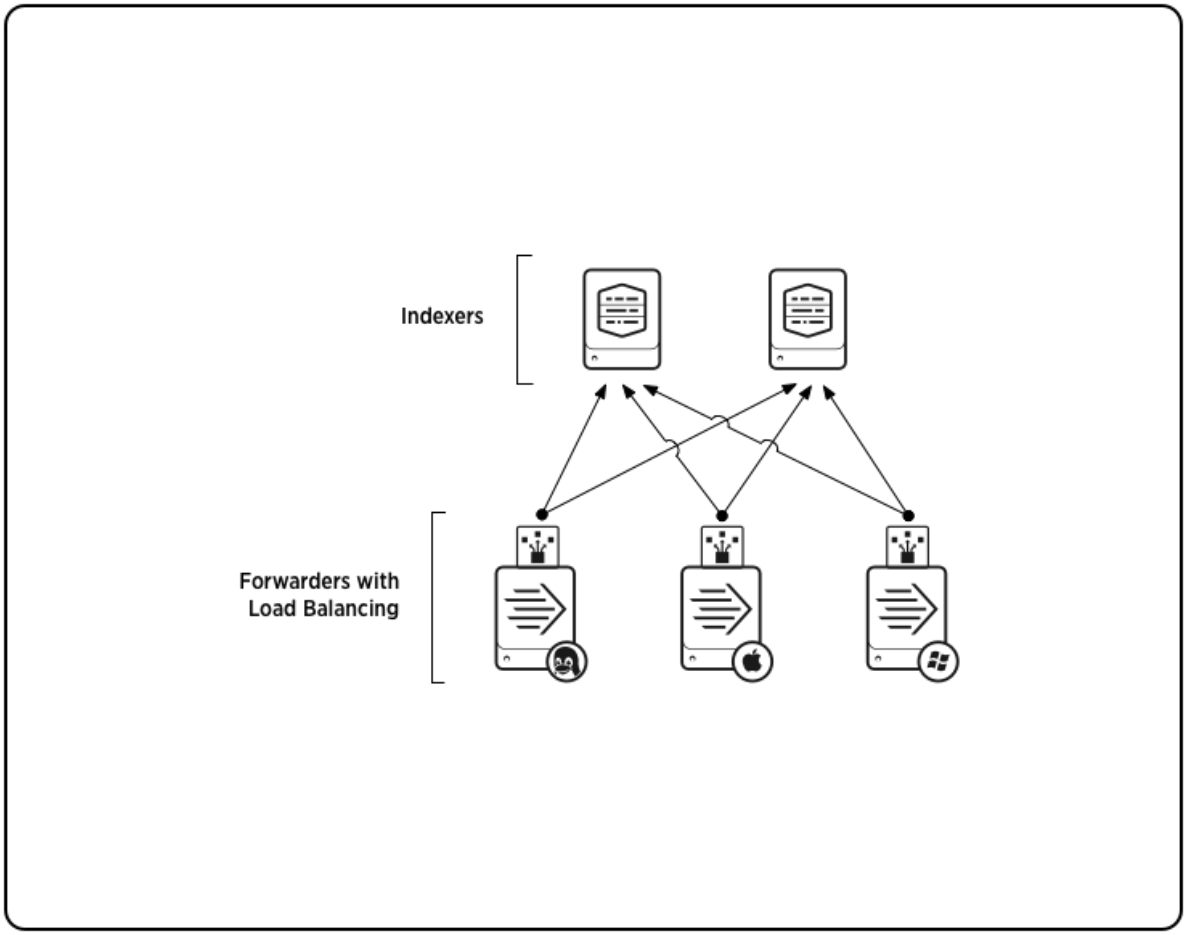

## *2.3 HEAD SEARCH*

To coordinate and consolidate search activities across multiple indexers, we separate out the functions of indexing and searching. This type of deployment is called distributed search and each indexer just indexes data and performs searches across its own indexes. A Splunk instance dedicated to search management, the search head, coordinates searches across the set of indexers, consolidating the results and presenting them to the user:

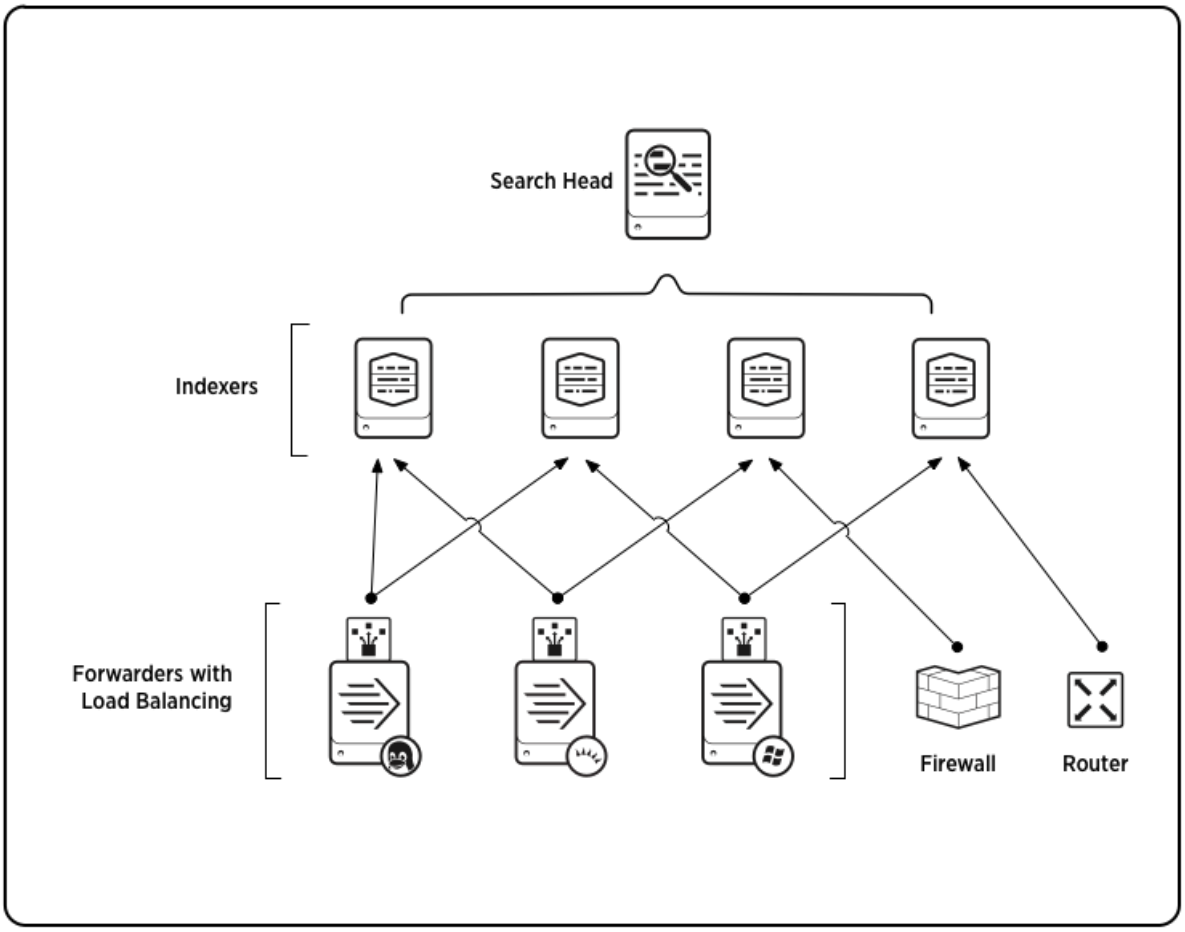


## 2.4 DEPLOYMENT SERVER

Both indexers and forwarders can also act as deployment servers. A deployment server distributes configuration information to running instances of Splunk via a push mechanism which is enabled through configuration.

# 3 DEPLOYMENT DETAILS

Constituent Splunk local server deployments should be carefully planned and configured to achieve the best performance. Search and index performance depends on the total volume of data being indexed and the number of active concurrent searches (scheduled or other) at any time.

Indexers, in addition to rapidly writing data to disk, do much of the work involved in running searches: reading data off disk, decompressing it, extracting knowledge and reporting. As a result, when scaling up data volumes, additional indexers should be added. These indexers will help handle larger volumes of data, reduce contention for resources during searches and accelerate search performance.

Your physical hardware should reflect these minimum requirements:

- Intel x86-64-bit chip architecture
- 2 CPU, 4 core per CPU (8 cores total), ~3Ghz per core
- 16 GB of RAM
- RAID 0 or 1+0, with a 64 bit OS installed
- 800 IOPS disk performance
- Standard 1Gb Ethernet NICs

Splunk is often constrained by disk I/O first, so always consider that first when selecting hardware. Note: RAID 0 configurations do not provide fault-tolerance. Be certain that a RAID 0 configuration meets your data reliability needs before deploying a Splunk indexer on a system configured with RAID 0.

## *3.1 CERT-EU SOLUTION ARCHITECTURE*

CERT-EU helps constituent to deploy Splunk solution and import data for various sources like Active Directory security log, proxy logs etc.

All data files are in constituent control. If constituent choose to use CERT-EU Splunk licence there should be a communication link between CERT-EU Splunk console and local indexers for licence communication. If constituent has its own Splunk license he can make lookups based on the list of IOCs provided by CERT-EU.

The following diagram shows a basic CERT-EU / Constituent solution architecture with several forwarders sending data to a single indexer and automatic search for malicious activities: