splunk>

Splunk for Cisco IronPort WSA: Install Guide

Table of Contents

Introduction

This manual covers the Splunk for Cisco IronPort WSA Product. This application is made up of a customized Splunk app and a Splunk server polling log data collected from an IronPort Web Security Appliance. The Splunk for Cisco IronPort WSA Reporting Application provides reports and dashboards designed to give insight into data from the IronPort Web Security Appliance (WSA).

Installation Guide

System Requirements

The installation guide will focus on an installation of Splunk running on Windows or Linux. Given the common availability of 64-bit compatible hardware, that will be considered a minimum requirement for Splunk instances. Also, for the purpose of this document, there is no support for virtualization of any core function of Splunk referenced within this document unless specifically mentioned.

Platform Requirements: Reference hardware can be commodity-grade, and must have the following minimum specifications to be eligible for Cisco support:

Intel x86-64-bit chip architecture with (2) CPU's, 4 core per CPU, 2.5-3Ghz per core

16GB RAM

(4) 300GB SAS hard disks at 10,000 rpm each in RAID10 (800 IOPS or better)

Standard 1Gb Ethernet NIC, optional 2nd NIC for a management network

Note: Splunk is often constrained by disk I/O first, so always consider that first when selecting the storage hardware. The file system will be assumed to be running on local disk volumes formatted as NTFS or EXT2/3. A separate OS volume should be created per industry best practices. The Splunk installation should reside on its own logical volume whenever possible.

Supported Browsers

As of June 1, 2011 the supported browsers are:

- Firefox 2 and 3.0.x
- Firefox 3.5 (with Splunk version 4.0.6 and later)
- Internet Explorer 6, 7 and 8
- Safari 3
- Chrome 9

An updated list may be found online at:
http://www.splunk.com/base/Documentation/latest/installation/SystemRequirements#Supported_browsers

Please see Recommended browsers and flash requirements on page 14 in this document for additional information regarding flash requirements.

Supported Operating Systems

This application is currently supported by Cisco on Redhat and Windows Operating systems. Splunk builds are available for other operating systems but they do not exist within the Cisco support matrix.

Sizing & Scaling Recommendations

The base configuration is a single-tier architecture with one server offering all 3 parts of the core functionality of a typical Splunk deployment: a search head, an indexer, and a monitor for data sources. If the estimated requirements for indexed data volume exceed 100k/Users (estimate: 100GB/day,) the Splunk infrastructure should be adjusted. By

adding another Splunk instance and adjusting the configuration, the new infrastructure would offer an increase in aggregate indexing and search performance (once the data is load-balanced), and an increase in storage and retention capacity. A dedicated forwarder server would also be added to the Splunk infrastructure and configured to monitor the WSA log files and forward the log data across multiple indexers using load balancing. To facilitate the implementation and configuration of an environment that exceeds 100k users, it is recommended that Cisco engage Splunk professional services on behalf of the IronPort WSA customer.

Based upon log volume estimates against an IronPort WSA device with 10k users, the amount of data collected is 10GB/day uncompressed. Once indexed, the data compresses to an estimated 2.5GB/day indexed storage used. The Splunk instance would retain approximately 200 days of indexed data based upon a volume size of 500GB.

IronPort users	Estimated log volume (2500 transactions/user/day)	Estimated indexed volume	Estimated retention (500GB volume)
10K users	10GB/day	2.5GB	200 days
50K users	50GB/day	13GB	40 days
100K users	100GB/day	25GB	20 days

Note that these are guidelines based upon estimated log volumes and mid-capacity drives in an array.

Splunk Install

Windows (No local PDF Server option)

Installing Splunk on Windows with the Graphical User Interface (GUI)-based installer.

Download and run the Splunk 4.2 Windows installer (http://www.splunk.com/download), click **Next** to continue until the licensing panel is reached. Select "I accept the terms in the license agreement" and click **Next** to continue installing.

When the **Destination Folder** panel is displayed, Splunk will be installed by default into `\Program Files\Splunk` on the system drive. Click **Change..** button to specify the separate data volume configured during the OS installation. Install into \\ $volume \ Splunk. Splunk's installation directory is referred to as `$SPLUNK_HOME` or `%SPLUNK_HOME%` throughout this documentation set.

The **Logon Information** panel gives the option to select the user that Splunk will run as. If Splunk is installed as the "Local System" user, Splunk will have access to all of the information on the local machine. This will not interfere with the AD/LDAP authentication configured later.

Click **Install** to proceed. The installer runs and displays the **Installation Complete** panel. The installation completes, Splunk starts, and Splunk Web launches in a supported browser if the box was checked. Or the web interface can be reached through a web browser at http://hostname:8000

Note: The first time Splunk Web is accessed after installation, login with the default username `admin` and password `changeme`. Choose 'skip' to ignore the change password prompt at this time.

Redhat Linux

Splunk 4.2 can be downloaded and installed on Redhat Linux as an RPM or a tarball package (http://www.splunk.com/download). The Splunk installation assumes the Splunk instance will run under root credentials. To run Splunk as a non-root user, first install Splunk as root. Then, **before starting Splunk for the first time**, change the ownership of the splunk directory to the desired user. The following are instructions to install Splunk and run it as a non-root user, splunk.

```
useradd splunk
groupadd splunk
chown -R splunk $SPLUNK_HOME/
```

If Splunk is run as a non-root user, make sure Splunk has the appropriate permissions to: read the files and directories it is configured to watch and write to Splunk's directory structure. Splunk's installation directory is referred to as `$SPLUNK_HOME` or `%SPLUNK_HOME%` throughout this documentation set.

To install Splunk with the RPM (defaults to /opt/splunk)

```
rpm -i --prefix=/path/splunk splunk_package_name.rpm
```

where the /path above refers to the separate data volume created.

To install Splunk with the tarball: place the file into /tmp folder and run:

```
tar –xvf /tmp/splunk_package_name.tar.gz /destination_path
```

to expand the tarball into /opt/splunk (default) or a /splunk directory created on a separate data volume.

The first time Splunk is started after a new installation, the license agreement must be accepted. To start Splunk and accept the license in one step:

```
$SPLUNK_HOME/bin/splunk start --accept-license
```

> **NB:** If you wish to have Splunk start automatically with the server you may use the enable boot-start command:
>
> `$SPLUNK_HOME/bin/splunk enable boot-start -user <splunk_user>`

Splunk Web can be found at at http://hostname:8000

Start Splunk for the first time: Windows or Linux

Set and record a new password for the Splunk local admin. This account will be used for the initial configuration and serve as a troubleshooting account should AD/LDAP authentication have issues.

Licensing Configuration

Splunk licenses are issued in capacities and are linked to the estimated users and WSA devices at a given account. The license will then be emailed to the installer and can be added into the Splunk instance.

1. In the Splunk web interface, navigate to **Manager > Licensing**
2. Click **Add license**
3. Either click **Choose file** and navigate to the license file and select it, or click **copy & paste the license XML directly...** and paste the text of the license file into the provided field
4. Click **Install**

Licenses and violations

If the initial configuration requires pre-populating Splunk with historical data, the Splunk instance may report a license violation if the historical data indexed exceeds the daily capacity of the license for a given site. A yellow banner will appear on Splunk Web “Daily indexing volume limit exceeded today. See the License Manager for details.”It is recommended that in such a case, a 90-day Eval license with a high index limit be installed into Splunk initially and used during the data import process. Once that process has been completed and tested successfully, the Eval license can be removed and an active production license should then be installed.

If licensed daily volume is exceeded on any one calendar day, there will be a violation warning. The message persists for 14 days. If there are more than 5 violations in a rolling 30-day period, search will be disabled. Search capabilities return when there have been fewer than 5 violations in the previous 30 days or when a new license with a larger volume limit is applied. Note: During a license violation period, Splunk does not stop indexing data. Splunk only blocks search access while the allowed number of license violations has been exceeded.

Further Reading:

Splunk License Installation: http://www.splunk.com/base/Documentation/latest/Admin/Installalicense

Splunk License Violations: http://www.splunk.com/base/Documentation/4.2.1/Admin/Aboutlicenseviolations

Enable SSL

To enable HTTPS through Splunk Manager, navigate to Manager: System settings: General Settings.

Select the Yes button underneath the “Enable SSL (HTTPS) in Splunk Web” setting.

Note: Restart Splunk to enable the new settings. Also, "https://" must now be appended to the URL used to access Splunk Web.

Authentication/Authorization configuration steps

Depending upon the environment, a choice should be made between deploying AD/LDAP for authentication or, on a basic access environment, using Splunk’s own local authentication.

By default, Splunk locally establishes 3 roles and creates one default user account: the ‘admin’ user. The roles in Splunk define the default app an assigned user can run, access to specific indexes, rights to change configurations, and so on. Local Splunk authentication supersedes any other authentication option configured. As such, the ‘admin’ account will be retained even if the authentication method is changed, and provides an account that can be used to configure, test, and troubleshoot a Splunk installation.

For example, if the customer wants a selection of users to have access to the Splunk Web, it would be easier to manage if the users were added to a splunk-specific group in the directory services and that group DN imported into Splunk. Evaluate the default roles in Splunk, and determine if one is a good fit for the rights that the group DN should have. If so, map the group DN to the default Splunk role (such as power) in Splunk Web: Manager: Access Controls: Configure LDAP role mapping. If not, create a new role in Splunk that has the desired rights and map that role to the group DN instead. Again, if the customer requirements are simply that a couple of chosen people can view Splunk data, then using local authentication may be sufficient.

Configure Active Directory/LDAP through the Splunk web

Before mapping the AD/LDAP settings in Splunk, figure out the user and group base DN, or distinguished name. The DN is the location in the directory where authentication information is stored. If group membership information for users is kept in a separate entry, enter a separate DN identifying the subtree in the directory where the group information is stored. If the AD/LDAP tree does not have group entries, the group base DN can be set the same as the user base DN to treat users as their own group. This requires further configuration, described later. It is recommended that the AD/LDAP Administrator be contacted for assistance.

Set up AD/LDAP via Splunk Web

First, set LDAP as the authentication strategy:

1. Click **Manager** in Splunk Web
2. Under **System configurations**, click **Access controls**
3. Click **Authentication method**
4. Select the **LDAP** radio button
5. Click **Configure Splunk to work with LDAP**
6. Click **New**
7. Enter an **LDAP strategy name** for the configuration (ex. Domain_config)
8. Enter the **Host** name of the AD/LDAP server. Be sure that the Splunk server can resolve the host name
9. Enter the **Port** that Splunk should use to connect to the AD/LDAP server

By default AD/LDAP servers listen on TCP port 389

LDAPS (LDAP with SSL) defaults to port 636

10. To turn on SSL, check **SSL enabled**

Important note: SSL must be enabled on your AD/LDAP server

11. Enter the **Bind DN**

This is the distinguished name used to bind to the AD/LDAP server

This is typically the administrator or manager user. This user needs to have access to all AD/LDAP user and group entries you want to retrieve

Leave blank if anonymous bind is sufficient

12. Enter and confirm the **Bind DN password** for the binding user

13. Specify the **User base DN**. You can specify multiple user base DN entries by separating them with semicolons

Splunk uses this attribute to locate user information

Note: You must set this attribute for authentication to work

14. Enter the **User base filter** for the object class you want to filter your users on

Note: This is recommended to return only applicable users. For example, (department=IT)

Default value is empty, meaning no user entry filtering

15. Enter the **User name attribute** that contains the user name

Note: The username attribute cannot contain whitespace. The username must be lowercase

In Active Directory, this is sAMAccountName

The value uid should work for most other configurations

16. Enter the **Real name attribute** (common name) of the user

Typical values are displayName or cn (common name)

17. Enter the **Group mapping attribute**

This is the user entry attribute whose value is used by group entries to declare membership

The default is dn for active directory; set this attribute only if groups are mapped using some other attribute besides user DN

For example, a typical attribute used to map users to groups is uid

18. Enter the **Group base DN**. You can specify multiple group base DN entries by separating them with semicolons

This is the location of the user groups in AD/LDAP

If your AD/LDAP environment does not have group entries, you can treat each user as its own group:

- Set groupBaseDN to the same value as userBaseDN. This means you will search for groups in the same place as users
- Next, set the groupMemberAttribute and groupMappingAttribute to the same attribute as userNameAttribute. This means the entry, when treated as a group, will use the username value as its only member
- For clarity, you should probably also set groupNameAttribute to the same value as userNameAttribute

19. Enter the **Group base filter** for the object class you want to filter your groups on

    Note: This is recommended to return only applicable groups. For example, (department=IT)

    Default value is empty, meaning no group entry filtering

20. Enter the **Group name attribute**

    This is the group entry attribute whose value stores the group name

    This is usually cn

21. Enter the **Group member attribute**

    This is the group attribute whose values are the group's members

    This is typically member or memberUid

Note: When you save the configuration, Splunk will attempt the AD/LDAP connection immediately. An error will be displayed as a red bar on the top of the Splunk Web page (ex. Encountered the following error while trying to update: In handler 'LDAP-auth': Error binding to LDAP: Can't contact LDAP server) A successful configuration or update will display a light blue bar stating a change or update was performed.

Map existing AD/LDAP groups to Splunk roles

Once you have configured Splunk to authenticate via your AD/LDAP server, map your existing AD/LDAP groups to any roles you have created in Splunk Web: “Manager: Access Controls: Authentication Method: Configure LDAP role mapping.” If you do not use groups, you can map users individually.

Note: You can map either users or groups, but not both. If you are using groups, all users you want to access Splunk must be members of an appropriate group. Groups inherit capabilities from the highest level role they're a member of.

All users that can login are visible in the **Users** page in Splunk Manager: Access Controls. All users imported via AD/LDAP will be labeled with “Authentication System: LDAP.” Remember, local authentication (Splunk) trumps AD/LDAP. A user with the same ‘name’ in both local and LDAP will always use the local password and assigned rights.

Test your AD/LDAP configuration

If you find that your Splunk install is not able to successfully connect to your AD/LDAP server, try these troubleshooting steps:

Check $SPLUNK_HOME/var/log/splunk/splunkd.log for any authentication errors.

Remove any custom values you've added for **userBaseFilter** and **groupBaseFilter**.

In linux, Perform an ldapsearch to confirm that the variables you are specifying will return the expected entries:

```
ldapsearch -h "<host>" -p "<port>" -b "<userBaseDN>" -x -D "<bindDN>" -W
"realNameAttribute"  ldapsearch -h "<host>" -p "<port>" -b "<groupBaseDN>" -x -D
"<bindDN>" -W "groupNameAttribute"
```

For Active Directory, verify the group name(s) and users by using the dsquery command:

Ex. dsquery group -name splunk* | dsquery user

Returns the DN of all users in any group name beginning with splunk.

Department membership query configuration

This query is run against AD/LDAP to output a .csv file that is used in the reports extensively to filter searches based upon Splunk roles that match to Groups in AD/LDAP. If the implementation utilizes AD/LDAP groups bound to roles in Splunk, and the customer plans to give views into some of the data based upon organizational roles, then the script that polls for changes in those selected group DN should be configured and enabled.

Using a text editor, open the file:

For Red Hat Enterprise Linux:

```
$SPLUNK_HOME/etc/apps/SplunkforCiscoIronportWSA/bin/discovery.py
```

For Windows:

```
X:\$SPLUNK_HOME\etc\apps\SplunkforCiscoIronportWSA\bin\discovery.vbs
```

Edit the first 4 fields at the top of the header to facilitate a proper connection to AD/LDAP: AD/LDAP host, service account, service account password, and the Group Base DN's used for AD/LDAP import earlier.

```
strComputer = 'ad_ldap_host'
strUser = 'cn=service_account,cn=Users,dc=my_directory,dc=net'
strPassword = 'service_account_password'
strGroupOUs = 'Group base DN;Group base DN;Group base DN'
```

Save the file.

Enable the scripted input. Using a text editor, open the file:

$SPLUNK_HOME/etc/apps/CiscoforIronportWSA/local/inputs.conf

Find the stanza for the appropriate scripted input, and enable it.

```
# membership script Windows
```

or

```
# membership script Linux
disabled = false
```

Note: The script is set to run every day by default. The interval is set in seconds and can be changed as per the deployment requirements. Verify that the user data has been populated in the:

$SPLUNK_HOME/etc/apps/CiscoforIronportWSA/lookups/departments.csv

file before attempting to view any reports.

Note: The *discovery.py* script for Red Hat Enterprise Linux depends on the installation of the *ldapsearch* tool. If not already installed, the tool can be added by executing the following command: *sudo yum install openldap-clients*

Department membership role configuration

Once the department membership query configuration is complete Splunk will automatically begin associating departments with user IDs in accesslogs. Roles may be created to advantage of this mapping by restricting specific users to only see reports for a specific department.

Add and edit roles using Splunk Web

In Splunk Web:

1. Click **Manager**

2. Click **Access controls**

3. Click **Roles**

4. Click **New** or edit an existing role

5. Specify new or changed information for this role. In particular, you can:

restrict what data this role can search with a search filter. E.g. if there exist a role called "sales" type "department=sales" into the search filter

6. Click **Save**

Cisco IronPort WSA App Installation

The Cisco IronPort WSA app will not be hosted on Splunkbase with other Splunk apps. Because of this, the WSA app installation package will be provided to the installer manually.

From Splunk Web: Manager: Apps: Install App from File. Browse to the WSA app zip/tar file you received earlier and select it. Once Splunk reports a successful import, restart Splunk (Manager: Server Controls: Restart.)

Log into Splunk Web, go to Manager: Apps and verify app the 'Splunk for Cisco IronPort WSA' app is visible and enabled

Configure IronPort WSA log source

The configuration for end-to-end management of the Cisco IronPort WSA logs will depend upon the customers' environment and requirements for managing their logs. The installer needs to consider the method and manageability of the log file process before designing a solution with the customers' requirements in mind. The IronPort WSA app is designed by default to accept log files placed into a local directory on the Splunk server. The app also assumes the IronPort hostname will be inline with the log file path and uses that hostname to differentiate sources when multiple IronPort devices are deployed. **NB: Finally, the inputs are configured to delete the log file placed into the folder once it has been read. It is important that Splunk NOT be used as the primary log storage for this reason. Th**ere are 2 sets of logs the app is interested in: Traffic Monitor and Access logs.

The WSA app configuration file defines a default structure: /$Input_base/host_name/accesslogs/ and /$Input_base/host_name/trafmonlogs/. The first-level folder name ($input_base) will need to be changed to match the chosen deployment and the /host_name/ changed to match the IronPort device name.

Note: ***DO NOT*** pre-load the IronPort WSA log files into the folders during the configuration stage.

Create and document the chosen folder structure, making certain that the IronPort host name appears as the 2nd folder in the path. In Splunk Web: Manager: Data Inputs: Files and Directories. Find any inputs labeled SplunkforCiscoIronportWSA and disable them. The inputs paths will be adjusted through the configuration file directly, and not through Splunk Web. Copy the file:

$SPLUNK_HOME/etc/apps/CiscoforIronportWSA/default/inputs.conf

To the folder:

$SPLUNK_HOME/etc/apps/CiscoforIronportWSA/local/

Using a text editor, open the file $SPLUNK_HOME/etc/apps/CiscoforIronportWSA/local/inputs.conf

There are two input stanzas, one for each log source. In a multi-IronPort deployment, there would be two inputs per IronPort host. The only field to be concerned with is the path located in the header.

```
[batch:///inputs_target/.../trafmonlogs/...*]
## l4tm logs
host_segment = 2
disabled = true
sourcetype = wsa_trafmonlogs
move_policy = sinkhole
crcSalt = <SOURCE>
```

For a Linux-based Splunk host, adjust the path using the appropriate slashes:

```
batch:///data_store/host_name/accesslogs
```

For a Windows-based system file path, edit the header with the appropriate slashes:

```
batch://X:\data_store\host_name\accesslogs
```

Update both inputs to the appropriate paths. Finally, change:

```
disabled = true
```

to

```
disabled = false
```

Save the file and restart Splunk. Though Splunk Web: Manager: Data Inputs: Files and Directories, verify that the inputs are listed, enabled, and have the correct path.

Please defer the initial loading of data until after a successful import of historical data.

Historical data import

Historical IronPort WSA data can be loaded and indexed into Splunk. The following instructions do not deal with live data or day-to-day functionality. This is a one-time process to include historical data. This application leverages summary indexes to improve performance. These summarizations are built regularly as part of normal operation. When historical data is loaded – we must manually tell Splunk to create this summary information for the hiostrical data.

This is a two-step process. The historical log files need to be placed into the appropriate host and log type folders created earlier. From there the logs will be pulled into and indexed by Splunk. Note: Any logs placed into those folders WILL be deleted after the data is indexed by default.

Once the import completes, a summarization process is manually run against the historical data in Splunk. To trigger the summarization process, from a command prompt run the file:

```
$SPLUNK_HOME/etc/apps/SplunkforCiscoIronportWSA/bin /summary.sh
```

For Windows:

```
X:\$SPLUNK_HOME\etc\apps\SplunkforCiscoIronportWSA\bin\summary.vbs
```

Type in (or browse) to the splunk folder and enter the local Splunk admin credentials when prompted. While the process is running, the description of the saved searches that the data is being summarized for will be visible. Please note that the summary is responsible for the data in most stock reports. Therefore, if Cisco IronPort WSA app is running you will see no results until the process completes.

The summary job is estimated to take about 4 minutes per 5M events (2GB of raw data) per summary job based upon the platform hardware recommendations. ex. A 10GB file representing 25M historical events is estimated to take 20 minutes to run against each summary job. There are 27 summary jobs used by the reporting built into SplukforCiscoIronportWSA app. So the historical summary can take up to 9 hours to complete.

The default for the summary script is to summarize up to 90 days of history.

Verify data is being imported: In Splunk Web, login as admin. Go to the search app. Go to Status; Index Activity; Index Activity Overview report. Look for summary index growth.

If the historical data import was run under a Splunk Eval license, install the Enterprise default license downloaded for the account and remove any non-Production licenses.

User Navigation

Set SplunkforCiscoIronportWSA app as the default app for all users/roles

Login to the Splunk Web as admin or an admin equivalent.

Go to Manager: Access Controls: Roles.

Identify the roles that will be utilizing the SplunkforCiscoIronportWSA app exclusively.

Click on the link to the role. The very first option is to set the Default app. Select SplunkforCiscoIronportWSA from the dropdown list and scroll to the bottom of the page and save. Cycle through and change the rest of the chosen roles. The results will be displayed on the main Roles page as “Default app.”

Verify users have their default app set appropriately by verifying the “Default app” field in Manager: Access Controls: Users.

Deployment customization options

PDF Report Server and email instructions

Splunk Web users can generate a scheduled PDF output from any dashboard, view, search or report. To enable this functionality, the PDF Report Server app must be downloaded from Splunkbase and installed into a Splunk instance on a single Linux host. In addition, an internal email server will be configured in Splunk to allow it to send the PDF reports.

Note: Splunk PDF reporting is not yet available on the Windows platform. Currently, there must be a Linux-based instance of Splunk running on the network to support scheduled PDF reporting. For a minimal installation: a standard Linux image, with an installation of Splunk configured as a forwarder (no indexing or web interface required) with the PDF Server app installed can serve multiple Splunk instances for PDF generation.

1. Download and install the PDF Report Server add-on. Click **Browse more apps** in Launcher, or download it separately from Splunkbase here

2. Ensure that the Xvfb X server, xauth and fonts for your Linux distribution are installed. These are included with most Linux distributions, but not installed by default. On Redhat/CentOS/Fedora, type:

```
yum install Xvfb xauth bitstream-vera-fonts
```

3. Launch Splunk Web on the Linux host and navigate to **Manager**

4. Navigate to **System Settings > Email Alert Settings**

5. Check the **Use PDF Report Server** box

6. Click **Save**

Under Mail server settings you can enter or update information related to the SMTP server that Splunk interacts with in order to send out alert emails. Identify the SMTP mail host server. Provide an authentication username/password if the SMTP server requires them. Optionally specify that Splunk use SSL or TLS when it communicates with the SMTP server.

Under Email format you can enter information about the format of the emails that Splunk sends. You can define the name that appears in the "sender" field (by default it is Splunk), and you can set up the format of the email subject line (by default it is Splunk Alert: $name$, where $name$ is the name of the search that the alert is based upon). You can also set at the Manager level the default email format for all alerts and whether or not alert emails provide inline results.

The local host is now configured to generate PDFs for Splunk Web users or can accept jobs from other Splunk hosts with PDF printing enabled.

Note: If the hostname of the Splunk Web instance that this PDF Report Server will talk to is not resolvable in DNS, enter its IP address or a hostname that resolves to that IP in the **Link hostname** field. This will ensure that Splunk Web can contact the PDF Report Server, and that links sent in emailed PDF reports work correctly. If the field is left empty, Splunk will try to autodetect the hostname.

Time configuration notes

Splunk internally operates in UTC/GMT/Zulu time and displays data based upon the local time of the server Splunk is running on. For example, if the Splunk instance is running on a server set to a timezone of America/Los Angeles, the

data will be normalized in the search results to PST/PDT time. The time displayed in the search results will always reflect the 'local' time set where the Splunk instance runs. There is currently no configuration option that would allow Splunk to generate reports showing a different time zone.

To maintain consistency across devices/inputs when displaying the search results, Splunk must know what timezone the source data was created in. By default, all Splunk inputs for the Cisco IronPort WSA logs are set to TZ = GMT. If the IronPort devices are not set to GMT, or are not set consistently across IronPort devices that will be sending data to Splunk, the installation requirements must be re-evaluated. If it's not otherwise specified, Splunk assumes that the data was created in the time zone where the Splunk instance resides.

To update the timezone, edit the props.conf file and update TZ=GMT to TZ=???

??? may be cross referenced from the zoneinfo list, a copy of which may usually be found here: https://secure.wikimedia.org/wikipedia/en/wiki/List_of_zoneinfo_timezones per

Custom web ports:

Should the installation require the Splunk Web service or the splunkd service to utilize a different port, the defaults can be changed.

To change the splunk web service port: From the %SPLUNK_HOME%\bin directory: `splunk set web-port ####`

To change the splunk core service port: From the %SPLUNK_HOME%\bin directory: `splunk set splunkd-port ####`

Establish IronPort log transfer into Splunk

A push from IronPort to Splunk should be established at 60 minute (max) increments into the folder structure established in the Historical data import section of this document (page 10). This is accomplished by creating a new log subscription (as depicted in Figure 1) using the parameters in Table 1. FTP is the supported transport solution.

Setting	Value
Log directory	accesslogs
Log Style	Squid
Custom Fields (optional)	%XK NB: %XK is needed to add the web reputation threat reason to the logs
Filename	<user preferred>
Maximum File Size	This should be set to produce logs NO GREATER THAN 1 hour increments. If logs contain data older than one hour – that data will not appear in reports not produced every hour – data older. Maximum Files size is recommended to be no more than 500Mb. Maximum Time Interval Between Transfers is recommended to be 3600 seconds (1 hours). The WSA will push the accesslogs to the ftp server at a minimum every hour, or immediately if they file exceeds 500Mb.

Table 1

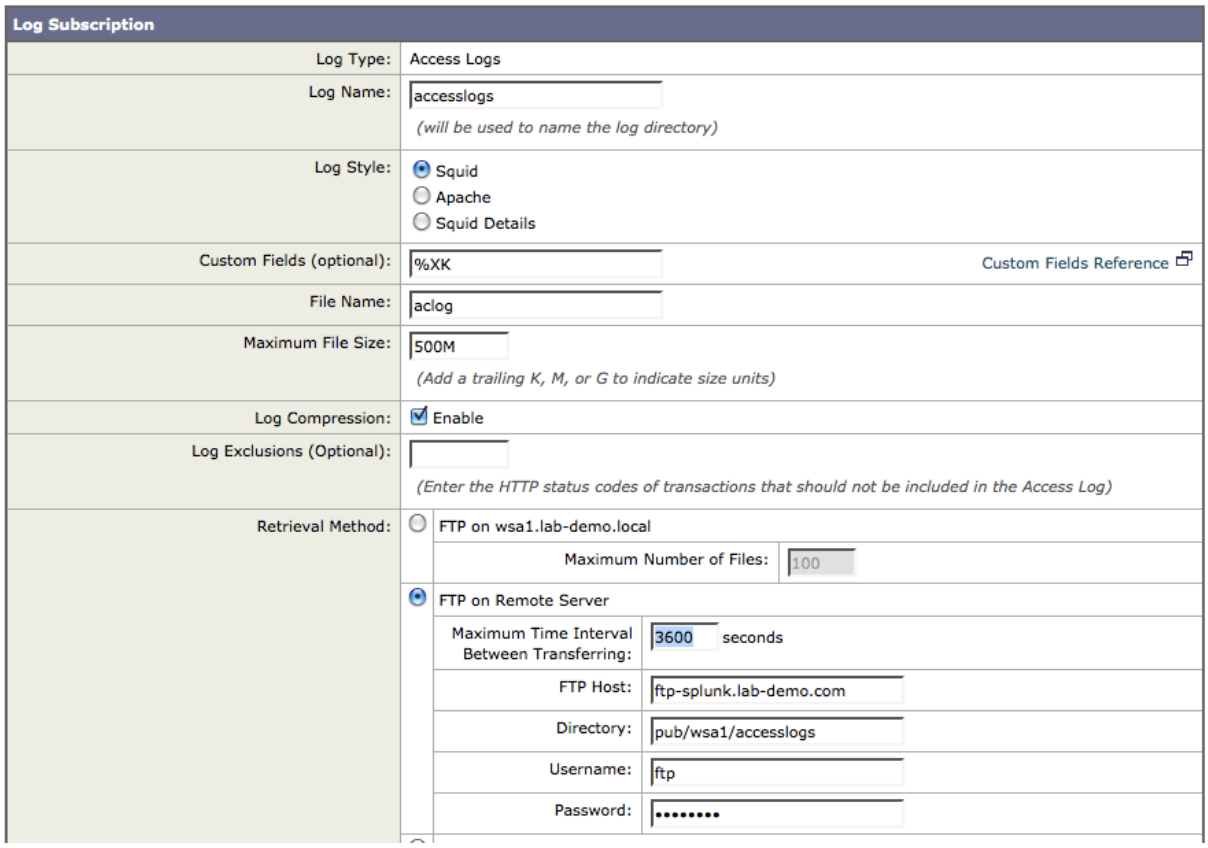

Figure 1

Setting	Value
Log directory	trafmonlogs
Log Style	L4TM
Filename	<user preferred>
Maximum File Size	This should be set to produce logs NO GREATER THAN 1 hour increments. If logs contain data older than one hour – that data will not appear in reports not produced every hour – data older. Maximum Files size is recommended to be no more than 100Mb. Maximum Time Interval Between Transfers is recommended to be 3600 seconds (1 hours). The WSA will push the accesslogs to the ftp server at a minimum every hour, or immediately if they file exceeds 100Mb.

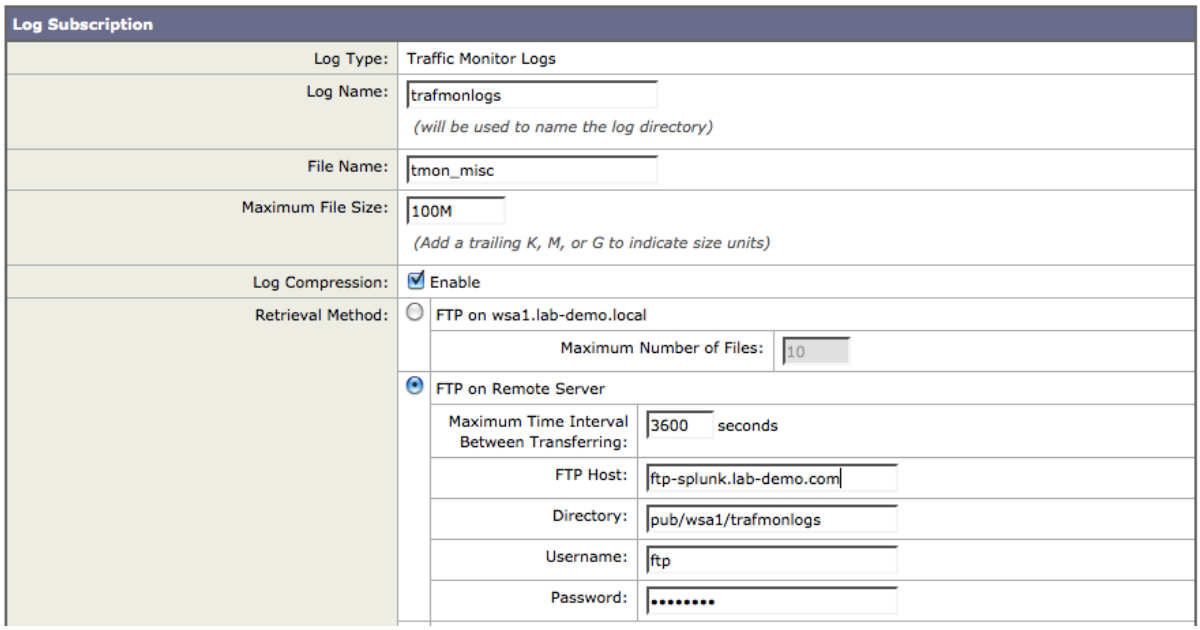

Figure 2

Recommended browsers and flash requirements

As of Splunk version 4.2, the supported browsers are: Firefox 2, 3.0.x, and Firefox 3.5; Internet Explorer 6, 7 and 8; Safari 3; and Chrome 9.

Flash Player 9 (or above) is required.

Configuration best practices

Verify Splunk services are set to restart automatically and test them. Verify the time zone set on the machine running the Splunk instance. Document the local admin account password (regardless of the chosen authentication method).

Vmware Whitepaper from Splunk

Cisco does not recommend nor support the deployment of this tool in a virtualized environment. However, in the event that a virtualized non-production deployment is still desired due a lab or demo environment, more information can be found at:
http://www.splunk.com/page/securelink/signup/Splunk_and_VMware_VMs_Tech_Brief

Under the section labeled Technical Briefs is a whitepaper on VMware deployments of Splunk. A splunk.com login is required.

Data retention advanced options

The topics of data backups and retention are complex and require a deeper understanding of how Splunk functions. Reading the following pages of the standard Splunk documentation is highly recommended to understand the options.

Backup of Splunk data: http://www.splunk.com/base/Documentation/latest/admin/Backupindexeddata

How to archive data: http://www.splunk.com/base/Documentation/latest/Admin/Automatearchiving

Once these have been reviewed, a recommendation can be made to the customer about implementing an archiving policy that will cover data retention while establishing whether access to historical data is supported in a given solution.