



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials (Security 401)"
at <http://www.giac.org/registration/gsec>

Discovering Security Events of Interest Using Splunk

GIAC GSEC Gold Certification

Author: Carrie Roberts, clr2of8@gmail.com

Advisor: Antonios Atlasis

Accepted: July 10, 2013

Abstract

Security events of interest can be discovered by analyzing several different sources of machine data, including logs. Applications and the servers they run on contain many valuable logs which detail the events that have occurred on them. By analyzing and correlating this data, important information about the attacks against these systems can be discovered. Splunk is a powerful tool for analyzing such data. It provides a high performance solution for analyzing large amounts of unstructured data from multiple sources. This paper includes procedures for setting up a Splunk server and forwarding data to it from multiple sources. Example searches and use of pre-built add on functionality is given. It is a concise, comprehensive guide for deploying and using a centralized system for intelligence gathering with a focus on detecting security events of interest.

1. Introduction

Servers and the applications that run on them are under attack by malicious users through a variety of techniques (Mitnik & Simon, 2006). Detecting intrusion or destructive attempts against these systems plays an important role in being able to protect them. The large size and complexity of these systems can make the task of monitoring them for security events of interest a daunting one. The actions on these systems can be, and typically are, recorded in various logs throughout the system (Skoudis & Liston, 2006). Aggregating all these logs into a central location and efficiently searching through them becomes difficult as the systems scale. Splunk is a software tool that can efficiently provide a solution for aggregating, indexing and visualizing large scale data.

Splunk can aggregate logs from many sources into one centralized location. It is used to analyze structured and unstructured machine data (Splunk Enterprise Product Data Sheet, 2012). Machine data in this case refers to computerized human readable data. Splunk indexes the data for high performance searches. It recognizes key value pairs automatically as well as many elements of common log files like ‘referrer’, ‘status’ and ‘user agent’. It utilizes an intuitive and familiar search language with constructs similar to those used to query databases. Splunk is scalable and designed for processing of “big data.”

This paper includes a Splunk architecture overview and instructions for setting up this architecture in an environment with multiple data sources and types. The examples focus on detecting security events of interest from various data sources. Example attack scenarios are presented with queries that can be used within Splunk to detect attack and intrusion attempts. A further discussion of Splunk Apps, which include pre-packaged functionality, is included. Splunk helps to provide a clear understanding of what is really going on in a system and makes it easy to spot anomalous behavior.

A wide variety of visualizations for the data can be created including graphs, pie charts, tables and gauges. These visualizations can be customized and placed on dashboards for easy to access summarized results (Splunk User Manual, 2013). Others can access these views via a URL or have a PDF version of the report emailed to them on a regular schedule.

Alerts can be created to notify system operators of any noteworthy events discovered through running searches. Searches can be saved and scheduled to run automatically. Setting up saved searches with alerts is an effective way to monitor the system as it operates.

Business metrics, customer behavior study, debugging and security investigations are some of the key uses for the Splunk application. The focus of this paper is to describe how to setup and use Splunk to discover a variety of security events of interest. An architecture overview is given followed by setup procedures for the example architecture and instructions for searching, reporting and alerting on the data of interest.

2. Splunk Architecture

The Splunk architecture consists of a Splunk Server and optional data forwarders to get data to the server. The Splunk server indexes and searches data and can be deployed as a distributed service to meet performance needs. The server includes an interactive web interface. The interface can be used to enter search commands, create and view dashboards, make configuration changes and more. Windows, Linux and Unix installers are available for both the server and forwarder software.

The forwarders are used to send data to the Splunk server from another machine. There is no dependency between the operating system that the server runs on and the operating system that the forwarder runs on. There are two types of forwarders, the universal forwarder and the heavy forwarder. The universal forwarder is a light weight version of Splunk with only the pieces needed to send data to the server. The heavy forwarder is a full Splunk implementation with some of the features disabled to make it smaller.

An architecture that includes a Splunk server, several data input sources and forwarders, is typical. The setup and use of the example architecture shown in Figure 1 will be covered in detail in the remainder of this paper. This includes a Splunk Server with its web interface running on a Linux server. Data aggregation is accomplished through the use of forwarders as well as custom scripts. One forwarder is configured to forward web application server logs from a Linux machine. The second forwarder is installed on a Windows 2008 server running a domain controller. Lastly, data is collected from an Amazon S3 storage location using a custom

script running on the Splunk server itself. The custom data is representing any other data source or structure that may exist in a system, possibly of a unique nature.




Figure 1. Example Splunk Architecture

When the data set to be indexed and searched becomes very large, aka “big data”, a distributed Splunk architecture is deployed. In this architecture, the core components of indexing and searching are separated. Multiple Splunk indexers are deployed and the forwarders are configured to automatically load balance between any and all of them (Splunk Distributed Deployment Manual, 2013). Each indexer only indexes and searches locally and a separately deployed search head is used to coordinate the searches between them and return the final result. Splunk supports this architecture to meet performance and scalability demands.

3. Setting up the Splunk Server

The Splunk server shown in Figure 1 includes the indexer, search mechanism and the web interface to it. The web interface is a convenient place for specifying searches and for viewing reports. It also allows configuration of the server itself. The installation of a Splunk Server is well documented on the splunk.com website and is covered only briefly here. The installation here is specific for the Linux install in the example architecture shown in Figure 1:

- 1) Download the Splunk server from <http://www.splunk.com/download> (splunk-5.0.2-149561-Linux-x86_64.tgz in this case)
- 2) Decompress the file from the command line:
 - a. `$tar -xzf splunk-5.0.2-149561-Linux-x86_64.tgz`
- 3) Start the Splunk server from the command line:
 - a. `$< splunk-dir>/splunk/bin/splunk start`
- 4) Configure Splunk to start at boot time:
 - a. `$sudo <splunk-dir>/bin/splunk enable boot-start`
- 5) View the web interface by visiting <http://localhost:8080> in a web browser, from the Splunk server machine, or replace localhost with the server IP address to access it from another computer.

After following instructions to change the password, the welcome page is shown. There are a few additional steps in order to setup the Splunk server to receive data from a forwarder.

- 1) While logged in as admin, click the “Manager” link in the upper right hand corner. Choose “Forwarding and receiving” in the Data section.
- 2) In the Receive data section next to “Configure receiving” select “Add new”.
- 3) Select an available port on the server to receive data on, for example 7878.
- 4) Restart the Splunk server from the command line:
 - a. `$<path-to-splunk-dir>/splunk/bin/splunk restart`

Alternatively, steps 1-4 above could be done on the command line as follows:

- 1) Listen on for forwarded data on a port:
 - a. `$<path-to-splunk-dir>/splunk/bin/splunk enable listen <port> -auth <username>:<password>`
- 2) Restart the Splunk server from the command line:
 - a. `$<path-to-splunk-dir>/splunk/bin/splunk restart`

The Splunk server is now functional and ready to index data, execute searches and receive data from forwarders.

4. Getting Data to the Splunk Server

The example architecture includes universal forwarders on two different types of machines. The forwarders send data to the Splunk server. The following steps will install and configure the forwarder on the Linux machine to forward web application logs and other server data.

- 1) Install the Universal forwarder onto the Linux web application server. The installer can be downloaded from <http://www.splunk.com/download/universalforwarder>. The “uname -a” command can be used on the server to determine the kernel version and architecture.
- 2) Click on the installer (splunkforwarder-5.0.3-163460-Linux-x86_64.tgz in this example) and login to Splunk.com and save the installer.
- 3) Decompress the file from the command line:
 - a. `$tar -xvf splunkforwarder-5.0.3-163460-Linux-x86_64.tgz`
- 4) Start the Splunk forwarder from the command line:
 - a. `$<splunkforwarder-dir>/bin/splunk start`
- 5) Read and accept license agreement.
- 6) Configure Splunk to start at boot time:
 - a. `$sudo <splunkforwarder-dir>/bin/splunk enable boot-start`
- 7) Point forwarder to the Splunk Server (receiver):
 - a. `$sudo <splunkforwarder-dir>/bin/splunk add forward-server hostname.domain:port`
 - b. Set the *hostname.domain* to the hostname and domain of the Splunk Server. The IP address of the server can be used if a hostname and domain have not been setup. Set port equal to the port number selected during the Splunk server setup (7878 in this case)
 - c. Enter Splunk username and password. The default username is “admin” and the default password is “changeme”.
- 8) Configure data to be forwarded (web server logs):
 - a. `sudo <splunkforwarder-dir>/bin/splunk add monitor /var/log/apache2/ -index main -sourcetype apache-logs`

- b. This command will forward everything in the /var/log/apache2 directory to the main index of the Splunk server with a sourcetype of “apache-logs”
- 9) Configure data to be forwarded (application logs)
- a. `sudo <splunkforwarder-dir>/bin/splunk add monitor /var/log/my-app/ -index main -sourcetype application-logs`
 - b. This command will forward everything in the /var/log/my-app directory to the main index of the Splunk server with a sourcetype of “application-logs”

To configure the forwarder to send logs and other data of interest from the Windows Server 2008 machine running the domain controller, the following steps are necessary.

- 1) Install the Universal forwarder onto the windows domain controller. The msi installer can be downloaded from <http://www.splunk.com/download/universalforwarder>. Double click on the installer, click next and then accept the license agreement and click next.
- 2) This setup will include a receiving indexer so Leave the Deployment Server information blank and click next.
- 3) Enter the Splunk server IP address or hostname for the receiving indexer (the Splunk server) and the port number the receiver is listening on, port 7878 in this example.
- 4) Optionally configure the certificate information, and then click Next. Forwarded data will be encrypted with the default Splunk certificate if none other is specified.
- 5) Select ‘Local Data Only’ which only allows collection from this host.
- 6) Select the Windows event logs to forward. These include application, security and systems logs as well as CPU load, memory, disk and network usage. Select all of them.
- 7) Click Next, then Install, then click Finish.

Lastly, for the example architecture of Figure 1, custom data is collected from an Amazon S3 storage location. This data is received on the Splunk server through a script which synchronizes files between the s3 storage location and the server using the s3cmd tool. The script is run every minute utilizing a cron job. The following steps result in the s3 data being regularly indexed on the Splunk server. These steps should be performed on the Splunk server.

- 1) Install the s3cmd tool used to pull files to the Splunk Server:

Carrie Roberts, clr2of8@gmail.com

- a. `$sudo apt-get install s3cmd`
- 2) Configure the s3cmd tool with the s3 access keys. Start the configuration process with the following command.
 - a. `$s3cmd --configure`
- 3) Create the scheduled job to pull files to Splunk server every minute:
 - a. `$crontab -e`
 - b. Enter and save the text shown here between the quotes to the crontab file `“* * * * * s3cmd sync --skip-existing s3://s3-bucket-name/custom-data/ /data/s3data/”`
 - c. Replace ‘s3-bucket-name’ with the name of the bucket that contains the data. Include the rest of the path to the data and be sure to include the trailing forward slash. The last parameter, `/data/s3data/`, is the directory on the Splunk server to put the data pulled from s3.
 - d. Note that creating a more advanced script which protects against running more than one of these commands at once is advised but not covered here.
- 4) Add the data source to Splunk
 - a. From the Splunk server web interface select “Manager” from the upper right hand corner, then “Data inputs” and click the “Add new” link on the “Files & directories” row.
 - b. Select “Browse Server” and select a data file from the directory to be monitored and click “Continue”
 - c. Start a new source type or select an existing one. The source type is used when restricting search results to a certain data set. It is also used to automatically recognize key/value pairs in the source data.
 - d. Make adjustments as needed to ensure that Splunk can locate the timestamp in the data, then click “Continue”
 - e. Accept the remaining configuration defaults.

The Splunk server is now receiving data in multiple formats, from multiple sources and automatically indexing it for searching and reporting on. This data includes applications logs, security logs, system logs, CPU load, memory and disk usage, network statistic and more. This data will include important information about the use of systems and should be analyzed and

learned from. This diverse dataset is now collected in a centralized location enhancing the ability to correlate data across systems.

5. Data Analysis Aids – Splunk Apps

Splunk provides a powerful search interface for understanding and visualizing data. The search interface will be discussed and demonstrated in a later section. First, it is useful to demonstrate some of the many pre-built searches and reports already available without needing to be trained on using the search language. These pre-built packages, or bundles of functionality, are called Splunk apps. Browsing and installing apps is done through the Splunk server web interface. There are hundreds of free and commercial apps available for both Linux and Windows (Do More with Splunk, 2013). A small sample of some of the free security related apps is listed in Figure 2.

Sample of Free Splunk Apps Available for Download	
Cisco Security Suite	Splunk for Zenprise Mobile Security Intelligence
Centrify Insight	Bee Ware I-suite Application Security Dashboard
ModSecurity	Windows Security Operations Center
TA-iptables	Splunk for Active Directory
Phishing Lookup	Splunk for QualysGuard
Security Onion	Vormetric Data Security for Splunk
Splunk for DNS	Splunk for F5 Security

Figure 2. Sample of Free Splunk Apps Available for Download

5.1. Installing Splunk Apps

The procedure for installing Splunk apps varies depending on whether there is an install button or a download button available on the app page. Both install methods are included below.

Install Splunk App directly:

- 1) From the Splunk server web interface, select “Find more apps . . .” from the App menu.
- 2) Search through available apps and press the install button for the desired app.

- 3) Enter Splunk.com website credentials and then restart Splunk

Download and install from file:

- 1) From the Splunk server web interface, select “Find more apps . . .” from the App menu.
- 2) Search through available apps and press the download button for the desired app and save the .tgz file after entering user credentials for Splunk.com.
- 3) Click “App” in the upper right hand corner of the web interface and then select “Manage Apps . . .”
- 4) Select “Install app from file”
- 5) Browse to the .tgz file that was downloaded and select “Upload”
- 6) Restart the Splunk server after installation.

There are many free and commercial apps available for download. One useful free app is the Splunk for Unix and Linux app.

5.2. Splunk for Unix and Linux App (Free)

The “Splunk for Unix and Linux” app is a comprehensive application that monitors many system metrics (Splunk for Unix and Linux, 2013). It includes many scripts that can be run on a configurable schedule to gather system data. It also provides dashboard views to visualize the data quickly and efficiently. It can monitor CPU usage, memory and disk usage, network usage, user actions and configuration files. The views provided can display per host, user, process, state, etc. as shown in Figures 3, 4 and 5.




Figure 3. Splunk view of CPU consumption by user




Figure 4. Splunk view of CPU consumption by host




Figure 5. Splunk view of CPU consumption by command

5.3. Windows Security and Operations Center App (Free)

The Windows Security and Operations Center by INFIGO IS is a free app for monitoring security information from Windows servers (Windows Security Operations Center, 2013). Once installed, the app can be accessed from the Splunk server web interface by clicking “App” in the upper right hand corner, then clicking the app. This app will report on the Windows server data that was forwarded to the Splunk server following the procedure described earlier.

Figure 6 shows the menu bar of the app with the Login events menu expanded. The app has a dashboard for successful and failed logins for active directory, NTLM and RDP. It includes a count of the number of login hosts per user. Note that logging of failed logins is not

on by default on Windows servers so this must be enabled in order to maximize the value of the dashboard information (Skoudis & Liston, 2006).



Figure 6. Windows Security and Operations Center App Menu

The User management menu includes links to graphs of added, deleted, locked, unlocked and disabled domain accounts as well as the summary information shown in Figure 7. The drop down menu is expanded to demonstrate how to select different time periods to view.




Figure 7. Summary of User Account Management

The Change Control menu includes links to graphs of domain policy changes, deleted log entries, system restarts, new added services and patch information. The Windows firewall menu includes links to dashboards of configuration changes, as well as allowed and blocked connections and binds. The Saved Searches menu includes links to all the searches used in

created all the dashboards previously mentioned. The Search menu item takes the user to the generic search page similar to what is seen when using the Search App that is included with Splunk by default. The Search App will be discussed in more detail in the Custom Searches section.

5.4. Splunk App for Active Directory (Free)

The Splunk App for Active Directory is provided by Splunk (Splunk App for Active Directory, 2013). It is similar in functionality to the Windows Security and Operations Center App (WSOC). It supports Windows Server 2003 up to Windows Server 2008 R2. One additional feature of this app is that it includes over 50 audit reports. The installation and configuration of this app is significantly more involved than the WSOC app as described in the README file for the app. Steps include enabling PowerShell, installing Perl and deploying technology add-ons to each domain controller.

5.5. Splunk App for Enterprise Security (Commercial)

Splunk has created a commercial app called “Splunk App for Enterprise Security” (Splunk App for Enterprise Security, 2013). This app was selected best Security Information and Event Management (SIEM) appliance in 2013 by SC Magazine (Splunk Enterprise Selected, 2013). The Enterprise Security app is full featured and its correlation searches include anomalous new services/processes/accounts and malware and brute force access detection (Splunk User Manual, 2013). Details of this app are not included in this document because it is not freely available for trial and analysis.

6. Detecting Security Events of Interest – Custom Searches

Detecting security events of interest is the main premise of an Intrusion Detection System (Jacobson, 2009). Splunk can be used as an intrusion detection system but it is recommended here only as an enhancement to an existing one. Splunk server comes pre-installed with the Search App. This is where custom searches can be done across all indexed data sources. It can be accessed through the Splunk server web interface by clicking on “App” in the upper right

hand corner, then selecting “Search”. Figure 8 shows the Search App. Custom searches can be entered in the text box which contains the word “search” in gray.





Figure 8. The Splunk Search App

The Splunk search language is intuitive and provides functionality similar to that found in a standard query language like SQL. Several search examples are given herein to introduce the search mechanism.

To limit search results to data in a particular index, the “main” index in this case, enter ‘index=main’ in the search box.



The search is executed by either pressing the enter key after entering the text or by pressing the image of the magnifying glass on the right end of the search bar. The search results appear in the area below the search bar and are displayed as they are being found, newest first. This is useful so that the most recent results can be inspected immediately without having to wait for the entire search to complete. A timeline showing the count of results is included by default and is shown circled in red in Figure 9. Hovering over the timeline displays count and time information for the given bar on the graph. Clicking on a bar in the timeline will limit the search results to that time period. The area circled in blue in Figure 9 shows the individual search results as found in the data files.




Figure 9. Splunk Search Results Window

The data returned from a search can be restricted to a specific time range by selecting the button just to the left of the magnifying glass which, by default, says “All time”. “All time” returns data with any time stamp. The menu provides quick access to commonly used time ranges but also allows custom time ranges to be specified. Figure 10 shows the time selection menu.




Figure 10. Splunk Search App Time Range Selection Menu

Time ranges can also be set directly in the search itself as shown in this example:

```
>index=main earliest=-1w latest=-1d
```

The search above is limited to a time starting 1 week ago (-1w) and ending 1 day ago (-1d).

The space character is interpreted as an “AND” operator. Searches are case insensitive by default. In the example below, the search will return results with the word “Network” and the word “Interface” regardless of case

```
>Network Interface
```

The following would return results for entries that contain the word “network” or “interface”

```
>Network OR Interface
```

Note that the “OR” command must be all upper case in order to be interpreted as a command instead of a search term. This is typical of search commands. The search command can create charts and tables of stats like averages and counts. The following search would find one type of login failure from the Windows event logs and output a table of number of hosts per user.

```
>sourcetype="WinEventLog:Security" (EventCode=4771 AND Keywords="Audit Failure") | stats distinct_count(host) by Account_Name
```

This search introduces the pipe character “|” which takes the output of one search and uses it as input to the next search. In this case, a list of login failures is output to the stats command. The stats command is using the distinct_count() method to determine the number of distinct hosts that each user failed to login from.

As Splunk indexes data, it automatically recognizes key/value pairs in the input file. It also recognizes many other standard input file formats like Apache logs. This makes searching for the data easier without needing to create custom expressions to extract the data. A raw log line from an apache access log is shown below.

```
192.168.0.178 - - [14/May/2013:16:42:47 -0700] "GET /testphp.php HTTP/1.1" 200 8877 "-"
"Mozilla/5.0 (X11; Ubuntu; Linux x86_64; rv:15.0) Gecko/20100101 Firefox/15.0.1"
```

It could be difficult to develop a search that can extract the HTTP method, GET or POST, for example, but Splunk provides help because it recognizes Apache log format. In this case, the following search can be used to find all GET requests.

```
>method=GET
```

Additional recognized fields include, but are not limited to, client IP, status code, user agent and referrer. If Splunk does not automatically recognize the key/value pairs as it should, it can be configured to do so.

Specific events of interest can be analyzed through these searches. The following section describes searches that can be used to detect security events of interest.

6.1. Searching for Cross-site Scripting Attempts

A common web service attack is cross site scripting (XSS). An XSS attack results when data controlled by the attacker is interpreted by the web browser as if it was code written by the web developer (Weinberger et al., 2011). A common example of this is a web page that reflects, or displays, input from one user in the browser of another user. A blogging website would be one such example. A commonly used form of XSS payload is enclosed in `<script>` tags as shown in the javascript shown here:

```
<script>alert("XSS")</script>
```

An attempt to submit such a payload to a web application can be detected with a Splunk query assuming that the payload is logged somewhere. In this example, the application logs all interactions in the `application.log` file. The following query would find the basic use of script tags submitted by users.

```
>source="/var/log/my-app/application.log" "<script>" OR "</script>"
```

The `"source="` tag tells Splunk in which files to perform the search. There are two search terms `"<script>"` and `"</script>"` which are searched for in a case insensitive manner. This means that the search would match `"<script>"`, `"<sCrIpT>"` and `"<SCRIPT>"` for example. The

capital “OR” has special meaning in Splunk and performs a logical OR, meaning that a match will be found if *either* search term is present.

While this is a good place to start, there are actually many varieties of the XSS attack. Each of these different attack techniques will require added search parameters. The Open Web Application Security Project (OWASP) XSS Filter Evasion Cheat Sheet (XSS Filter Evasion Cheat Sheet, 2013) lists many techniques that attacker can use to be more conspicuous. These techniques include the use of image source tags, with and without quotes, and varying case as shown in these examples:

```
<IMG SRC="javascript:alert('XSS');">
```

```
<IMG SRC=javascript:alert('XSS')>
```

```
<IMG SRC=JaVaScRiPt:alert('XSS')>
```

Other evasion techniques including the use of grave accents, malformed tags, encoding in various character sets and more are explained on the XSS Filter evasion cheat sheet. A more comprehensive Splunk query for detecting XSS attempts is included below.

```
>source="/var/log/my-app/application.log" "&#" OR "script" OR "`" OR "cookie" OR "alert" OR "%00"
```

This query is looking for any occurrences of “&#” which can be indicative of an attempt to use alternative character encoding to bypass XSS filters. It also searches for the basic “script” tag which will match both the script tag and the javascript tag. The use of the grave accent is included in the search as well. A common exploitation of a XSS vulnerability is to access the browser cookies via a call to “document.cookie”. Therefore, a search for the word “cookie” may highlight an XSS attempt. Attackers often use an “alert” box when searching for vulnerabilities so the term “alert” has been included in the search. Lastly “%00” will match a null character in a URL, another indicator of malicious activity.

Sometimes special characters such as the tab character, “\t”, are inserted to subvert detection. The following searches can be used to find embedded tab, new line and carriage return control characters respectively.

```
>source="/var/log/my-app/application.log" | regex "\t"
```

```
>source="/var/log/my-app/application.log" | regex "\n"
```

```
>source="/var/log/my-app/application.log" | regex "\r"
```

The pipe character, “|”, takes the output of one search, in this case everything in the application log, and uses it as input to the next search. This search uses a regular expression, as indicated by the “regex” term, in order to specify the special control characters.

If the application logs all user input, these queries will give good insight into suspicious activity against a web application.

6.2. Searching for SQL Injection Attempts

Another attack which is on the OWASP top ten list of most critical web application security risks is injection (Category: OWASP Top Ten Project, 2013). SQL injection is one example, and will be the focus of the search queries given here as examples. An example attack string is shown below from the OWASP literature.

```
http://example.com/app/accountView?id=' or '1'=1
```

The critical piece of the attack is the use of the tick marks and the “or” and “=”. In another case, the attacker may chain together commands to cause damage to the database, for example:

```
user'; drop table users; --
```

Important aspects of this attack are the use of the tick mark again, as well as the semi-colon and the drop table command. These types of attacks are difficult to distinguish from normal traffic because both the tick mark and the equal sign are used commonly. Adjustment to the search may be needed in order to ignore known good uses of these characters. A proposed search for detecting SQL injection attempts is given here.

```
>source="/var/log/my-app/application.log" (' AND =) OR (' AND ;) OR (drop table) OR --
```

The query above searches for the following:

- 1) The use of the tick mark and the equals sign, or . . .
- 2) The use of the tick mark and a semi-colon, or . . .

- 3) The use of the drop table command, or . .
- 4) The use of --, which is an SQL comment indicator used to ignore the rest of the input

Although these queries only investigate some of the most common SQL injection attempts, they will help the system operator to understand the nature and frequency of injection attempts against the application.

6.3. Searching for Open Ports

Monitoring the server's use of CPU, memory and the network are all of interest from a security perspective. One example of an interesting statistic to monitor is the listening ports on the system. The netstat command can be used as follows to output the listening tcp and udp ports on a system.

```
netstat -ln | egrep "^tcp|^udp"
```

This can be included in a script that writes the output to a log file that can be indexed by Splunk. The script can be scheduled to run on a regular basis using a cron job. Splunk also includes its own scheduler for running scripts. This can be configured in the "Add data" section of the web interface. A graphical visualization of the ports can be created so that it is easy to spot an anomaly, or something out of the norm. Alternatively an alert can be created for immediate notification of unexpected results.

7. Creating Alerts in Splunk

Splunk searches and reports help system administrators to learn about normal system behavior as well as alert them of error conditions. From within the Search App an alert can be created for the current search by selecting "Alert ..." from the "Create" drop down menu as shown in Figure 11.

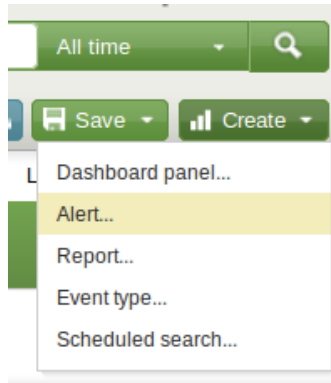


Figure 11. Splunk Menu for Creating an Alert

After selecting create alert, enter a name for it and select the desired scheduling option.

The scheduling options include:

- 1) Trigger in real-time whenever a result matches
- 2) Run on a schedule once every...
- 3) Monitor in real-time over a rolling window of...

Figure 12 shows example settings for a rolling window alert that will alert when a single user has more than 3 failed logins in any 10 minute window of time.

 A screenshot of the 'Create Alert' configuration window in Splunk. The window has a title bar 'Create Alert' with a close button. Below the title bar are three tabs: '1 Schedule' (active), '2 Actions', and '3 Sharing'. The 'Schedule' tab is selected. The form contains the following fields:

- Name:** A text input field containing 'Too many failed logins for an individual user.'
- Schedule:** A dropdown menu set to 'Monitor in real-time over a rolling window of...'. Below it is a numeric input field with '10' and a unit dropdown menu set to 'minute'.
- Trigger if:** A dropdown menu set to 'Number of results'. Below it is a dropdown menu set to 'Is greater than' and a numeric input field with '0'.

 At the bottom of the form are two buttons: 'Cancel' and 'Next >>'.

Figure 12. Configuration of a Splunk Alert

An example query for the failed login attempts is listed below. This query was built from a similar query used within the Windows Security and Operations Center app for finding failed login counts per user.

```
>sourcetype="WinEventLog:Security" ("EventCode=675" OR ("EventCode=672" AND  
Type="Failure Audit")) OR (EventCode=4771 AND "Audit Failure") NOT (User_Name="*$"  
OR Account_Name="*$") NOT Failure_Code=0x19 | stats count by Account_Name | where  
count > 3
```

The options for alerting are to send an email, run a script and/or show the alert in the Splunk alerts manager view. The alerts can be throttled so that an alert condition doesn't overwhelm the system or administrator with alerts.

8. Conclusion

Detecting security events of interest on servers and applications is a critical part of operating securely. Splunk is a powerful tool for aggregating, analyzing, visualizing and reporting on that data. Several security applications, both free and commercial, are available within Splunk to extend its effectiveness. Any number of custom searches, dashboards and alerts can be created to meet specific monitoring needs. Splunk helps system administrators to understand and correlate events across the entire infrastructure from one centralized location. Through its indexing and distributed architecture it can effectively search in a large scale data environment.

9. References

- Category: OWASP Top Ten Project. (2013, June 17). Retrieved June 20, 2013 from The Open Web Application Security Project:
https://www.owasp.org/index.php/Category:OWASP_Top_Ten_Project
- Do More with Splunk. (2013). Retrieved July 10, 2013 from <http://splunk-base.splunk.com/apps>
- Jacobson, D. (2009) *Introduction to network security*. Boca Raton, FL: Chapman & Hall/CRC.
- Mitnick, K. D., & Simon, W. L. (2006). *The art of intrusion, the real stories behind the exploits of hackers, intruders and deceivers*. Indianapolis, IN: Wiley.
- Skoudis, E., & Liston, T. (2006) *Counter Hack Reloaded: A step-by-by step guide to computer attacks and effective defenses*. Upper Saddle River, NJ: Pearson.
- Splunk App for Active Directory (2013). Retrieved July 10, 2013 from <http://splunk-base.splunk.com/apps/51338/splunk-app-for-active-directory>
- Splunk App for Enterprise Security (2013). Retrieved July 10, 2013 from <http://splunk-base.splunk.com/apps/22297/splunk-app-for-enterprise-security>
- Splunk Distributed Deployment Manual. (2013). Distributed Splunk overview. Retrieved June 20, 2013 from
<http://docs.splunk.com/Documentation/Splunk/5.0.2/Deploy/Distributedoverview>
- Splunk Enterprise Product Data Sheet. (2012) The Platform for Machine Data. Retrieved June 20, 2013 from
http://www.splunk.com/web_assets/pdfs/secure/Splunk_Product_Datasheet.pdf
- Splunk Enterprise Selected Best SIEM Appliance in 2013 SC Awards. (2013, Mar 12). Retrieved June 20, 2013 from <http://www.splunk.com/view/splunk-enterprise-selected-best-siem-appliance-in-2013-sc-awards/SP-CAAHRP>
- Splunk for Unix and Linux (2013). Retrieved July 10, 2013 from <http://splunk-base.splunk.com/apps/22314/splunk-for-unix-and-linux>

Splunk User Manual. (2013, June 8). Splunk App for Enterprise Security 2.4. Retrieved June 20, 2013 from <http://docs.splunk.com/Documentation/ES/latest/User/High-leveldashboards>

Weinberger, J., Saxena, P., Akhawe, D., Finifter, M., Shin, R., and Song, D. (2011). A Systematic Analysis of XSS Sanitization in Web Application Frameworks. In V. Atluri & C. Diaz (Eds.), *Computer Security – ESORICS 2011* (pp. 150-171). New York, NY: Springer.

Windows Security Operations Center (2013). Retrieved July 10, 2013 from <http://splunk-base.splunk.com/apps/24435/windows-security-operations-center>

XSS Filter Evasion Cheat Sheet. (2013, June 5). Retrieved June 20, 2013 from The Open Web Application Security Project:
https://www.owasp.org/index.php/XSS_Filter_Evasion_Cheat_Sheet