FortiSIEM - User Guide

VERSION 4.10.0

FORTINET DOCUMENT LIBRARY

http://docs.fortinet.com

FORTINET VIDEO GUIDE

http://video.fortinet.com

FORTINET BLOG

https://blog.fortinet.com

CUSTOMER SERVICE & SUPPORT

https://support.fortinet.com

http://cookbook.fortinet.com/how-to-work-with-fortinet-support/

FORTIGATE COOKBOOK

http://cookbook.fortinet.com

FORTINET TRAINING SERVICES

http://www.fortinet.com/training

FORTIGUARD CENTER

http://www.fortiguard.com

FORTICAST

http://forticast.fortinet.com

END USER LICENSE AGREEMENT

http://www.fortinet.com/doc/legal/EULA.pdf

FEEDBACK

Email: techdocs@fortinet.com

Dec 13, 2017

FortiSIEM 4.10.0 User Guide

Revision 5

TABLE OF CONTENTS

Change Log

Date	Change Description
2017-09-09	Initial version of FortiSIEM User Guide for 4.10.0.
2017-09-12	Revision 2 with updates in the section 'Upgrade'.
2017-09-25	Revision 3 with 'Common Vulnerabilities and Exposures' table added under 'Resolved Issues'.
2017-09-26	Revision 4 with two new sections added 'License notes' and 'Upgrade notes'.
2017-10-17	Revision 5 with updated section *Managing FortiSIEM > Integrating with External CMDB and Helpdesk Systems*. Installation and Upgrade information is available in a separate manual under http://docs.fortinet.com/fortisiem/admin-guides.

What's new in 4.10.0

This section describes the new and enhanced features in FortiSIEM version 4.10.0.

Refer to FortiSIEM 4.10.0 Release Notes for more information.

FortiSIEM license key generation has been improved. Starting September 11, 2017, FortiCare will generate licenses ONLY using the new license key generation algorithm and will only work with FortiSIEM 4.10.0 and later releases.
This has the following implications:

- **If you want a new FortiCare license key, for any reason (for example, to extend maintenance and support or buy additional device license), you must upgrade to 4.10.0 first and then apply the license. This is only specific to 4.10.0 release since license generation keys are modified. Future releases will likely not require this.**
- **If you are upgrading to FortiSIEM 4.10.0, you must get a newly generated license key before starting the upgrade process.**

Note the following FortiSIEM Cluster Upgrade process:

1. Upgrade Supervisor.
2. Upgrade Workers.
3. Apply 4.10.0 license key.
4. Upgrade Collectors.

Features

FortiSIEM 4.10.0 release includes the following new features:

Unlimited Unmanaged CMDB Devices

In prior releases, every device in CMDB costs a device license. A newly discovered device may not be added to CMDB if the device license has exceeded.

Starting with 4.10.0 release, the concept of 'Managed' and 'Unmanaged' CMDB device is introduced. A device directly sending log to FortiSIEM or being monitored by FortiSIEM for availability/performance/configuration change is considered a Managed Device. The device license only affects Managed Devices, in other words, you can have unlimited unmanaged devices in CMDB. During discovery, the user can choose the discovered devices to be either managed or unmanaged. If a managed device is newly discovered but the device license is exceeded, the device will be entered in the CMDB but as an Unmanaged device. The user can flip a device state from Unmanaged to Managed and vice versa.

This feature enables customers to capture their entire infrastructure in FortiSIEM CMDB without necessarily paying for a large license at the beginning. Customers can then incrementally convert devices from Unmanaged to Managed and increase device license. This provides the customer greater deployment flexibility.

EPS Bursting Using Unused EPS

Prior to 4.10, FortiSIEM dropped events that exceeded 110% of licensed EPS. Starting with the 4.10 release, FortiSIEM will add a 'reservoir' of events to reduce the possibility of events being dropped. This allows greater EPS bursts without event loss.

The process is as follows:

1. FortiSIEM will keep track of Unused EPS during a day. FortiSIEM will add 50% of unused EPS from the previous day at midnight to the EPS reservoir. Unused EPS is the difference between Licensed EPS and used EPS as measured throughout the 24 hour day. Over time, the reservoir approaches average unused EPS in a 24 hour day. This provides a significant buffer to handle EPS bursts.
2. FortiSIEM will let customer EPS burst up to five times Licensed EPS using the currently accumulated Unused EPS. Note that borrowing future EPS is not allowed as it may lead to EPS starvation.

To benefit from this EPS Bursting, the customer must have compute, storage IOPS and storage space infrastructure available to process the extra EPS. The system must be provisioned for five times Licensed EPS to handle potential event surge.

Enhancements

FortiSIEM 4.10.0 release includes the following enhancements:

- Optimized Dynamic EPS Re-allocation Algorithm based on incoming EPS at the Collectors. Previously, this was based on guaranteed EPS. This design change allows quicker re-adjustment and reduces event loss.
- HTML5 Incident and Dashboard is redesigned for better viewing via tablets.
- Query Engine is made robust to handle event database index corruptions.
- Data Manager and Indexing engine are now more robust and efficient and provides sustained EPS insertion. A new report now captures the event database insert rate.
- Optimized Domain Generation Algorithm (DGA) by including white list and algorithmic optimizations. This algorithm can detect connectivity to malware domains that are likely to have dynamically generated domain names.
- Fortinet-centric dashboard is added in Flex version. Similar dashboard already exists in the HTML5 version of FortiSIEM 4.9.0.

New Rules

FortiSIEM 4.10.0 introduces new rules under the following categories:

- Windows Sysmon log based attack detection
- Linux syslog based attack detection
- FortiGate, FortiSandbox and FortiMail Rules
- FortiSIEM own analytics

Windows Sysmon based attack detection

These rules need Windows sysmon logs to be collected via FortiSIEM Advanced Windows Agent.

1. **Windows Malicious Service Installed** - Detects known malicious service installs that only appear in cases of lateral movement, credential dumping, and other suspicious activity.
2. **Windows DSRM Password Change** - Detects password change of the Directory Service Restore Mode (DSRM) account. This is a local administrator account on Domain Controllers. Attackers may change the password to gain persistence.
3. **Windows Suspicious Logon Failures** - Detects suspicious logon failures for the following reasons - account disabled, time of day violation, forbidden logon, error during logon, account locked out.
4. **Windows Server Console Logon** - Detects interactive console logon to windows server.
5. **Windows Backup Catalog Deleted** - Detects suspicious windows backup catalog deletions.If you delete the backup catalog for a computer, you will not be able to access the backups created on that computer using the Windows Server Backup snap-in.
6. **Windows Password Dumper Activity** - Detects process handle on LSASS process with certain access mask and object type.
7. **Windows Administrator User Reconnaissance Activity** - Detects attempts to enumerate administrative users via commands such as "net user administrator /domain" and "net group domain admins /domain".
8. **Windows LSASS Process Access** - Detects process access to LSASS which is typical for malware. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
9. **Windows Remote Thread in LSASS** - Detects remote thread creation to the lsass.exe process - likely an attempt to collect passwords. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
10. **Windows PowerShell Download from URL** - Detects a PowerShell was launched containing download commands in its command line string. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
11. **Windows PowerShell Opening External Connections** - Detects a PowerShell opening external connections to external IP addresses. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
12. **Windows PowerShell using Suspicious Parameters** - Detects a PowerShell was launched containing download commands in its command line string. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent
13. **Windows Suspicious PowerShell Invocations** - Detects suspicious powershell innovation from scripting parent processes. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
14. **Malicious HTML Applications Spawning Windows Shell** - Detects a Windows command line executable started from Malicious HTML Applications. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
15. **Windows Office Macro Spawning shell** - Detects a Windows command line executable started from Microsoft Office applications. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
16. **Suspicious Windows Driver Load from Temp Directory** - Detects a windows driver load from a temporary directory. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
17. **Windows Code Execution in Non-Executable Folder** - Detects attempt to run code from uncommon folders. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
18. **Windows Code Execution in Webserver Root Folder** - Detects a suspicious program execution in a web service root folder. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.

19. **Windows Command Line Processes Started by MMC** - Command line processes started by MMC could be a sign of lateral movement using MMC application COM object. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
20. **Windows Net.exe Processes Started** - Net.exe Processes started. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
21. **Windows Network Connection from Suspicious Program Locations** - Detects suspicious network connections from suspicious program locations. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
22. **Windows BITSAdmin download** - Detects a BITSAdmin tool downloading a file. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
23. **Windows DHCP Callout DLL installation** - Detects the installation of a Callout DLL via CalloutDlls and CalloutEnabled parameter in Registry, which can be used to execute code in the context of the DHCP server. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
24. **Windows DNS ServerLevelPluginDll Install** - Detects the installation of a plugin DLL via ServerLevelPluginDll parameter in Registry, which can be used to execute code in context of the DNS server. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
25. **Windows WScript or CScript Dropper** - Detects wscript/cscript executions of scripts located in user directories. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
26. **Windows Certutil Decode in AppData** - Detects a Microsoft certutil execution with the decode sub command. Sometimes this is used to decode malicious code with the built-in certutil utility. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
27. **Windows Command With Suspicious URL and AppData String** - Detects suspicious command line execution with URL and AppData string in parameters. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
28. **Suspicious Control Panel DLL Load** - Detects suspicious Rundll32 execution from control.exe as used by Equation Group and Exploit Kits. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
29. **Windows Suspicious Command Line Reconnaissance Activity** - Detects suspicious command line based reconnaissance activity. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
30. **Windows Suspicious Regsvr32 Activity** - Detects suspicious Regsvr32 activity - this executable is used to register and unregister OLE controls such as DLLs and ActiveX controls in the Windows registry. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
31. **Windows Scheduled Tasks in User Session** - Detects the creation of scheduled tasks in a user session. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
32. **Windows Suspicious File Execution By Scripts** - Detects suspicious file execution by wscript and cscript executables
33. **Windows Suspicious Password Hash Retrieval** - Detects suspicious commands that could be related to activity that uses volume shadow copy to steal and remotely retrieve hashes from the NTDS.dit file. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
34. **Windows Suspicious WMI execution** - Detects suspicious WMI commands to get information from the system. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
35. **Windows User Account Control Bypass via Event Viewer** - Detects User Account Control bypass method using Windows event viewer. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
36. **Windows User Account Control Bypass via sdclt** - Detects User Account Control bypass method using Windows backup tool sdclt.

37. **Windows Java running with remote debugging** - Detects a JAVA process running with remote debugging allowing more than just localhost to connect. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
38. **Windows Web shell Reconnaissance** - Detects certain command line parameters often used during reconnaissance activity via web shells. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
39. **Windows Web Servers spawning command shell** - Detects web servers spawning command shell - could be the result of a successfully placed web shell or an other attack. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
40. **Windows Wannacry ransomware activity** - Detects Wannacry ransomware activity. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
41. **Windows NotPetya ransomware activity** - Detects NotPetya ransomware activity. This rule requires Windows sysmon events collected via FortiSIEM Advanced Windows Agent.
42. **Windows Application Audit log cleared** - Detects an admin or an application has cleared the specified event log. Clearing the event logs may indicate a malicious activity so the admin should make sure that this is indeed a legit action.

Linux syslog based attack detection

1. **Suspicious Linux log entries** - Detects suspicious entries in Linux syslog.
2. **Shellshock Expression in Log Files** - Detects Shellshock expression in Linux syslog.
3. **Suspicious Linux SSHD Errors** - Detects suspicious SSH / SSHD error messages that indicate a fatal or suspicious error that could be caused by exploit attempts.
4. **Suspicious Linux VSFTPD Errors** - Detects suspicious VSFTPD error messages that indicate a fatal or suspicious error that could be caused by exploit attempts.

FortiGate, FortiSandbox and FortiMail Rules

These rules will trigger based on log received from these devices.

1. **Web Traffic to FortiSandbox Malicious URLs** - Detects HTTP traffic where the URL matches the external FortiSandbox Malicious URL list.
2. **Web Traffic to FortiGuard Malicious URLs** - Detects HTTP traffic where the URL matches the external FortiGuard Malicious URL list.
3. **FortiSandbox detects high/medium risk file malware** - FortiSandbox claims that a file has malware.
4. **FortiSandbox detects unknown risk file malware** - FortiSandbox claims that a file has malware but the risk is unknown and so needs to be investigated further.
5. **FortiSandbox detects Botnet** - FortiSandbox has detected a botnet.
6. **FortiGate detects Botnet** - FortiGate detected a botnet.
7. **FortiSandbox detects Malware URL** - FortiSandbox claims that a file has malware.
8. **FortiSandbox detects Network Attack** - FortiSandbox has detected a network attack.
9. **FortiSandbox detects multiple attacks from same source** - FortiSandbox detects multiple attacks from the same source.
10. **FortiSandbox detects multiple hosts with infected files** - FortiSandbox detects multiple hosts infected with the same infected file indicating an outbreak.
11. **Virus found in mail** - FortiMail and other mail gateways found a virus in mail attachment.
12. **Host Quarantined by NAC** - Detects hosts quarantined by FortiGate NAC.

FortiSIEM own analytics

These rules will trigger based on network traffic or web proxy logs that directly contain destination FQDN or destination IP addresses can be mapped to destination FQDN:

1. **Dynamically generated host name - malware likely**- Detects algorithmically generated host name in network traffic - malware often uses algorithmically generated host names to communicate.

FortiSIEM Basics

These topics provide an overview of the FortiSIEM solution, including its component and various deployment configurations.

Features and Architecture

FortiSIEM provides an all-in-one, seamlessly integrated and service-oriented IT infrastructure monitoring solution that covers performance, availability, change, and security monitoring aspects of network devices, servers, and applications. It is offered in two versions:

- A VMware based virtual appliance, which you can deploy as a single appliance or a cluster of virtual appliances in a highly available, scaled-out grid architecture. This is what we refer to as FortiSIEM Enterprise.
- Software-as-a-Service (SaaS), where you deploy a Collector virtual on-premises for a customer, and all of the customer data is transmitted to FortiSIEM data center. This is what we refer to as FortiSIEM Multi-Tenant, since collector deployments are commonly used by organizations such as Managed Service Providers to monitor the services of their customers.

Some of the features of the FortiSIEM monitoring solution include:

- Intelligent Device Discovery
- Analytics
- Business Services
- Architecture

Intelligent Device Discovery

The first step in the monitoring process is IT infrastructure discovery. FortiSIEM has a fast and intelligent discovery engine that can automatically crawl an IT infrastructure and discover network devices, servers, and applications in depth. The user needs to provide appropriate credentials, a discovery IP address range, and optionally a starting router IP address for faster discovery.

A wide range of information is discovered including hardware information, serial numbers and licenses, installed software, running applications and services, and router configuration. The discovered devices are automatically categorized into detailed functional groups, such as Routers/Switches, Firewalls, and Network IPS, and this information is maintained within an integrated configuration management database (CMDB). Some special relationships are also discovered, for example WLAN Access Points to WLAN Controllers, VMware guests to physical hosts, etc. The CMDB is kept up to date through user-defined scheduled discoveries and FortiSIEM listening to changes as part of performance monitoring.

A novel aspect of FortiSIEM discovery is that those aspects of a device that can be monitored are also discovered at the same time. For example, given SNMP, WMI, and JDBC credentials for a Windows server, FortiSIEM might discover the following:

- System performance metrics that can be collected by SNMP, for example CPU, memory utilization, and disk space utilization
- System performance metrics that can be collected by WMI, for example Disk I/O utilization, memory swap rates, and process utilization
- Application specific metrics that can be collected by WMI, for example IIS, DNS, DHCP, and Exchange metrics
- Event logs that can be collected by WMI
- Database logs that can be pulled from the server by JDBC

You simply approve the discovered results and monitoring begins. This approach reduces human error, since FortiSIEM learns from the true device configuration state.

Analytics

FortiSIEM uses a unified event-based framework to analyze all data including logs, performance monitoring data. Logs can either be sent to FortiSIEM via Syslog, SNMP traps, or other common log shipping methods, or FortiSIEM can periodically access the system and collect the logs. Performance monitoring data is collected by periodically probing the system. The data is parsed, indexed, and stored in a proprietary flat-file based database. In contrast, the CMDB information is stored in a PostgreSQL relational database. FortiSIEM unified data management architecture combines the two databases and presents a single view to the user.

FortiSIEM provides a broad range of metrics. First, it is possible to search all data based on keywords or in a structured way using the attributes parsed by AcceOps. The search can be done in real time, in which the data streaming in from devices is displayed, or the search can be based on historical data. Historical data is referred to as Reports in FortiSIEM, and can be scheduled to run at intervals you set. A large number of reports are provided in a categorized fashion, based on device type, and also based on functionality such as availability, performance, change and security. Two novel aspects of FortiSIEM metrics include unification and drill-down capabilities. With unification, all the data is analyzed and presented the same way, whether it be real time search, reports, rules or performance, availability, or change or security data. By using drill-down you can start from a specific context, such as Top Authentication Failed Users, and iteratively select attributes to further analyze data and get to the root cause of a problem. As an example, the investigation of Top Authentication Failed users could follow a drill-down of pick user and time range -> Top Destination IP, Ports for specific user and time range -> pick destination IP and port -> Query all raw messages.

FortiSIEM also uses rules for real-time alerting - a real-time event correlation engine analyzes all data and triggers alerts based on these rules. FortiSIEM ships with 500+ broad rules that cover a broad range of inter-related performance, availability, change and security scenarios. Rules can vary from simple text search and threshold conditions, to comprehensive logic supporting full Boolean operators and nested sub-patterns referencing multiple elements including thresholds and defined services. Thresholds can be static or dynamically derived from profiled network, system resource and user activity. You can add new rules, and customize existing ones, as described in Creating Rules using GUI.

Business Services

A business service lets you view FortiSIEM metrics and prioritize alerts from a business service perspective. A business service is defined within FortiSIEM as a smart container of relevant devices and applications serving a business purpose. Once defined, all monitoring and analysis can be presented from a business service perspective. It is possible to track service level metrics, efficiently respond to incidents on a prioritized basis, record business impact, and provide business intelligence on IT best practices, compliance reporting, and IT service improvement. What is also novel about FortiSIEM is how easily a business service can be defined and maintained. Because FortiSIEM automatically discovers the applications running on the servers as well as the network connectivity and the traffic flow, you can simply choose the applications and respective servers and be intelligently guided to choose the rest of components of the business service. This business service discovery and definition capability in FortiSIEM completely automates a process that would normally take many people and considerable effort to complete and maintain.

Architecture

The FortiSIEM virtual appliance solution operates as a turnkey, guest host application running within the most popular hypervisors with the option of using NFS or local storage. The implementation process is flexible and can be accomplished in phases to support a variety of distributed and hybrid-cloud implementations The FortiSIEM

virtual appliance is placed on a network where it can obtain operational data, as well as establish sessions with the infrastructure. Remote sites can use the FortiSIEM Collector client to locally discover, collect, compress and securely transmit of operation data back to the FortiSIEM virtual appliance. FortiSIEM' scale-out architecture allows for virtual appliance clustering to increase processing capacity and availability. Additional virtual appliances can be added on-the-fly with nominal configuration, which will automatically distribute workload across cluster members to extend event analysis throughput and to reduce query response time.

Supervisors, Workers, Collectors, and Organizations

An FortiSIEM deployment can be configured using either a single virtual appliance, or with multiple virtual appliances that play different roles within the deployment. The **Supervisor** virtual appliance is the primary component in both standalone and cluster deployments, and all deployments begin with the set up and configuration of the Supervisor. As described in Supervisor and Worker Cluster Deployment for Enterprises, there may be situations in which the single appliance cannot monitor all the data and devices in your infrastructure, and so you can deploy **Worker** virtual appliances to take up the extra load. Finally, you may encounter situations in which you need to deploy **Collectors** for the purpose of gathering data that will be processed by Supervisors and Workers. As described in Supervisor with Collectors Deployment for Enterprises and Supervisor and Worker Cluster Deployment for Multi-Tenancy, these are most likely situations where you need to monitor IT infrastructure for different sites, as in the case of a large or distributed enterprise, or for different organizations, as in the case of multi-tenant installations for Managed Service providers (MSPs). For these situations each **Organization** is defined separately within FortiSIEM, so you can tailor your monitoring, analytics, and reports to meet the specific needs of that organization.

Deployment Options

FortiSIEM architecture of workers, collectors, and supervisors offers a number deployment options for enterprises at any level of scale, as well as deployment options for managed service providers who need Service Provider solutions. Topics in this section describe these deployment options in detail, including use cases for each deployment type as well as node and server configurations for each deployment type.

- Enterprise Deployment Options
- Multi-Tenant Deployment Options for Managed Service Providers or Multiple Organizations

Enterprise Deployment Options

For FortiSIEM, an Enterprise deployment is one in which there is a single organization for which data is gathered and analyzed, and the virtual appliances are located entirely on-premises for that organization.

- Standalone Supervisor Deployment for Enterprises
- Supervisor and Worker Cluster Deployment for Enterprises
- Supervisor with Collectors Deployment for Enterprises
- Matrix of Enterprise Deployment Configuration Options

Standalone Supervisor Deployment for Enterprises

This is the simplest possible deployment option, in which a single Supervisor handles all the work of monitoring, processing, and analyzing data. You can configure the Supervisor to use local or NFS storage, depending on your event data storage requirements, as described in Using NFS Storage with FortiSIEM:

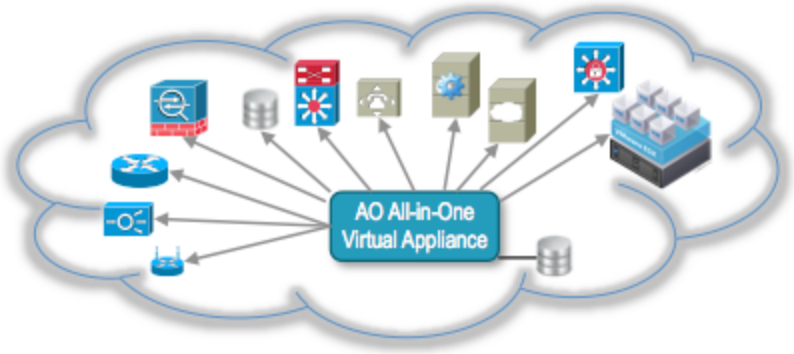

Supervisor and Worker Cluster Deployment for Enterprises

As the number of monitored devices, or the analyzed event rate, grows, one Supervisor may not be able to handle the load. In that case, you can deploy a cluster of Supervisor and Worker virtual appliances that share data over NFS. In a cluster deployment, the Supervisor and Worker nodes have specific functions:

- Discovery always runs on the Supervisor node.
- Logs can be sent to either the Supervisor or Worker nodes, and parsing occurs on the node where the event is received.
- Performance monitoring jobs are distributed by the Supervisor node across all Supervisor and Worker nodes following a load distribution algorithm.
- Users connect to the Supervisor via the FortiSIEM interface, and the Supervisor node runs the Application server, PostgreSQL (containing CMDB) and SVN database.
- Adhoc user queries, preset continuously running reports, and rules are handled by the cluster in a collaborative manner.
- Worker nodes are stateless, and can be seamlessly added or removed from the cluster as needed as the number of monitored devices or the rate of events grows, or if better query performance is required.

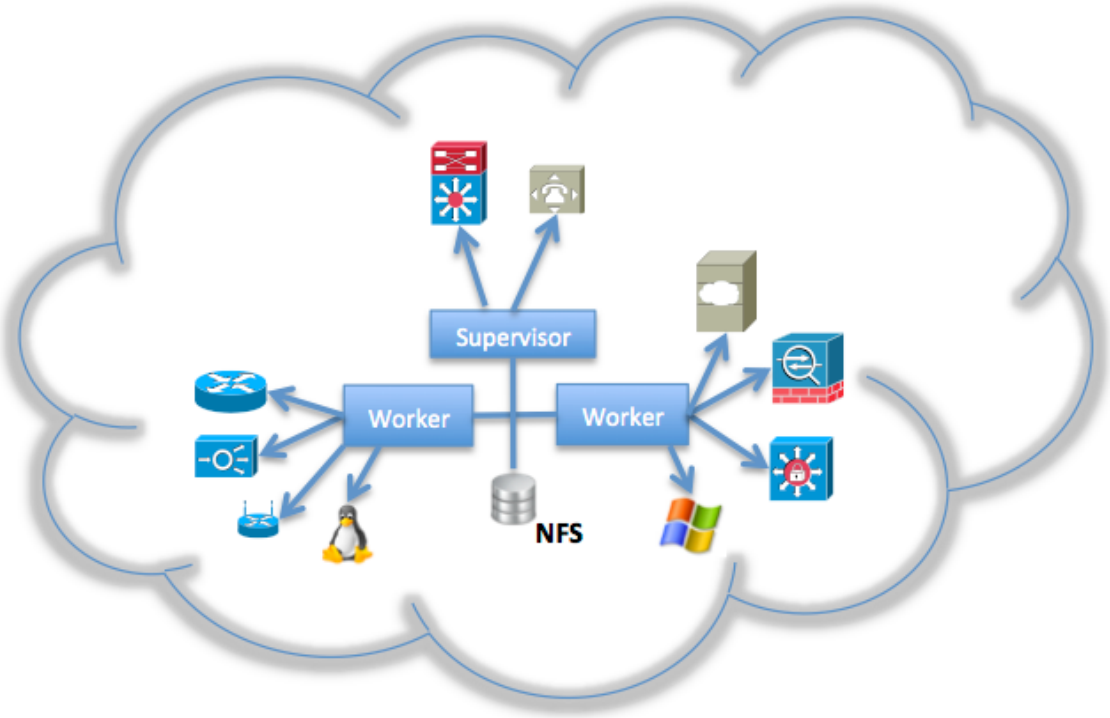

Supervisor with Collectors Deployment for Enterprises

There are two cases where a single Supervisor may not be enough for your deployment:

- There are monitored devices behind a firewall that will not allow monitoring protocols like Windows Management Instrumention (WMI) to be used from the Supervisor
- The Supervisor can only reach the monitored devices through a high latency network like a Wide Area Network (WAN), in which case monitoring like protocols like Simple Network Management Protocol (SNMP) or WMI do not work well

In these cases, you can deploy **Collectors** to monitor the devices, and they will communicate to the Supervisor over HTTP(S). The Collectors communicate with the devices, collect and parse events and logs, compress them, and then send them to the Supervisor for monitoring and analysis. Collectors also can buffer the events, in case transmission to the Supervisor is interrupted. As shown in the diagrams, you can use Collectors in a deployment with a single Supervisor, or in a deployment that also includes Workers.

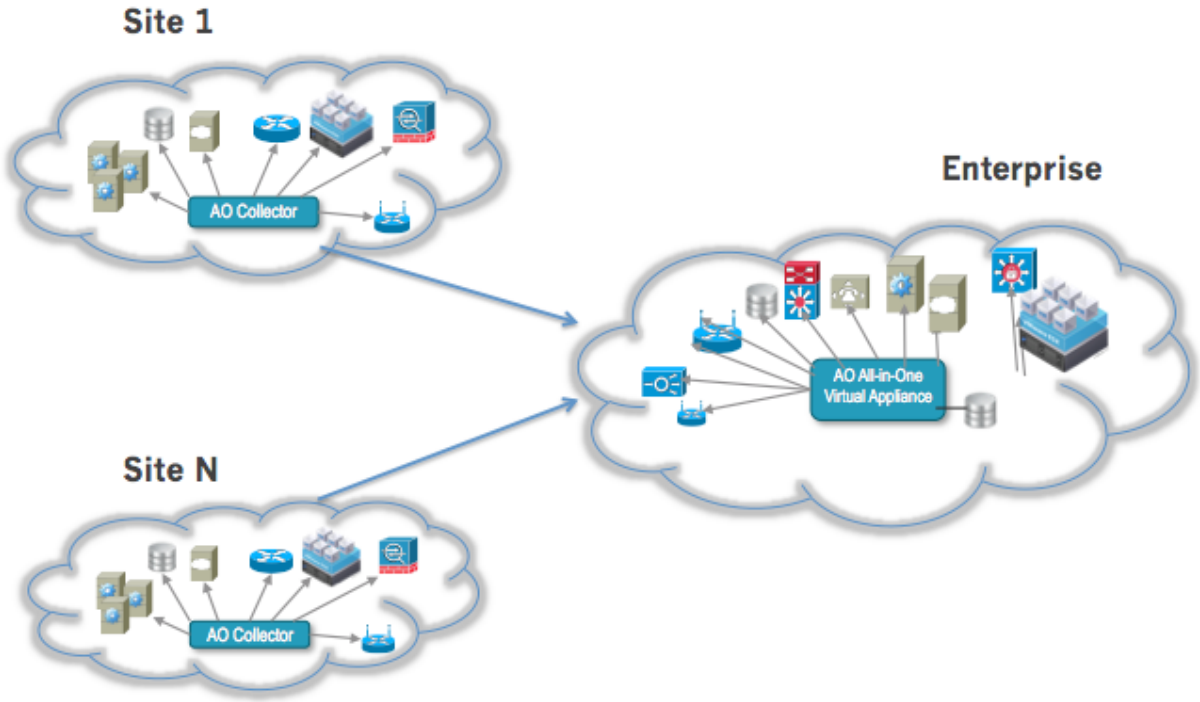

FortiSIEM deployment with a single Supervisor and Collectors

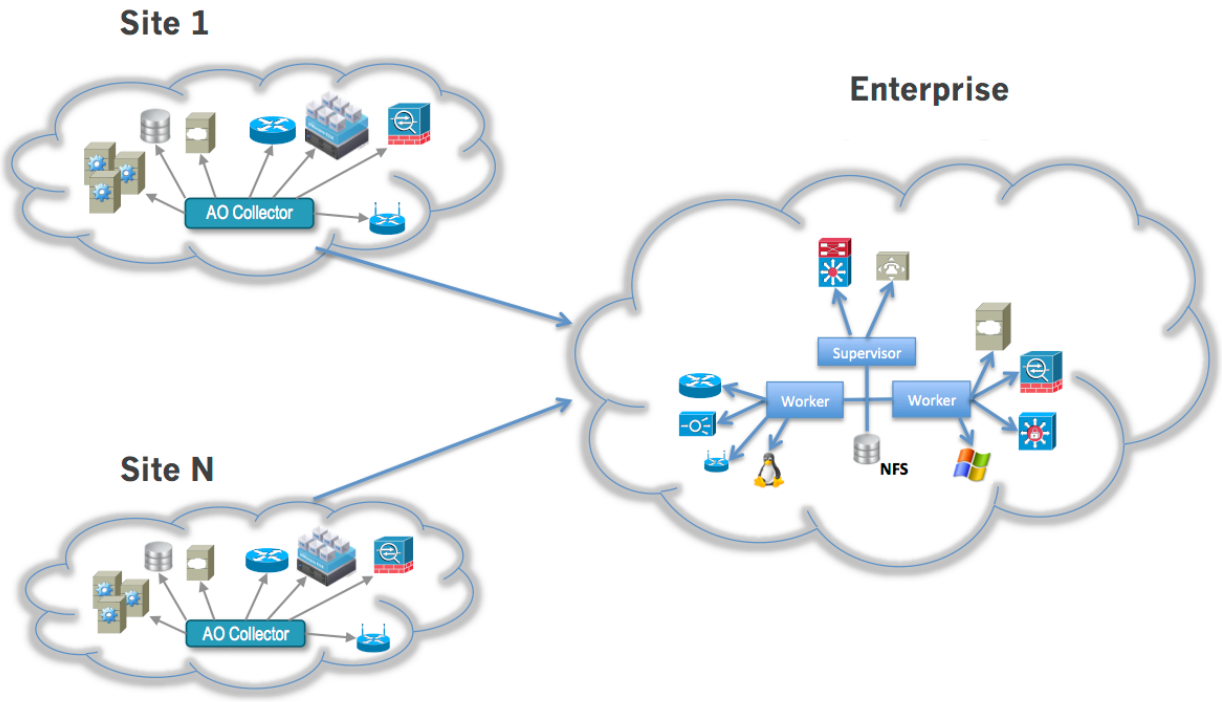

FortiSIEM deployment using a Single Supervisor + 2 Workers + 2 Collectors.

Matrix of Enterprise Deployment Configuration Options

This matrix shows the components required for each enterprise deployment option.

Deployment Option	Supervisor Node	Worker Node	Collector Node	NFS Server	Report Server	Visual Analytics Server	Description
Single Supervisor Node	x						This is the most basic single site enterprise deployment.
Supervisor Node with Collectors	x		x				This is also an enterprise deployment covering multiple sites. Data collection is simplified by deploying a collector for the satellite sites.
Enterprise Cluster	x	x		x			This is the scalable enterprise deployment. An NFS Server is required in the data sharing architecture between Supervisor and Worker nodes.
Enterprise Cluster with Collectors	x	x	x	x			This deployment adds collectors to the mix and is the most comprehensive enterprise deployment.
Supervisor Node with Tableau Visual Analytics	x				x	x	This is the most basic single node enterprise deployment, with added capability for Visual Analytics with Tableau
Supervisor Node with Collectors and Tableau Visual Analytics	x		x		x	x	This is also an enterprise deployment covering multiple sites with added capability for Visual Analytics with Tableau. Data collection is simplified by deploying a collector for the satellite sites.

Deployment Option	Supervisor Node	Worker Node	Collector Node	NFS Server	Report Server	Visual Analytics Server	Description
Enterprise Cluster with Tableau Visual Analytics	x	x		x	x	x	This is the scalable enterprise deployment with added capability for with added capability for Visual Analytics with Tableau. An NFS Server is required in the data sharing architecture between Supervisor and Worker nodes.
Enterprise Cluster with Collectors and Tableau Visual Analytics	x	x	x	x	x	x	This deployment adds collectors to the mix and is the most comprehensive enterprise deployment, with added capability for Visual Analytics with Tableau.

Multi-Tenant Deployment Options for Managed Service Providers or Multiple Organizations

While a common use case for FortiSIEM is the monitoring of IT infrastructure for a single enterprise, Managed Service Providers (MSPs) and large enterprises with multiple organizations can also use FortiSIEM to monitor IT infrastructure at the customer or organization level, either by splitting IP addresses to correspond to the customer or organization, or by deploying Collectors for each customer or organization and managing the monitoring and analysis of their data from a centralized Supervisor.

- Standalone Supervisor Deployment for Multi-Tenancy
- Supervisor and Worker Cluster Deployment for Multi-Tenancy
- Matrix of Multi-Tenancy Deployment Configuration Options

Standalone Supervisor Deployment for Multi-Tenancy

FortiSIEM allows users to create organizations, to and manage the entire IT infrastructure monitoring life cycle from data collection, storage, analytics and alerting for an organization that organization as a separate entity from other organizations. There are several use cases for this this multi-tenant model.

- Hosting service providers that host multiple customers in their own data center
- Managed service providers that manage a customer's data centers from their own data center
- Large enterprises that want to manage separate parts of the organization as individual customers

The simplest multi-tenancy deployment involves a single Supervisor, with organizations defined through the splitting of IP address ranges. For example:

- 10.1.1.0/24 = Customer 1
- 10.1.2.0/24 = Customer 2

During the discovery process, FortiSIEM will tag a device with the right customer ID based on the IP address definition.

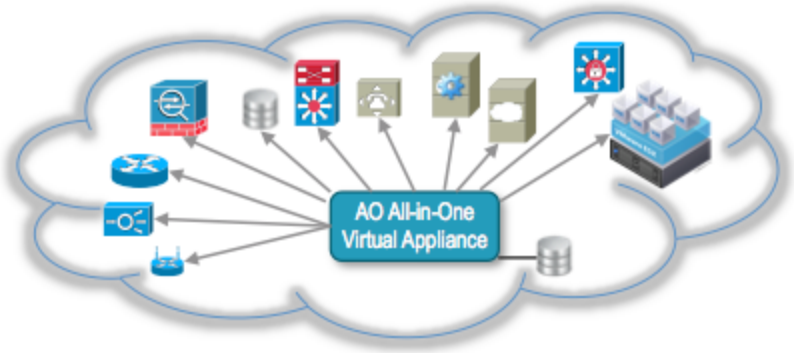

Supervisor and Worker Cluster Deployment for Multi-Tenancy

As the number of monitored devices, or the analyzed event rate, grows, one Supervisor may not be able to handle the load. In that case, you can deploy a cluster of Supervisor and Worker virtual appliances that share data over NFS. In a cluster deployment, the Supervisor and Worker nodes have specific functions:

- Discovery always runs on the Supervisor node
- Logs can be sent to either the Supervisor or Worker nodes, and parsing occurs on the node where the event is received
- Performance monitoring jobs are distributed by the Supervisor node across all Supervisor and Worker nodes following a load distribution algorithm.
- Users connect to the Supervisor via the FortiSIEM interface, and the Supervisor node runs the Application server, PostgreSQL (containing CMDB) and SVN database
- Adhoc user queries, preset continuously running reports, and rules are handled by the cluster in a collaborative manner
- Worker nodes are stateless, and can be seamlessly added or removed from the cluster as needed as the number of monitored devices or the rate of events grows, or if better query performance is required

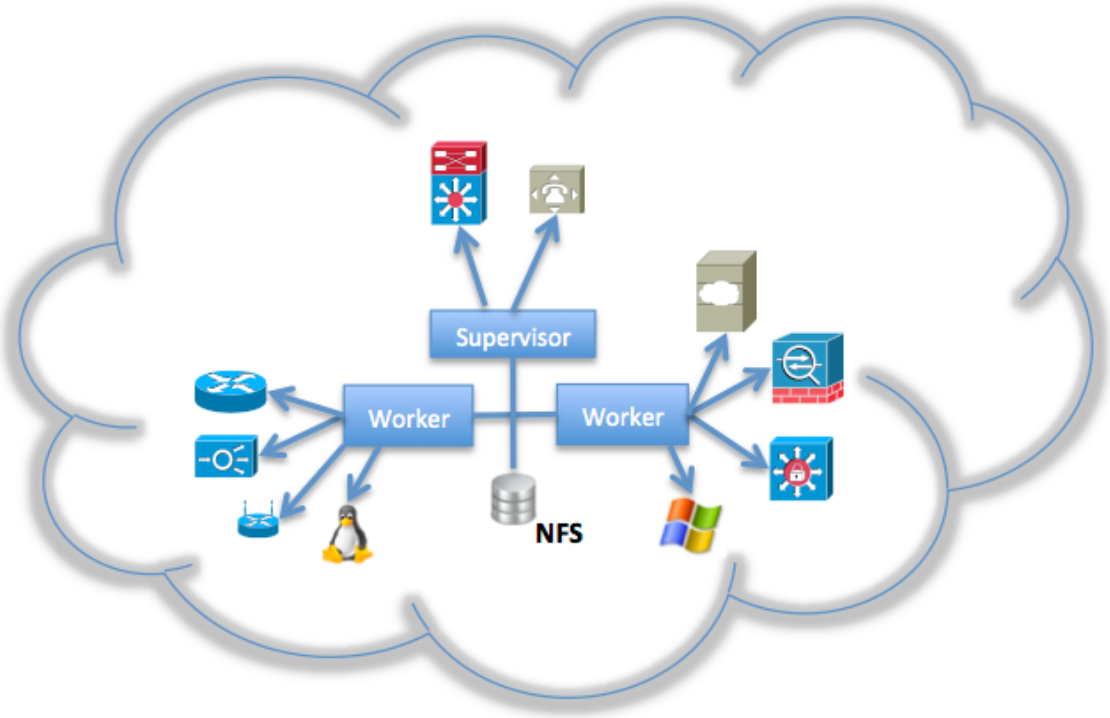

In these deployments, you can define organizations by splitting IP address ranges. For example:

- 10.1.1.0/24 = Customer 1
- 10.1.2.0/24 = Customer 2

During the discovery process, the FortiSIEM Supervisor node will tag a device with the correct customer ID based on the IP address definition.

Matrix of Multi-Tenancy Deployment Configuration Options

This matrix shows the components required for the each multi-tenancy deployment option.

Deployment Option	Supervisor Node	Worker Node	Collector Node	NFS Server	Report Server	Visual Analytics Server	Description
Single Multi-Tenant Supervisor Node	x						This is the most basic single site multi-tenant deployment, primarily suitable for hosting providers. Organizations are created by splitting up the IP address space.
Multi-Tenant Supervisor Node Collectors with	x		x				This is a service provider deployment covering multiple sites. Data collection is simplified by deploying a collector for the satellite sites. You can add organizations by assigning a collector to an organization, or by splitting up the IP address space.
Multi-Tenant Cluster	x	x		x			This is a scalable service provider deployment suitable for deployments with large compute and storage needs. An NFS Server is required in the data sharing architecture between Supervisor and Worker nodes. Organizations are created by splitting up the IP address space.

Deployment Option	Supervisor Node	Worker Node	Collector Node	NFS Server	Report Server	Visual Analytics Server	Description
Multi-Tenant Cluster with Collectors	x	x	x	x			This deployment adds collectors to the configuration and is the most comprehensive service provider deployment. You can add organizations by assigning a collector to an organization, or by splitting up the IP address space.
Multi-Tenant Supervisor Node with Tableau Visual Analytics	x				x	x	This is the most basic single site multi-tenant deployment, with added capability for Visual Analytics with Tableau.
Multi-Tenant Supervisor Node with Collectors and Tableau Visual Analytics	x		x		x	x	This is a service provider deployment covering multiple sites, with added capability for Visual Analytics with Tableau. Data collection is simplified by deploying a collector for the satellite sites.
Multi-Tenant Cluster with Tableau Visual Analytics	x	x		x	x	x	This is a scalable service provider deployment ,with added capability for Visual Analytics with Tableau. An NFS Server is required in the data sharing architecture between Supervisor and Worker nodes.

Deployment Option	Supervisor Node	Worker Node	Collector Node	NFS Server	Report Server	Visual Analytics Server	Description
Multi-Tenant Cluster with Collectors and Tableau Visual Analytics	x	x	x	x	x	x	This deployment adds collectors to the configuration and is the most comprehensive service provider deployment, with added capability for Visual Analytics with Tableau.

Export-Restrictions

FortiSIEM Export Control Classification Number is D5002.

Our Encryption Registration Number is available upon request.

Our product can not be exported, distributed, sold or used in: Cuba, Iran, Sudan, Syria and North Korea.

Configuring FortiSIEM

This chapter describes the following:

Initial System Configuration

Before you can initiate discovery and monitoring of your IT infrastructure, you will need to configure several general settings, add users, and add organizations for Service Provider deployments..

- Setting Up the Email Gateway
- Setting Up Routing Information for Reports and Incident Notifications
- Setting Up User Roles
- Adding Users for Enterprise Deployments
- Managing Organizations for Multi-Tenant Deployments
- Adding Users to Multi-Tenant Deployments

Setting Up the Email Gateway

Before you can set up notifications, you have to set up the email gateway that your system will use for all alerts and system notifications.

1. Log into your Supervisor node.
2. Go to **Admin > General Settings > Email Settings**.
3. Enter the **Email Gateway Server**.
4. Enter any additional account or connection information.
5. Click **Save**.

Setting Up Routing Information for Reports and Incident Notifications

Topics in this section describe how to set up email addresses to send alerts to when a scheduled report runs, and distribution information for notifications associated with incidents. You can also automate the sending of tickets to a Remedy system when an incident occurs. These are all general settings, in that you don't need to have any rules or reports defined before you configure them. For information on configuring specific notification policies for rules and incidents, see Incident Notifications.

- Setting Up Email Alert Routing for Scheduled Reports
- Setting Up SNMP Traps for Incident Notifications
- Setting Up XML Message Routing for Incident Notifications
- Setting Up Routing for Remedy Tickets

Related Links

- Scheduling Reports
- Incident Notifications

Setting Up Email Alert Routing for Scheduled Reports

You can schedule reports to run and send email notifications to specific individuals. This setting is for default email notifications that will be sent when any scheduled report completes.

1. Log into your Supervisor node.
2. Go to **Admin > General Settings > Analytics**.
3. Click **+**.
   If you haven't configured your email gateway yet, you will see an error message.
4. Select **SMS** or **Email** for the delivery method.
5. Enter the email address or SMS number.
6. Click **OK**.
7. Click **Save All** when you are done.

Sending Alerts to the Console

Select **Send an alert to console** if you also want to send alerts to the console. Alerts are always displayed in the Incidents tab, while the alerts sent to the console are immediately displayed but without any grouping by rule name, incident source, incident target, or other detail information.

Empty Reports

Sometimes a report may be empty because there are no matching events. If you don't want to send empty reports to users, select **Do not send scheduled emails if report is empty**. If you are running a multi-tenant deployment, and you select this option while in the Super/Global view, this will apply only to Super/Global reports. If you want to suppress delivery of empty reports to individual organizations, you will have to configure this option in the organizational view.

Related Links

- Setting Up the Email Gateway
- Scheduling Reports

Setting Up SNMP Traps for Incident Notifications

You can define SNMP traps that will be notified when an event triggers an incident.

1. Log in to your Supervisor node.
2. Go to **Admin > General Settings > Analytics**.
3. Enter the **SNMP Trap IP Address**.
4. Enter the **SNMP Community String** that will authorize sending the trap to the SNMP trap IP address.
5. Select the **SNMP Trap Type**.
6. Select a **Protocol**.
7. Click **Test SNMP** to check the connection.
8. Click **Save All**.

Related Links

- Incident Notifications

Setting Up XML Message Routing for Incident Notifications

You can configure FortiSIEM to send an XML message over HTTP(s) when an a incident is triggered by a rule.

1. Log in to your Supervisor.
2. Go to **Admin > General Settings > Analytics**.
3. For **HTTP(S) Server URL**, enter the URL of the remote host where the message should be sent.
4. Enter the **Username** and **Password** to use when logging in to the remote host, and then **Reconfirm** the password.
5. Click **Test HTTP** to check the connection.
6. Click **Save All**.

Setting Up Routing for Remedy Tickets

You can set up Remedy to accept notifications from FortiSIEM and generate tickets from those notifications. These instructions explain how to set up the routing to your Remedy server.

1. Log in to your Supervisor node.
2. Go to **Admin > General Settings > Analytics**.
3. For **WSDL**, enter the URL of the Remedy Server.
4. Enter the **Username** and **Password** associated with your Remedy server, and then **Reconfirm** the password.
5. Click **Test Remedy** to test the connection.
6. Click **Save All**.

Setting Up User Roles

FortiSIEM has a wide operational scope - it provides performance, availability, and environmental alerts, as well as change and security monitoring for network devices, servers and applications. It is difficult for one admin to monitor across the entire spectrum of available information. In addition, devices may be in widely distributed geographical and administratively disjointed locations. Role-based access control provides a way to partition the FortiSIEM administrative reponsibilities across multiple admins.

A role defines two aspects of a user's interaction with the FortiSIEM platform:

- Which **user interface elements** a user can see and the ability to use the associated Read/Write/Execute permissions. As an example, the built-in **Executive** role can see only the dashboard, while the **Server Admin** role cannot see network devices. Role permissions can be defined to the attribute level in which, for example, a **Tier1 Network Admin** role can see network devices but not their configurations.
- What **data** can the user see. For example, consider a **Windows Admin** role and a **Unix Admin** role. They both can run the same reports, but the Windows admins sees only logs from Windows devices. This definition can also be fine-grained, for example one Windows admin sub-role can be defined to see Windows performance metrics, while another Windows admin sub-role can see Windows authentication logs.

Topics in this section explain how to use the Default roles that come with FortiSIEM, and how to define new ones.

Default Roles

To perform any action with FortiSIEM, a user must be assigned a role with the required permissions. The roles listed in this table are default roles. You can create custom roles and permissions by following the instructions in the topic Creating Custom User Roles.

Role	Permissions
Full Admin	Full access to the GUI and full access to the data. Only this role can define roles, create users and map users to roles.
Network Admin	Full access to the network device portion of the GUI and full access to logs from network devices
System Admin	Full access to the Server/Workstation/Storage part of the GUI and full access to logs from those devices
Server Admin	Full access to the Server part of the GUI and full access to logs from those devices
Windows Server Admin	Full access to the Windows Server part of the GUI and full access to logs from those devices
Unix Server Admin	Full access to the Unix Server part of the GUI and full access to logs from those devices
Security Admin	Full access to Security aspects of all devices
Storage Admin	Full access to the Storage device part of the GUI and full access to logs from those devices
DB Admin	Full access to the database servers part of the GUI and full access to logs from those devices
Helpdesk	Access to the Admin, CMDB, and Dashboard tabs, with view and run permissions for the the Analytics and Incidents tabs
Read Only Admin	View access to all tabs and permission to run reports
Executive	View access to the Business Service dashboard and personalized My Dashboard tabs, but reports can be populated by logs from any device

Creating Custom User Roles

1. Log in to your Supervisor node.
2. Go to **Admin > Role Management.**
3. Click **New**.
4. Enter a **Role Name** and **Role Description**.
5. Enter the **Data Conditions** for this role.
   This restricts access to the event/log data that is available to the user, and will be appended to any query that is submitted by users with this role. This applies to both Real-Time and Historical searches, as well as Report and Dashboard information.
6. Enter the **CMDB Report Conditions** for this role.
   This restricts access to the reports for devices, users, and monitors that are available to the user with this role.
7. Select the **UI Access Conditions** for this role.
8. This defines the user interface elements that can be accessed by users with this role. By default, child nodes in the tree inherit the permissions of their immediate parent, however you can override those default permissions by explicitly editing the permission of the child node. Options for these settings are:

Setting	Description
Full	**No access restrictions**
Edit	**The role can make changes to the UI element**
Run	**The role can execute processes for the UI element**
View	**The role can only view the UI element**
Hide	**The UI element is hidden from the role**

Hiding Network Segments

If a **Network Segment** is marked as hidden for a user role, then users with that role will not be able to see any of the devices whose IP addresses fall within that network segment, even if the CMDB folder(s) containing those devices have not been hidden.

Explicit v. Effective Permissions

When a permission icon is shown within a grey box, that means that the permission was explicitly set. If the icon is shown without a border, then it represents a node's effective permission. If a permission has not been set for a node, then its effective permission is that of its nearest parent in the tree.

Adding Users for Enterprise Deployments

Adding users to enterprise deployments involves first deciding if you are going to use external authentication, or local authentication credentials defined within each user profile. You can then add users on an individual basis, or, if you are using LDAP authentication, you can discover users within Active Directory over LDAP. For mutt-tenant deployments you can add individual users to an organization as described in these topics, but if you need to add users who have a role in more than one organization (Global users), see the topics under Adding Users to Multi-Tenant Deployments.

- Setting Up External Authentication
- Adding a Single User
- Adding Users from Active Directory via LDAP
- Adding Users from Okta
- Adding 2-factor Authentication via Duo Security

Setting Up External Authentication

You have three options for setting up external authentication for your FortiSIEM deployment. The first option, LDAP, is discussed in detail in Adding Users from Active Directory via LDAP. The other options, RADIUS and Okta, follow the same authentication set up process.

Multiple Authentication Profiles

If more than one authentication profile is associated with a user, then the servers will be contacted one-by-one until a connection to one of them is successful. Once a server has been contacted, if the authentication fails, the process ends, and the user is notified that the authentication failed.

1. Log into your Supervisor node.
2. Go to **Admin > General Settings > External Authentication**.
3. Click **Add**.
4. If you are setting up authentication for an organization within a multi-tenant deployment, select the **Organization**.
5. Select the **Protocol**.
6. Complete the protocol settings.

Protocol	User-Defined Settings
LDAP	**Access IP** **Select Set DN Pattern to open a text field in which you can enter the DN pattern if you want to override the discovered pattern, or you want to add a specific LDAP user.** **See Adding Users from Active Directory via LDAP for more information about configuration settings for LDAP.**
RADIUS	**Access IP** **Shared Secret** **Select CHAP if you are using encrypted authentication to your RADIUS server**
Okta	**Certificate** **See Configuring Okta Authentication for more information.**

7. Click **Test**, and then enter credentials associated with the protocol you selected to make sure users can authenticate to your deployment.
   You can now associate users to this authentication profile as described in Adding a Single User.

Configuring Okta Authentication

To use Okta authentication for your FortiSIEM deployment, you must set up a SAML 2.0 Application in Okta, and then use the certificate associated with that application when you configure external authentication

1. Log into Okta.
2. In the **Applications** tab, create a new application using **Template SAML 2.0 App**.
3. Under **General Settings**, configure these settings:

Post Back URL	Post Back URL
Destination	**https://<FortiSIEMIP>/phoenix/okta**
Recipient	**FortiSIEM**
Audience Restriction	**Super**
authnContextClassRef	**PasswordProtectedTransport**
Request	**Uncompressed**

4. Click **Save**.
5. In the **Sign On** tab, click **View Setup Instructions**.
6. Click **Download Certificate**.
7. Follow the instructions in Setting Up External Authentication and enter the downloaded certificate for Okta authentication.

Adding a Single User

1. Log in to your Supervisor node.
2. Go to **CMDB > Users**.
3. Click **New**.
4. Complete the **User Name** and user profile information.
5. For **System Administrator**, select **Yes**.
6. Select a **Default Role** for the user.
   See the topic Default Roles for a list of default roles and permission. You can also create new roles as described in Creating Custom User Roles, which will be available in this menu after you create them.
7. For **System Account Enabled**, select **Yes**.
8. For **Session Timeout**, enter the number of minutes after which an inactive user will be logged out.
9. For **User Lockout**, enter the number of minutes the user will be unable to log into the system after three successive authentication failures.
10. For **System Password Reset**, enter the number of days after which a user's current password for logging in to the system will automatically expire.
    If left blank, the user's password will never expire.
11. For **Password**, select **Local** or **External**.
    If you select **Local**, enter and then reconfirm the user password. See Setting Up External Authentication for more information about using external authentication.
    Multiple Authentication Profiles: If more than one authentication profile is associated with a user, then the servers will be contacted one-by-one until a connection to one of them is successful. Once a server has been contacted, if the authentication fails, the process ends, and the user is notified that the authentication failed.
12. Click **Save**.

Related Links

- Default Roles
- Creating Custom User Roles

Adding Users from Active Directory via LDAP

If you want to add users to your FortiSIEM deployment from an Active Directory server over LDAP, you must first add the login credentials for your server and associate them to an IP range, and then run the discovery process on the Active Directory server. If the server is discovered successfully, then all the users in that directory will be added to your deployment. You then need to set up an authentication profile, which will become an option you can associate with users as described in Adding a Single User.

- Create Login Credentials and Associate with an IP Address
- Discover the Active Directory Server and Users

Page Size Limit for FortiSIEM Versions Prior to 4.3.1

There is a page size limit for each LDAP search result in the Active Directory server, which is often set to 1000. If any OU has more than 1000 users and the default limit is not increased, then all the users may not be discovered in FortiSIEM - this issue has been addressed in 4.3.1 by using the paged control LDAP search API. For information on how to change this limit, see this Microsoft KB article.

Create Login Credentials and Associate with an IP Address

1. Log in to your Supervisor node.
2. Go to **Admin > Setup Wizard > Credentials**.
3. Enter a **Name**.
4. For **Device Type**, select **Microsoft Windows**.
5. Select your **Access Protocol.**
   FortiSIEM supports these LDAP protocols:

Protocol	Port
LDAP	**Non-secure version on port 389**
LDAPS	**Secure version on port 636**
LDAP Start TLS	**Secure version on port 389**

6. For **Used For**, select **Microsoft Active Directory**.
7. For **Base DN**, be sure to enter the root of the LDAP user tree.
8. Enter the **NetBIOS/Domain** for your LDAP directory.
9. Enter the **User Name** for your LDAP directory.
   For user discovery from OpenLDAP, specify the full DN as the user name. For Active Directory, use your server login name.
10. Enter and confirm the **Password** for your **User Name**.
11. Click **Save**.
    Your LDAP credentials will be added to the list of **Credentials**.
12. Under **Enter IP Range to Credential Associations**, click **Add**.
13. Select your LDAP credentials from the list of **Credentials**.
14. Enter the **IP range** or host name for your Active Directory server.
15. Click **OK**.
    Your LDAP credentials will appear in the list of credential/IP address associations.
16. Click **Test Connectivity** to make sure you can connect to the Active Directory server.

Discover the Active Directory Server and Users

1. Go to **Admin > Discovery**.
2. Click **Add**.
3. For **Name**, enter **Active Directory**.
4. For **Include Range**, enter the IP address or host name for your Active Directory server.
5. Leave all the default settings, but clear the **Discover Routes** option.
6. Click **OK**.
   Active Directory will be added to the list of discoverable devices.
7. Select the Active Directory device and click **Discover**.
8. After discovery completes, go to **CMDB > Users** to view the discovered users.
   You may need to click **Refresh** for the user tree hierarchy to load.

Adding Users from Okta

Create an Okta API Token

1. Log in to Okta using your Okta credentials.
2. Got to **Administration > Security > API Tokens**.
3. Click **Create Token**.
   You will use this token when you set up the Okta login credentials in the next section. Note that this token will have the same permissions as the person who generated it.

Create Login Credentials and Associate Them with an IP Address

1. Log in to your Supervisor node.
2. Go to **Admin > Setup Wizard > Credentials**.
3. Enter a **Name**.
4. For **Device Type**, select **Okta.com**.
5. For **Access Protocol**, select **Okta API**.
6. Enter the **NetBIOS/Domain** associated with your Okta account.
   For example, `FortiSIEM.okta.com`.
7. For **Pull Interval**, enter how often, in minutes, you want FortiSIEM to pull information from Okta.
8. Enter and reconfirm the **Security Token** you created.
9. Click **Save**.
   Your LDAP credentials will be added to the list of **Credentials**.
10. Under **Enter IP Range to Credential Associations**, click **Add**.
11. Select your Okta credentials from the list of **Credentials**.
12. Enter the **IP range** or host name for your Okta account.
13. Click **OK**.
    Your Okta credentials will appear in the list of credential/IP address associations.
14. Click **Test Connectivity** to make sure you can connect to the Okta server.

Discover Okta Users

If the number of users are less than 200, then Test Connectivity will discover all the users.

Okta API has some restrictions that does not allow FortiSIEM to pull more than 200 users. In this case, follow these steps:

1. Login to **Okta.**
2. Download user list CSV file (OktaPasswordHealth.csv) by visiting **Admin** > **Reports** > **Okta Password Health**.
3. Rename the CSV file to "all_user_list_%s.csv". (%s is the placeholder of token obtained in Create an Okta API Token - Step 3, e.g. 'all_user_list_00UbCrgrU9b1Uab0cHCuup-5h-6Hi9ItokVDH8nRRT.csv')
4. Login to **FortiSIEM Supervisor node**:
   a. Upload csv file `all_user_list_%s.csv` to this directory `/opt/phoenix/config/okta/`
   b. Make sure the permissions are admin and admin (Run "`chown -R admin:admin /opt/phoenix/config/okta/`")
   c. Go to **Admin** > **Setup Wizard** > **Enter IP Range to Credential Associations**. Select the Okta entry and run **Test connectivity** to import all users.

Adding 2-factor Authentication via Duo Security

Obtain keys for FortiSIEM to communicate with Duo Security

1. Sign up for a Duo Security account: signup. This will be admin account for Duo Security.
2. Log in to Duo Security Admin Panel and navigate to **Applications**
3. Click **Protect an Application.** Locate **Web SDK** in the applications.
4. Get **API Host Name**, **Integration key**, **Secret key** from the page. You will need it when you configure FortiSIEM.
5. Generate **Application key** as a long string. This is a password that Duo Security will not know. You can choose any 40 character long string or generate it as follows using python

```
import os, hashlib

print hashlib.sha1(os.urandom(32)).hexdigest()
```

Create and Manage FortiSIEM users in Duo Security

This determines how the 2-factor authentication response page will look like in FortiSIEM and how user will respond to the second factor authentication challenge

1. Log in to Duo Security as admin user
2. Choose the **Logo** which will be shown to users as they log on
3. Choose the super set of 2-factor **Authentication Methods**.
4. **Optional** - you can create the specific users that will logon via FortiSIEM. If the users are not pre-created here, then user accounts will be created automatically when they attempt 2-factor authentication for the first time.

Add 2-factor authentication option for FortiSIEM users

1. Create a 2-factor authentication profile
    a. Go to **Admin > General Settings > External Authentication**.
    b. Click **Add.**
        a. Enter **Name**
        b. Set **Organization** to be the scopre of the users who will be authenticated.
            1. For AO-VA, specify System.
            2. For AO-SP, specify System if this will be used globally. Else specify a specific organization
        c. Set **Protocol** as Duo
        d. Set **IP/Host** from API hostname from Step 4 in "Obtain keys for FortiSIEM to communicate with Duo Security"
        e. Set **Integration key**, **Secret key**from Step 4 in "Obtain keys for FortiSIEM to communicate with Duo Security"
        f. Set **Application key** from Step 5 in "Obtain keys for FortiSIEM to communicate with Duo Security"
        g. Click **Save**
2. Add the 2-factor authentication profile to an user:
    a. Go to **CMDB** > **User**.
    b. Select a specific user.

c. Check **Second Factor** checkbox.
d. Select the 2-factor authentication profile created in Step 1.
e. Click **Save**.

Login to FortiSIEM using 2-factor authentication

Before logging in to FortiSIEM with 2-factor authentication, make sure that the three steps are completed.

1. Obtain keys for FortiSIEM to communicate with Duo Security.
2. Create and Manage FortiSIEM users in Duo Security.
3. Add 2-factor authentication option for FortiSIEM users.

Follow these steps:

1. Logon to FortiSIEM normally (first factor) using the credential defined in FortiSIEM - local or external in LDAP.
2. If the 2-factor authentication is enabled, the user will now be redirected to the 2-factor step
   a. If the user is not created in Duo system (by Duo admin), a setup wizard will let you set some basic information like phone number and ask you download the Duo app.
   b. If the user already exists in FortiSIEM, then follow the authentication method and click **Log in**
3. The user will be able to log in to FortiSIEM

Managing Organizations for Multi-Tenant Deployments

Organizations can be created with or without Collectors. If you are using Collectors in a clustered deployment that includes Workers, please make sure you have followed the instructions in Configuring Worker Settings before you have registered your Collectors with the Supervisor in order to make sure your Collectors properly upload information to the Workers.

1. Log in to your Supervisor node as a Super/Global users.
2. Go to **Admin > Setup Wizard > Organization**.
3. Click **Add**.
4. Enter information for the organization.
5. If your organization uses Collectors, click **New** under **Collectors.**
6. Complete the Collector information.
   For **Guaranteed EPS**, enter the events per second from this collector that FortiSIEM will accept. See the topic Dynamic Distribution of Events per Second (EPS) across Collectors for more information. For **Start Time** and **End Time**, enter the dates for which the Collector license is valid.
7. Click **Save**.
8. For **Max Devices**, enter the maximum number of devices discovered by this collector that the system will accept.
9. Click **Save**.

Deleting Organizations

1. Log into your Supervisor node as a Super/Global user.
2. Go to **Admin > Setup Wizard > Organizations**.
3. Write down the **ID** of the organization you want to delete.
4. Go to **Admin > Collector Health**.
   Note the **IP Address** and **Collector Name** of any Collectors associated with the organization you want to delete.
5. Log out of your Supervisor node.
6. SSH into the Collector hosts for the organization as `root`.
7. Using phTools, stop the Collector processes.
8. Power down the Collector.
9. Log back into your Supervisor node as an Admin user for the organization you want to delete.
10. Go to **CMDB > Devices.**
11. Delete all devices in both the **Device View** and the **VM View**.
12. Go to **CMDB > Device View > Users**, and delete all users except for the default admin account under which you are currently logged in.
13. Go to **Admin > Setup Wizard > Synthetic Transaction Monitoring** and delete all STM tests.
14. Log out of your Supervisor node, and then log back in as the Super/Global user.
15. Go to **Admin > Collector Health.**
16. Delete the organization's Collectors.

    Issues with Deleting Collectors Because of In-Memory Processes

    You may encounter issues with deleting Collectors if there are processes in memory on the Supervisor that are related to Collector status that are updated to the CMDB. If you encounter these issues, please contact FortiSIEM Support.
17. Delete the organization.
18. Log out of your Supervisor node.
19. SSH into the Supervisor host machine as `root`.
20. In the `/data` directory, delete the `eventdb` database for that organization.

Finding the Right EventDB Database

You can tell which EventDB belongs to the organization you want to delete based on the organization ID that you wrote down in Step 3. For example, if the organization ID is 2005, you would look for /data/eventdb/CUSTOMER_2005 as the database to delete. Be careful that you don't delete the EventDB for a continuing organization.

Dynamic Distribution of Events per Second (EPS) across Collectors

In Service Provider deployments, the service provider is licensed a certain amount of EPS. The service provider distributes these EPS among the various collectors during collector setup by setting the **Guaranteed EPS**. Because an organization can have multiple collectors, the guaranteed EPS for an organization is the sum total of guaranteed EPS for all collectors belonging to that organization. This total must be no more than the total EPS licensed to the service provider. The remaining EPS (the difference between the service provider EPS and the total EPS across all collectors), if any, is allocated to the super-local organization, the service provider's core system, if that needs to be monitored. To monitor this system, FortiSIEM recommends creating a new organization to monitor the service's own network, and to install another Collector to monitor that organization.

The redistribution algorithm uses three metrics for each Collector.

Guaranteed EPS	Defined during the collector configuration process while setting up an organization, FortiSIEM ensures that the collector can always send EPS at this rate. This is a constant that never changes during the operation of the algorithm, unless you edit the Collector definition.
Incoming EPS	This is the EPS that the Collector sees. This changes continuously. You can view this metric for a Collector in **Admin > Collector Health**.
Allocated EPS	This is the EPS that is currently allocated to the Collector by the redistribution algorithm. Y ou can view this metric for a Collector in **Admin > Collector Health** .

Each Collector periodically reports Incoming EPS to the Supervisor, which then determines the Allocated EPS and pushes this control down to the collectors. Allocated EPS is set to Guaranteed EPS initially, but if for some Collector, Incoming EPS is greater than Allocated EPS, the Supervisor examines all Collectors and determines excess capacity as sum total of max (0,Allocated - Incoming) for all Collectors. If there is a Collector with excess capacity, its Allocated EPS is reduced and the excess amount is given to the Collector that needs the excess EPS. If the collector that gave up EPS, that is, Allocated EPS is less than Guaranteed EPS, subsequently needs the EPS, then EPS is taken away from the collectors with Allocated greater than Guaranteed and given back. This continuous readjustment is centrally coordinated by the Supervisor node.

How Devices are Added to Organizations

When you initiate device discovery for organizations, the way in which those devices are added to organizations depends on whether you are using Collectors in your deployment.

- For organizations with Collectors, discovery is carried out by the Collector, and the Collector assigns devices to the organization with which it is associated. If organizations have an overlapping IP range, deploying Collectors and assigning them to a specific IP range and organization will ensure that the device is added to the correct organization.
- For organizations without Collectors, discovery is carried out by the Supervisor. In this case, the **Include/Exclude IP Range** you defined when you set up the organization is used to add the device to the organization.
    - If a device matches only one defined organization IP Range, then it is assigned to that organization
    - If a device matches multiple defined IP Ranges, then it is assigned to the Super organization

You can change a device's assigned organization manually, and FortiSIEM will automatically update the Include/Exclude IP Range for that organization. This updated IP range definition will then be used in the next discovery process. However, this may create confusing IP range definitions for the organization, so you may want to re-define the organization's IP range and rediscover devices.

Adding Users to Multi-Tenant Deployments

Two kinds of admin users can be added

- users belonging to a specific organization or super-local
- users belonging to super-global

Adding specific organization users

This can be done from the specific organization admin account or from the super global account.

- Logon as an appropriate administrator - two possibilities
  - logon as admin user for that organization **or**
  - logon as super-global and then switch user to that organization.
- Follow the steps for AO-VA case described here. Note that for Active Directory based discovery, the Active Directory server has to belong to that specific organization. If the Active Directory server belongs to super-local, then the users also belong to super and would not be visible for that organization.

FortiSIEM provides a short-cut to add admin users for multiple organizations in one shot

- Logon as super-global.
- Manually create the user as described in the manual user creation mode here.
- Choose the Default role.
- Choose the permitted organizations and also override the default role for a specific organization if needed. In the example below, user1 is the Network Admin for every organization but System Admin for O-eng.

Adding super-global users

Super-global users are often need for managing multiple organizations, and can be created from the super-global account. There are two cases depending on whether organizations have collectors or not.

For the organizations-with-collector-only case, users must be created manually.

- Logon as super-global.
- Manually create the user as described in the manual user creation mode here.
- Choose the Default role.

- Choose the permitted organizations. Override the default role for each specific organization, if needed. In the example below, user1 is the Network Admin for every organization but System Admin for O-eng.

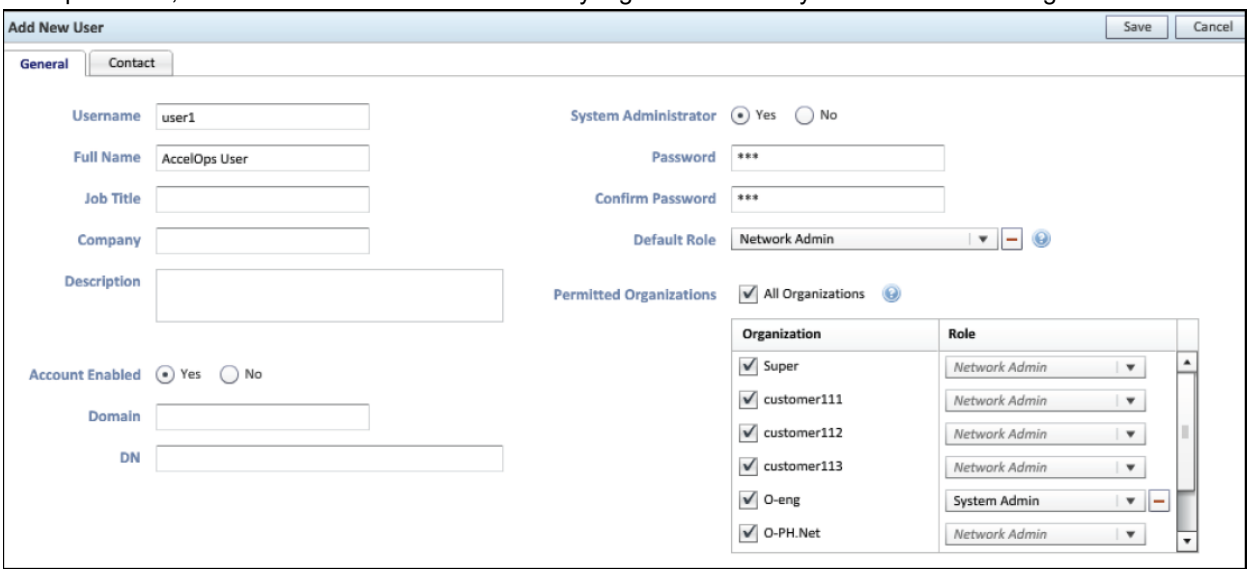

For the organizations-without-collector case, if the Active Directory Server belongs to super-local, then the discovered users would be visible from the super-global view and any of these users can be made FortiSIEM user. In this case the steps are

- Logon as super-global
- Create the user as described here - both manual and discovery-based approaches can be used
- Choose the Default role
- Choose the permitted organizations. And if needed, override the default role for specific organizations. In the example below, user1 is the Network Admin for every organization but System Admin for O-eng.

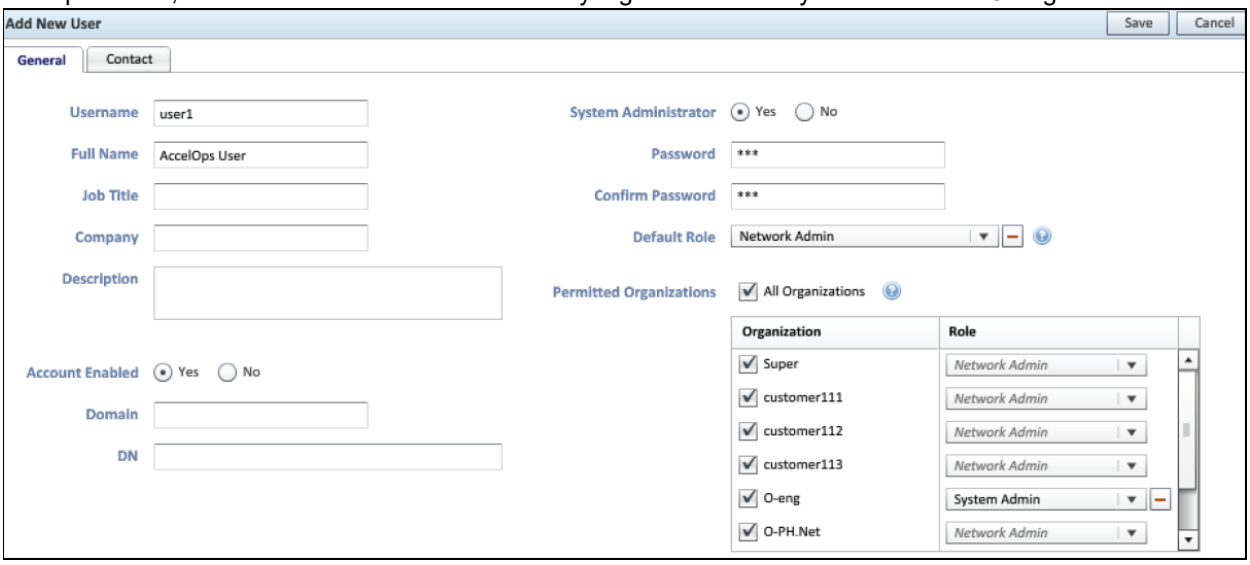

Adding Users to Organizations

Adding users to organizations for Service Provider deployments follows the same processes described in Adding Users for Enterprise Deployments, though if you want to discover users in an Active Directory server over LDAP, the Active Directory server has to belong the organization where you want to add the user.

1. Log in to your Supervisor node either as the Admin user for the organization where you want to add the user, or log in as a Super/Global user to add the user to more than one organization.
2. Create the user as described in Adding a Single User, or follow the instructions in Adding Users from Active Directory via LDAP.
3. If you have logged in as the Super/Global user, select the organizations where you want to add the user, overriding any Default Roles for the organization as necessary.

Adding Super/Global Users to Organizations with Collectors

In Service Provider deployments, you may need to create Super/Global users who have roles within multiple organizations. If your deployments include organizations with collectors, you must add add the users individually.

1. Log in to your Supervisor node as a Super/Global users.
2. Create the individual user as described in Adding a Single User, choosing the appropriate **Default Role**.
3. Select the **Permitted Organizations** the user is allowed to access, overriding any default role settings as necessary.
4. Click **Save**.

Adding Super/Global Users to Organizations without Collectors

For the organizations-without-collector case, if the Active Directory Server belongs to super-local, then the discovered users would be visible from the super-global view and any of these users can be made FortiSIEM user. In this case the steps are

- Logon as super-global
- Create the user as described here - both manual and discovery-based approaches can be used
- Choose the Default role
- Choose the permitted organizations. And if needed, override the default role for specific organizations. In the example below, user1 is the Network Admin for every organization but System Admin for O-eng.

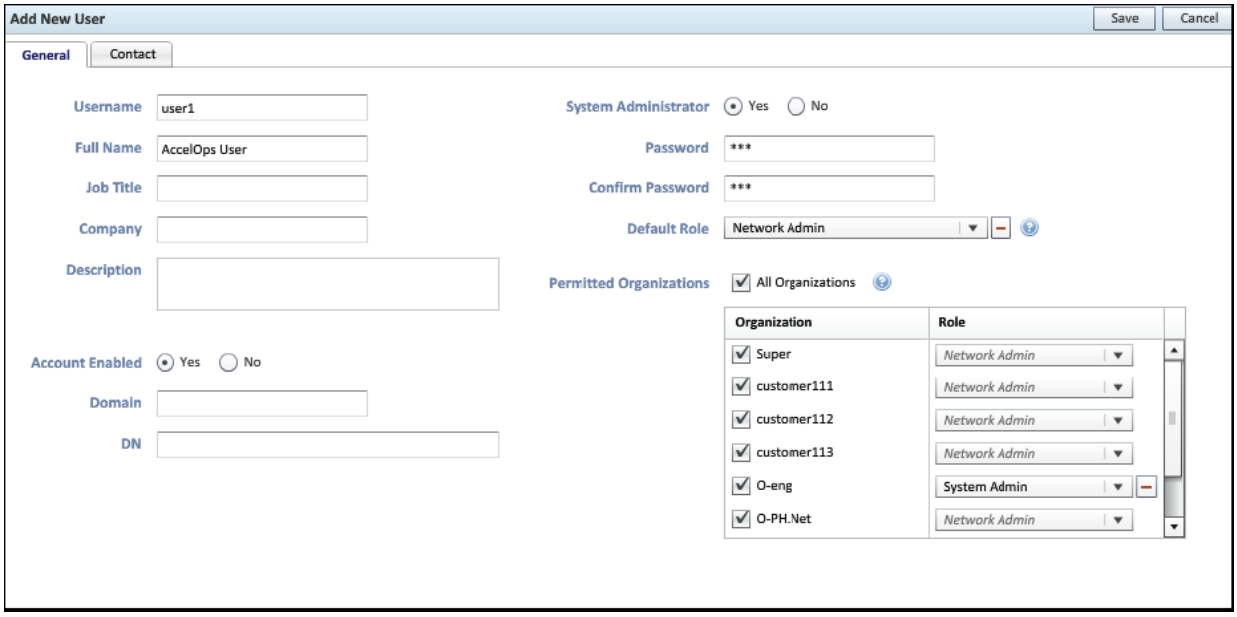

Discovering Infrastructure

FortiSIEM can automatically discover the devices, applications, and users in your IT infrastructure and begin monitoring them. You initiate device discovery by providing the credentials that are needed to access the infrastructure component, and from there FortiSIEM is able to discover information about your component such as the host name, operating system, hardware information such as CPU and memory, software information such as running processes and services, and configuration information. Once discovered, FortiSIEM will also begin monitoring your component on an ongoing basis.

Though FortiSIEM is able to automatically manage device discovery, the pulling of event information such as logs and IPS events from your device, and establishing what aspects of your device functionality it can monitor, you can also manually configure the way FortiSIEM interacts with your infrastructure by creating custom event pulling methods and monitoring profiles for your devices.

Check Device Configuration Before Initiating Discovery

Before you begin the process of device discovery, you should make sure your devices are properly configured for discovery and monitoring by FortiSIEM. Refer to *FortiSIEM External Systems Configuration Guide* for more information.

WMI or SNMP for Discovery of Windows Devices

Windows servers can be discovered by either SNMP or WMI. or both. SNMP provides installed software information, while WMI provides all application metrics, detailed system metrics, and logs.

Ping-Only Discovery for Basic Up/Down Status

If you only need to monitor a device for up/down status, then select the **Ping Discovery Only** option when setting the **Range Definition** for discovering the device. The device will be listed in the CMDB and consume one device license.

- Discovery Settings
- Discovery for Multi-Tenant Deployments
- Setting up CyberArk
- Setting Access Credentials for Device Discovery
- Discovering Devices
- Discovering Amazon Web Services (AWS) Infrastructure
- Discovering Microsoft Azure Infrastructure
- Approving Newly Discovered Devices
- Inspecting Event Pulling Methods for Devices
- Inspecting Changes Since Last Discovery
- Discovery Range Definition Options
- Scheduling a Discovery
- Adding Devices to the CMDB Outside of Discovery
- Decommissioning a device

Discovery Settings

Before you initiate discovery, you should configure the Discovery Settings in your Supervisor.

1. Log in to your Supervisor node.
2. Go to **Admin > General Settings > Discovery**.
3. Configure the settings as required for your deployment.
   See Setting Device Location Information for information on how to manually enter locations for devices, or to upload a CSV file of device locations.

Setting	Description
Virtual IPs	Often a common virtual IP address will exist in multiple machines for load balancing and failover purposes. When you discover devices, you need to have these virtual IP addresses defined within your discovery settings for two reasons: • Listing the virtual IP addresses ensures that two or more devices with the same virtual IP will not be merged into one device during device discovery, so each of the load-balanced devices will maintain their separate identity in the CMDB • The virtual IP will not be used as an access IP during discovery, since the identity of the device when accessed via the virtual IP is unpredictable Click the **Edit** icon to enter a Virtual IP address, and then click **+** to add more.
Excluded Shared Device IPs	An enterprise often has servers that share credentials, for example mail servers, web proxies, and source code control servers, and a large number of users will authenticate to these servers to access their services. Providing a list of of the IP addresses for these servers allows FortiSIEM to exclude these servers from user identity and location calculations in the **Analytics > IdentityandLocation** report. For example, suppose user U logs on to server M to retrieve his mail, and server M authenticates user U via Active Directory. If server M is not excluded, the **Analytics > Identity and Location Report** will contain two entries for user U: one for the workstation that U logs into, and also one for server M. You can eliminate this behavior by adding server M to the list of Server IPs with shared credentials.
Allow Incident Firing On	With this setting you can control incident firings based on approved device status. If you select **Approved Devices Only,** then FortiSIEM will use this logic to determine if an incident is triggered: • If an incident reporting device is not approved, the incident does not trigger • If an incident reporting device is approved, then there are two possible cases: (a) at least one Source, Destination or Host IP is approved and the incident triggers, or (b) none of the Source, Destination or Host IPs are approved and the incident does not trigger If you select **Approved Devices Only**, then when the discovery process completes, you will need to approve devices, as described in Approving Newly Discovered Devices, before incidents are triggered.

Setting	Description
CMDB Device Filter	This setting allows you to limit the set of devices that the system automatically discovers from logs and netflows. After receiving a log from a device, the system automatically discovers that device, and then adds it to CMDB. For example, when a Netflow analysis detects a TCP/UDP service is running on a server, the server, along with the open ports, are added to CMDB. Sometimes you may not want to add all of these devices to CMDB, so you can create filters to exclude a specific set of devices from being added to CMDB. Each filter consists of a required **Excluded IP Range** field and an optional **Except** field. A device will not be added to CMDB if it falls in the range defined in the Excluded IP Range field. For example, if you wanted to exclude the `172.16.20.0/24` network from CMDB, you would to add a filter with `172.16.20.0-172.16.20.255` in its **Excluded IP Range** field. The **Except** field allows you to specify some exceptions in the excluded range. For example, if you wanted to exclude the `172.16.20.0/24` network without excluding the `172.16.20.0/26` network, you would add a filter with `172.16.20.0-172.16.20.255` in the **Excluded IP Range** field, and `172.16.20.192-172.16.20.255` in the **Except** field. Click **Add** to add a new CMDB Device Filter, then click **Apply**.
Application Filtering	This setting allows you to limit the set of applications/processes that the system automatically learns from discovery. You may be more interested in discovering and monitoring server processes/daemons, rather than client processes, that run on a server. To exclude client processes from being discovered and listed in the CMDB, enter these applications here. An application/process will not be added to CMDB if it matches one of the entries defined in this table. Click **Add,** then enter the **Process Name** and any **Parameters** for that process that you want to filter. M atching is exact and case-insensitive based on Process Name and Parameter. If Parameter is empty, then only Process Name is matched.

Setting Device Location Information

In the **Admin > General Settings > Discovery** screen, you can set device locations based on IP range and organization. You can do this manually for each organization or IP range, or upload a CSV file that contains location information. This information can then be applied to devices already in the CMDB, or during the discovery process, to set their location.

Manually Creating Location Information

1. Log into your Supervisor node.
2. Go to **Admin > General Settings > Discovery**.
3. Under **Location**, click **Add**.
4. For Multi-Tenant deployments, enter the **Organization** you want to associate with the IP range and devices.
5. Enter the **IP/IP Range** you want to associate with the location.
   This can be in either CIDR notation, such as `192.168.64.0/24`, or range notation, such as `192.168.64.0-192.168.64.255`.
6. Enter the **Display Name** you want to use for this location.
   For example, **San Jose Office**, **Northern California Campus**, etc.
7. Enter any additional location information that is relevant for your location.
8. Click **OK**.
9. In the **Location Definition** dialog, select **Update Manual Devices** if you want to update devices that have had their locations set manually in the CMDB.
10. Click **OK**.
    The location information will appear in the **Location** pane.
11. Select a location in the **Location** pane, and then click **Apply** to associate all devices in the CMDB with that IP/IP range to that organization and location.
    A dialog will indicate how many devices have been updated.
12. Click **OK**.
13. Go to **CMDB > Devices** and check that your device locations have been updated.

Uploading Location Information from a CSV File

Prerequisite

Before you can upload it, you must first create a CSV file with this format.

Comma-separated IP address, Range, or Subnet	Location Display Name	Update Manual Devices (False/True)	Geographic Information ("region:;country:;state:;city:;building:;floor:;latitude:;longitude:;")

Example

"10.1.1.1/24,20.1.1.1-20.1.1.10"	San JoseDatacenter USA	true	
"30.1.1.10"	Fremont Datacenter USA	true	"region:North America;country:United States;state:California; city:Fremont;building: 10;floor:4;latitude:38.1747222;longitud e:-121.2775;"

Procedure

1. Log into your Supervisor node.
2. Go to **Admin > General Settings > Discovery**.
3. Under **Location**, click **Import**.
4. Browse to your CSV file and select it.
5. Click **Upload**.

Discovery for Multi-Tenant Deployments

In Service Provider deployments with organizations, the discovery process differs depending on whether or not you are using Collectors. This is because of the way in which IP addresses are used to establish the relationship between devices and organizations.

- If you are using Collectors, IP address overlap between organizations is allowed
- If you are not using Collectors, then each organization must have a unique IP address

These two requirements determine which administrative account you will use for discovery.

- For organizations with collectors, you must initiate discovery using the administrative account associated with the organization. Every device discovered by a collector is automatically assigned to the organization that the collector belongs to.
- For organizations without collectors, you must initiate discovery using the Super/Global administrative account. Devices for all organizations are discovered at the same time, and are assigned to organizations based on the IP address assignments you set up for the organization.
  .
  - If a device matches only one organization's IP address assignment, then it is assigned to that organization
  - If a device matches multiple organization definitions, then it is assigned to the Super/Global organization. These would typically be devices that are part of the Super/Global organization's network backbone.

Related Links

- How Devices are Added to Organizations
- Managing Organizations for Multi-Tenant Deployments

Setting up CyberArk

This section specifies how FortiSIEM can be configured to fetch credentials from CyberArk.

Installing CyberArk Provider in FortiSIEM

Refer to "**Credential Provider and ASCP Implementation Guide**" for more details on CyberArk Credential Provider installation.

1. Login to FortiSIEM as root.
2. Run the rpm command to begin the installation: `rpm -i CARKaim-<version>-<build number>.x86_64.rpm`

The installation runs automatically and does not require any interactive response from the user. When the installation is complete, the following message appears: "Installation process completed successfully."

Configuring CyberArk Provider in FortiSIEM

Refer to "**Credential Provider and ASCP Implementation Guide**" for more details on CyberArk Credential Provider installation.

1. Login as root.
2. Open the Vault.ini file and specify the parameters of the Vault that will be accessed by the Provider.
3. Run **CreateCredFile** to create a credential file for the administrative user that will create the Vault environment during installation.
   `Createcredfile <filename> Password –Username <username> -Password <password>`
4. Run the **CreateEnv** utility that was copied to the bin folder during Provider installation
   `CredFilePath <CredFilePath> -VaultFilePath <VaultFilePath> -AppProviderUser <ProviderUserName> -LicensedProducts <AIM\OPM\ALL> [-AppProviderConfSafe <ConfigurationSafeName>] [-MainAppProviderConfFilePath <MainConfigurationFilePath>] [-OverrideExistingConfFile <Y\N>] [-PIMConfigurationSafe <PIMConfigurationSafeName>] [-AppProviderUserLocation <ApplicationUsersLocation>]`
5. Check the log file **/var/tmp/aim-install-logs/CreateEnv.log** to make sure that the Provider environment was created successfully
6. Start the CyberArk Application Password Provider service manually as a privileged user
7. Add **/opt/CARKaim/sdk/** in **/etc/ld.so.conf**
8. Run **ldconfig**

Configuring CyberArk for communication with FortiSIEM

Refer to the Privileged Account Security Implementation Guide for more information about adding and managing privileged accounts.

1. Login to CyberArk Password Vault Web Access (PVWA) Interface as an user allowed to managed applications (it requires Manage Users authorization).
2. Add FortiSIEM as an Application
   a. Go to **Applications** and click **Add Application.**
   b. Set **Name** to FortiSIEM
   c. In the **Description**, specify a short description of the application that will help you identify it (e.g. FortiSIEM SIEM)

   d. In the **Business owner** section, specify contact information about the application's Business owner.
   e. In the lowest section, specify the **Location** of the application in the Vault hierarchy. If a Location is not selected, the application will be added in the same Location as the user who is creating this application.
   f. Click **Add**; the application is added and is displayed in the Application Details page
3. Check **Allow extended authentication restrictions** – this enables you to specify an unlimited number of machines and Windows domain OS users for a single application
4. Specify the application's (FortiSIEM) **Authentication** details. This information enables the Credential Provider to check certain application characteristics before retrieving the application password.
   a. In the **Authentication** tab, click **Add**; a drop-down list of authentication characteristics is displayed.
   b. Specify the OS user as **"admin"** and Click **Add.**
   c. Specify the application path as "/opt/phoenix/bin". Make sure **Path is folder** and Allow internal scripts to request credentials… check boxes are checked
   d. Do not specify a hash
   e. In the **Allowed Machines** tab, click **Add** and specify the IP/host name of the FortiSIEM Supervisor, Workers and Collectors
5. Authorize FortiSIEM to retrieve accounts.
   a. Go to **Policies** > **Access Control (Safes)**
   b. For every **Safe**, Click on **Members**.
   c. Click on **Add Safe Member**
   d. Search for FortiSIEM. An entry will already exist. Select that entry.
   e. Check **Retrieve accounts.**
   f. Click **Add**

Now FortiSIEM should be ready to retrieve passwords from CyberArk via Test Connectivity and Discovery.

Setting Access Credentials for Device Discovery

Before you can discover devices, you need to provide the access protocol and credentials associated with the IP address or range where your devices are located. FortiSIEM will then use this information to access your devices, pull information from them, and begin monitoring them.

Access Protocols Required for Discovery

SNMP, VM SDK (for VMware vCenter), or WMI (for Windows devices) must be one of the access protocols for which you provide credentials in order for the devices associated with an IP address or range to be discovered. If your device does not use one of these protocols, then you must configure it to communicate with FortiSIEM as described in *FortiSIEM External Systems Configuration Guide*. As described in those topics, you may also need to set up additional configurations within your devices to send logs and other information to FortiSIEM.

Associate Credentials Only with the IP Address Where They Will be Used

Credentials should only be associated with IP addresses where they can be used. Assigning multiple credentials to IP addresses where they are not used will trigger discovery operations for each credential, and the system will wait for a timeout to occur for each credential before it moves to the next one. This will cause the discovery process to require much more processing time and processing power from the FortiSIEM system. You can, however, associate the same credential (for example, a generic SNMP access credential) to multiple IP addresses where it will be used to communicate with a device over that protocol.

Before starting the discovery process, credentials need to be defined and then associated to specific IP addresses.

CyberArk configuration

If CyberArk is going to be used for Test Connectivity and Discovery, then first follow the steps defined here.

Define Credentials

1. Log into your Supervisor node.
2. Go to **Admin > Setup Wizard > Discovery**.
3. Under **Enter Credentials**, click **Add**.
4. Enter a **Name** for the credential.
5. Select a **Device Type** to associate with the credential.
6. Select the **Access Protocol** for which you want to enter credentials.
   Note that the Device Type selection determines which Access Protocols are available.
   Change the default destination ports only if needed
7. Choose **Password Configuration** method
   1. **Manual** - means that you have to define credentials in FortiSIEM
   2. **CyberArk** - means AccelPs will fetch credentials from CyberArk
8. If you choose **Password Configuration** as **Manual,** then enter the credentials required for the Access Protocol.
9. If you choose **Password Configuration** as **CyberArk,** then choose CyberArk parameters
   1. **AppID** must be set to FortiSIEM
   2. Specify **Safe, Folder, Object**: This is the CyberArk Vault Safe, Folder, Object where the credential is defined.

3. Specify **User Name**: This is the User Name of the credential
4. Specify **Platform (Policy ID):** This is the platform related property for the credential. Specify this only if this property is also set in CyberArk. The match will be case sensitive.
5. Specify **Database**: This is a property for the database credential. Specify this only if this property is also set in CyberArk. The match will be case sensitive.
6. Check **Include Address for Query**: If checked, FortiSIEM will query the CyberArk credential by IP or host name. Specify this if CyberArk credential objects are specified by IP.

10. Click **Save**. The credentials you created will be added to the list.

Specify Device to Credential Mapping

1. Under **Enter IP Range to Credential Associations**, click **Add**.
2. Select the credential you just created from the list.
   Note that you can add multiple credentials to the same IP/host information in this step by clicking **+**.
3. Enter an **IP address**, **IP range**, or **Host Name** to associate with the credential.
   Formats for IP Information

   You can provide the IP information in several formats:
   A single IP address e.g. 192.168.1.1
   A range of IP addresses in the format <IP Address>-<IP Address>

   A IP subnet specified in CIDR notation <IP Address>/<Maskbits> e.g. 192.168.0.0/16 to specify the IP range 192.168.0.0-192.168.255.255

   A combination of the three formats separated by comma, e.g. 192.168.1.1, 192.168.1.2, 10.10.0.0/16, 10.11.0.0-10.11.255.255
   A host name
4. Click **OK**.

Test Connectivity

You need to perform a **Test Connectivity** to make sure that the credentials are correct.

1. Select the IP/credential association you just created, and click **Test Connectivity**. A ping will be performed first to make sure that the host is alive. If ping is disabled in your network, then choose **Test Connectivity without ping**.
   A dialog will show you the results of your connectivity tests. Note that the connectivity tests can take several minutes, so you may want to use the **Run in Background** option.

Discovering Devices

Prerequisites

- Make sure you have configured the Discovery Settings for your deployment
- Set up the Access Credentials for your devices so FortiSIEM can communicate with them

Procedure

After you have set up the access protocols for your devices as described in Setting Access Credentials for Device Discovery, you are ready to discover devices in your IT infrastructure.

1. Log in to your Supervisor node.
   Discovering Devices for Multi-Tenant Deployments

   If you have a Service Provider FortiSIEM deployment that uses Collectors and you and want to discover devices for a specific organization, rather than the Global organization, log into your Supervisor node as an admin user for that organization. See Discovery for Multi-Tenant Deployments for more information about how discovery works for Service Provider deployments with and without Collectors.
2. Go to **Admin > Setup Wizard > Discovery**.
3. Click **Add**.
   You can also schedule single or recurring discovery processes as described in Scheduling a Discovery.
4. In the **Range Definition** dialog, set the options for this discovery.
   See Discovery Range Definition Options for more information about the options available in this dialog.
5. Click **OK**.
   Your range definition will be added to the list.
6. Select your range definition, and then click **Discover**.
   A discovery dialog will show you the progress of your discovery. For long-running discoveries, you can use the **Run in Background** option.
7. When discovery completes, the results will be displayed in the dialog. Click **Errors** to view any errors.
   Possible Causes of Discovery Errors

   If there are errors during the discovery process, the Errors screen will inform you of their severity, impact, and potential resolution. Some possible reasons for errors include:

   - A device is not online or not reachable via ping. FortiSIEM will attempt to ping devices before initiating a full discovery to save time.
   - A device is not responding to SNMP or WMI requests, or there is a firewall blocking these requests from FortiSIEM
   - The SNMP/WMI credentials are incorrect
   - WMI may not have been set up correctly on the server. See the appropriate topic in *FortiSIEM External Systems Configuration Guide* for how to configure WMI for your device.

   Approving Newly Discovered Devices

   If you selected **Approved Devices Only** for the discovery setting **Allow Incident Firing On**, as described in Discovery Settings, then you will need to approve your newly discovered devices before incidents will be triggered for those devices. See Approving Newly Discovered Devices for more information.

Discovering Amazon Web Services (AWS) Infrastructure

You can configure FortiSIEM to communicate with your device, and then initiate discovery of the device. For more information, refer to sections 'Discovering Infrastructure' and 'Setting Access Credentials for Device Discovery' under 'Chapter: Configuring FortiSIEM'.n AWS follows the same basic process described in Setting Access Credentials for Device Discovery and Discovering Devices, but requires a different approach to associating credentials to IP addresses, since AWS uses dynamic, rather than static, IP address assignment. The generic **AWS SDK** credential is used to discover Amazon Machine Instances (AMIs) and associated information such as host name, instance ID, and instance state, while credentials for generic versions of WMI, SMTP, and other access protocols are used to discover associated devices as you would for any other discovery process.

- Setting Access Credentials for AWS Instances
- Associating the AWS Host with Credentials

If you have not already configured Access Keys and permissions on AWS, please follow the steps outlined in AWS Access Key IAM Permissions and IAM Policies.

Setting Access Credentials for AWS Instances

1. Log into your Supervisor node.
2. Go to **Admin > Setup Wizard > Discovery**.
3. Under **Enter Credentials**, click **Add**.
4. Enter a **Name** for the credential.
5. For **Device Type**, select **Amazon AWS SDK**.
6. For **Access Protocol**, select **AWS SDK**.
7. For **Region**, enter the region where your AWS instance is located.
8. Enter the **Access Key ID** and **Secret Access Key** associated with your AWS instance.
9. Click **Save**.

Associating the AWS Host with Credentials

After you've defined all the credentials associated with the access protocols used by devices in your AWS instance, you need to associate those credentials to the AWS host. In other deployment configurations, you would associate credentials with IP addresses corresponding to your device locations, but since AWS uses dynamic IP addressing, you need to associate all your credentials to the same host.

1. Under **Enter IP Range to Credential Associations**, click **Add**.
2. For **IP/Host Name**, enter **amazon.com**.
3. Click **+**, and add the **AWS SDK** credential, as well as any other generic credentials you've created.
4. Click **OK**.
5. Click **Test Connectivity** to make sure you can reach your instance and that all credentials are entered correctly before you initiate discovery.
   Both the connectivity test and the discovery process will try to connect to the Amazon instances first, and from there will try to connect to the private IPs of discovered instances using the other access protocols.
6. You can now initiate discovery of your instances and associated devices as described in Discovering Devices, but for **Discovery Type**, select **AWS Scan**.

If discovery is successful, your discovered instances and devices will be added to **Admin > Setup wizard > Monitor Change/Performance**, and in **CMDB > Devices**, you will see an **Amazon EC2** directory, which will

include your discovered instances. If you have defined other access credentials, the discovered devices will also appear in that directory, as well as under **CMDB > Server**. You can query these devices from either directory.

Discovering Microsoft Azure Infrastructure

Discovering Microsoft Azure Cloud infrastructure follows the same basic process described in Setting Access Credentials for Device Discovery and Discovering Devices, but requires a different approach to associating credentials to IP addresses, since Azure uses dynamic, rather than static, IP address assignment.

- Create a Certificate file for communicating to Azure Management Server
- Setting Access Credentials for Microsoft Azure Discovery
- Associating Microsoft Azure with Credentials
- Discovering Microsoft Azure Compute Nodes

Create a Certificate file for communicating to Azure Management Server

1. Create a pem file.

```
openssl req -x509 -nodes -days 365 -newkey rsa:2048 -keyout azure-cert.pem
-out azure-cert.pem
```

2. Create the cert file.

```
openssl x509 -outform der -in azure-cert.pem -out azure-cert.cer
```

3. Login to the Azure old portal, upload the .cer to the Settings ->"Management Certificates" section.

Setting Access Credentials for Microsoft Azure Discovery

1. Log into your Supervisor node.
2. Go to **Admin > Setup Wizard > Credentials**.
3. Under **Enter Credentials**, click **Add**.
4. Enter a **Name** for the credential.
5. For **Device Type**, select **Microsoft Azure Compute**.
6. For **Subscription ID**, enter .
7. Upload the **Certificate File**, enter the region where your AWS instance is located.
8. Enter the **Access Key ID** and **Secret Access Key** associated with your AWS instance.
9. Click **Save**.

Associating Microsoft Azure with Credentials

After you've defined all the credentials associated with the access protocols used by devices in your Microsoft Azure instance, you need to associate those credentials.

1. Log into your Supervisor node.
2. Go to **Admin > Setup Wizard > Credentials**.
3. Under **Enter IP Range to Credential Associations**, click **Add**.
4. For **IP/Host Name**, enter **azure.com**.
5. Click **+**, and add the **Microsoft Azure Compute** credential created in "Setting Access Credentials for Microsoft Azure Discovery", as well as any other generic credentials you've created.
6. Click **OK**.

7. Click **Test Connectivity** to make sure you can reach your instance and that all credentials are entered correctly before you initiate discovery.

Discovering Microsoft Azure Compute Nodes

After you've defined and tested all the credentials, you can proceed to discovery.

1. Log into your Supervisor node.
2. Go to **Admin > Setup Wizard > Discovery**.
3. Click **Add**
4. For **Discovery Type**, select **Azure Scan.**
5. Click **Save.**
6. Select the entry just created and click **Discover.**

If discovery is successful, your discovered instances will be added to **Admin > Setup wizard > Monitor Change/Performance** and **CMDB > Devices > Microsoft Azure Cloud > Azure Compute.**

Approving Newly Discovered Devices

When devices are discovered by FortiSIEM, monitoring of them begins automatically, and incidents for those devices will trigger automatically based on the rules associated with that device. However, you can configure the Discovery Settings so incidents will be triggered only for devices you approve. If you select **Approved Devices Only** for **Allow Incident Firing On**, then you will need to approve devices before incidents will be triggered for those devices, but they will still be monitored and added to the CMDB.

1. Log in to your Supervisor node.
2. Go to **Admin > Discovery Results**.
3. Select a discovery result.
4. Click **View Changes**.
5. Expand the folder **Discovery Delta**.
6. Expand the folder **New Devices**.
7. Select the devices you want to approve, and click **Approve Selected**.
   You can approve all the new devices by selecting the **New Devices** folder, and then click **Approve All**.

Related Links

- Discovery Settings

Inspecting Event Pulling Methods for Devices

Once you have discovered and approved the devices in your IT infrastructure, you should verify that the FortiSIEM `perfMonitor` module is polling them over the correct access protocol and pulling event information from them. If you are having issues collecting performance metrics from your devices, you should begin troubleshooting by first checking the status of the event pulling method for the device.

1. Go to **Admin > Setup Wizard > Pull Events**.
2. Review the **Event Pulling Status** for each of your discovered devices.

Status	Description
Successful	If event information is being pulled from the device, you will see the name of the event pulling method rendered in plain black text.
Added but Not Monitored	If the name of the event pulling method has a **Star** icon next to it, event information can be successfully pulled from the device, but the `perfMonitor` module has not yet initiated monitoring.
Paused	A **Pause** icon indicates that event information is not being pulled from the device because it failed the verification check at the beginning of the monitoring cycle. This is usually caused by an issue with the access protocol credentials. The credential was valid when discovery succeeded, and so the event pulling method was able to monitor the associated metrics, but the `perfMonitor` module failed on the credential at a later time. You should check the access protocol credentials associated with the devices and event pulling methods, and then re-initiate discovery of the device.
Failed	An **Alert** icon and the name of the event pulling method in red indicates that adding that event pulling method for the device failed.

3. Click **Show Errors** to view a more detailed description of any errors associated with an event pulling method.
4. Click **Edit** to change any of the event pulling methods associated with a device.
5. Click **Apply** to apply any changes to your event pulling methods.
6. Click **Test Pull Events** to test any changes you make.

Inspecting Changes Since Last Discovery

After you run discovery for the first time, FortiSIEM keeps track of changes to your discovered devices during subsequent discovery runs, including new devices, changed devices, and failed devices.

1. Log in to your Supervisor node.
2. Go to **Admin > Discovery Results**.
3. Select a discovery result.
4. Click **View Changes**.
5. Expand the folder **Discovery Delta**.
6. Move your mouse cursor over a folder or item until a blue Information icon appears, and then click on the icon to view basic information about the item.

Discovery Range Definition Options

When you set the range definition for your discovery processes, several options are available for how you want the discovery process to run.

Discover Routes	Selected by default, if you clear this option then discovery will not use the route table to find next hop devices. This can be useful if your network includes border routers, which can significantly impact the time required for the discovery process.	
Discovery Type	Four types of scans are available for the discovery process:	
	Smart Scan	Smart Scan is an optimized search method in which only the live devices in the network are searched. To use Smart Scan, you first provide a root device (typically the first hop Layer 3 router). FortiSIEM then discovers the root device and learns of its first hop neighbors from the ARP table. These devices are then discovered using existing credentials, and their one hop neighbors are subsequently discovered. This continues until no more devices are discovered. Often a single Layer 3 router, switch, or firewall is sufficient to discover the entire network. However, if a firewall that can block SNMP is installed, then devices on either side of the firewall need to be provided as root devices. Smart Scan is usually faster than Range Scan, but in rare cases discovery can miss a device when it is quiet and not present in the ARP table of adjacent devices.
	Range Scan (default)	In contrast to Smart Scan, Range Scan is a brute force method in which FortiSIEM attempts to discover all the devices in the IP ranges you provide. With Range Scan, FortiSIEM will first attempt to ping a device, and if that succeeds, discovery will proceed.
	AWS Scan	AWS Scan is used to discover devices in Amazon Web Services. See Discovering Amazon Web Services (AWS) Infrastructure for more information.
	L2 Scan	L2 Scan is used to update the Layer 2 connectivity information used in the Identity and Location report. It does not discover system and application monitors, installed and running software, or users and groups, and, in contrast to the other scan methods, it does not update the CMDB and executes more quickly.
Do Not Ping Before Discovery	To save time, FortiSIEM first attempts to reach devices by ping before initiating discovery. You should select this option if ping has been disabled for your network, otherwise discovery will fail.	
Include Powered Off VMs	By default, only powered on VMs are discovered.	
Include VM Templates	By default, VM templates are not discovered.	
Include/Exclude Device Types	Click the Edit icon to select devices that you want to include or exclude from the discovery process. Note that if you have entries for both of these option, the discovery process will prioritize included devices over excluded ones.	
Include/Exclude Domains (AWS Only)	Enter the domains you want to include or exclude from the discovery process.	

Include/Exclude Ranges	Enter the IP addresses or host names you want to include or exclude from the discovery process.
Include/Exclude Zones (AWS Only)	Enter the zones you want to include or exclude from the discovery process.
Only Discover Devices not in CMDB	If you select this option, discovery will only find those devices whose IP addresses do not match the address of any device in CMDB. To make an exception to this rule, specify a list of IP addresses in the Exclude Ranges field. The primary use case for this is for indirect device discovery such as VCenter-based VM discovery, or WLAN controller-based access point discovery. By specifying the VCenter IP address in the Exclude Ranges field, new guest VMs can always be discovered even if the VCenter is already in the CMDB.
Ping Only Discovery	Select this option if you are only interested in discovering whether a device or service is up or down.
Root IPs	For Smart Scan only, provide the root IPs from which you want the Smart Scan to start.

Scheduling a Discovery

Discovery can be a long-running process when performed on a large network, or over a large IP range, and so you may want to schedule it to occur when there is less load on your network or during off hours. You may also want to set up a schedule for the process to run and discover new devices on a regular basis.

1. Log in to your Supervisor node.
2. Go to **Admin > Setup Wizard > Discovery**.
3. Click **Schedule**.
4. Click the **+** icon.
5. Select from the available ranges.
   You can select multiple ranges and set the order in which discovery will run on them by using the up and down arrows.
6. Set the time at which you want discovery to run.
7. For a one-time scheduled discovery, enter a **Date** for the discovery to run.
8. For recurring discoveries, select how often (hourly, daily, weekly, monthly), you want discovery to run, and then enter other scheduling options.
9. Click **OK**.

Adding Devices to the CMDB Outside of Discovery

There are situations in which you may want to add devices to the Configuration Management Database (CMDB) outside of the discovery procedure. For example, FortiSIEM needs access to devices over SNMP or WMI to discover them, but you may have devices in your infrastructure that don't utilize these access protocols. The IP addresses for those devices will still be contained in traffic logs, and rules may need to incorporate that device. In order to make sure that logs are parsed correctly and rules function as expected, you need to make sure that these undiscovered devices are associated with an IP address. Adding a device directly to the CMDB lets you provide the information necessary for FortiSIEM to recognize the device, including associating it with an IP address or range.

Adding Devices to Device Groups

When you add a device to the CMDB manually, make sure to choose the group, such **Firewall**, **Printers**, or **Storage**, in the **Device View** where you want to add it. If you only add it to the top-most **Devices** group, it will not be added to the topology map correctly.

1. Log into your Supervisor node.
2. Click **CMDB**.
3. In the **Device View**, select **Devices**, then select the sub-category where you want to add the device.
4. In the summary pane, click **New**.
5. For **Summary**, **Contact**, **Interfaces**, and **Properties**, enter information for the new device.

Entering Interface Information

When you enter the interface information for the device, make sure to provide the correct IP address and network mask for the interfaces. FortiSIEM will use this network information to generate the Network Segments for the device.

6. Click **Save** when you're done adding the device information.

Related Links

- Adding a Synthetic Monitoring Test to a Business Service

Decommissioning a device

Decommissioning a device lets you re-assign the IP address to a new device but still keep the old device in CMDB for historical purposes.

To decommission a device:

1. Go to **CMDB** > **Devices**
2. Select the device.
3. Click on the menu under **Name** and select **Decommission**.
4. Provide a **Reason** and Select **OK** to decommission the device
5. Consequences of decommissioning
    a. Device will be moved to **CMDB** > **Devices** > **Decommission** folder
    b. Device will be removed from maintenance calendars
    c. Performance monitoring will stop
    d. A new device with the same IP can be discovered

To re-commission the device:

1. Go to **CMDB** > **Devices > Decommission**
2. Select the device.
3. Click on the menu under **Name** and select **Recommission**.
4. The device will be moved back to the folder where it was when it was decommissioned.
5. Performance monitoring will resume

Creating Dynamic CMDB Group Policies

This setting allows you to write rules to put devices in CMDB Device Group and Business Service Groups of your choice. When a device is discovered, the policies defined here are applied and the device is assigned to the group (s) defined in the matching policies.

To create a new CMDB Group Policy:

1. Go to **Admin > General Settings > Discovery > CMDB Group**.
2. Click Add.
3. For matching conditions - enter the following information
    a. **Organization** - the organization which this rule applies to
    b. **Vendor** - the matching device vendor - select from the list
    c. **Model** - the matching device model - select from the list
    d. **Host Name** - matching device host name via regular expression match
    e. **IP Range** - matching device access IP - format is single IP, IP range, CIDR
4. For Actions (Add To) - enter the following information
    a. **Groups** - specify the groups which the matching devices will be added to
    b. **Biz Services** - specify the business services which the matching devices will be added to

This device grouping does not overwrite the CMDB Device group assigned during discovery. The grouping defined here is in addition to the discovery defined CMDB group.

Conditions are matched in ANDed manner

Both the actions are taken, that is, if both a Group and a Business Service is specified, then the device will be added to both the specified Group and Business Service.

To apply one or more CMDB Group policies,

1. Select one or more policies and click **Apply** or Click **Apply All** to apply all policies.
2. Once a policy is saved, then next discovery will apply these policies. That means, discovered devices will belong to the groups and business services defined in the policies.

Configuring Monitoring

Once FortiSIEM discovers your devices, they will monitored continuously, and you can use the data collected to analyze the performance of your infrastructure. You can also configure FortiSIEM to send notifications when events that meet specific conditions occur in your infrastructure.

You can disable the collection of metrics for specific devices, disable devices for monitoring, and change the polling interval for metric collection. Some devices need to be configured to send logs to FortiSIEM, as described in the topics under *FortiSIEM External Systems Configuration Guide*. You can also configure FortiSIEM to monitor important ports, processes, and interfaces, and set up monitoring tests that use synthetic transaction to make sure that critical services are up and running.

- Device Monitoring Settings
- Managing Monitoring of System and Application Metrics for Devices
- Setting Up Synthetic Transaction Monitoring Tests

Device Monitoring Settings

While FortiSIEM constantly monitors and reports on your IT infrastructure, there are several settings you can use to refine reporting on critical interfaces, important processes and ports, and disk utilization.

- Adding Important Interfaces
- Adding Important Processes
- Adding Important Ports
- Excluding Disks from Disk Capacity Utilization Monitoring

Adding Important Interfaces

This setting allows you to always get interface utilization reports on a set of important network interfaces across all device types.

Behavior prior to Release 4.8.1

FortiSIEM continuously monitored **every** network interface on every server and network device for utilization and up/down. By marking an interface as Critical, (a) up/down was monitored for Critical interfaces only but (b) all interfaces were monitored for utilization and critical ones were marked in the event. Since all interfaces were monitored for utilization - a large number of events may be generated, as certain devices such as voice gateways can have many logical interfaces.

Behavior in Release 4.8.1 onwards

FortiSIEM would monitor **only** the interfaces marked as Critical in this tab. **So it is important for to define ALL critical interface at once.**

Important Interface Setup after 4.8.1 Upgrade

The behavior of interface monitoring has dramatically changed since 4.8. So it is very important to follow these steps.

1. Create a list of all Important interfaces
2. Go to **Admin > General Settings > Monitoring > Important Interfaces**
3. Click **Enable.** This will stop all interface monitoring**.**
4. Click **Add.**
5. Select either **Device View** or **Interface View**.
6. Select a device to view and select its interfaces, or select an interface.
7. Click **OK** to add the selected interface to the list. The **Critical** and **Monitor** boxes would be automatically checked.
8. Check the **WAN** box if applicable. If checked, the interface utilization events would have isWAN = "yes" attribute. You can use this to run a report for all WAN interfaces.
9. Click **Apply All**. Now FortiSIEM will start monitoring only the selected interfaces in this tab will be monitored.
10. If you want to disable this behavior and return to ALL interface monitoring (as in releases prior to 4.8), then click **Disable**.

Adding Important Processes

This setting allows you to always get process resource utilization reports and up/down alerts on a set of important processes across all device types.

Behavior prior to Release 4.8.1

FortiSIEM continuously monitors **every** process on every server and network device for resource utilization and up/down. By marking an process as Critical, (a) up/down was monitored for Critical processes only, but (b) all processes were monitored for resource utilization. Since there are a large number of processes across all device types, a large number of events can be generated.

Behavior in Release 4.8.1 onwards

FortiSIEM would monitor **only** the processes marked as Critical in this tab. **It is important to define ALL critical processes at once.**

Important Process Setup after 4.8.1 Upgrade

The behavior of process utilization monitoring has dramatically changed since 4.8. So it is very important to follow these steps.

1. Create a list of all Important processes.
2. Go to **Admin > General Settings > Monitoring > Important Processes.**
3. Click **Enable.** This will stop monitoring all processes**.**
4. Click **Add.**
5. Enter a **Process Name** and any **Parameters**, and then click **OK**.
6. Click **Apply All**. FortiSIEM will start monitoring only the selected processes in this tab.
7. If you want to disable this behavior and return to ALL process monitoring, then click **Disable**.

Adding Important Ports

Always reporting the UP/DOWN status for every TCP/UDP port on every server can consume a significant amount of resources. FortiSIEM will report the UP/DOWN status only for the ports you add to the **Important Ports** list. Matching is exact based on port number and IP protocol.

1. Go to **Admin > General Settings > Monitoring**.
2. Under **Important Ports**, click **Add**.
3. Enter the **Port Number** and select the **Port Type**.
4. Click **OK**.
5. Click **Apply All**.

Excluding Disks from Disk Capacity Utilization Monitoring

You can exclude disks from disk capacity utilization monitoring. Disk capacity utilization events will not be generated for devices matching the device name, access IP, and disk name that you provide. Incidents will not trigger for these events, and the disks will not show up in summary dashboards.

Exclude Stable, Almost-Full Disks

Use this list to exclude read-only disk volumes or partitions that do not grow in size and are almost full. This will prevent from these servers from always showing a **CRITICAL** status in dashboards.

1. Go to **Admin > General Settings > Monitoring**.
2. Under **Excluded Disks**, click **Add**.
3. Select a device to view its disks, and then select the disk you want to exclude from monitoring.
4. Click **OK**.
5. Click **Apply All**.

Managing Monitoring of System and Application Metrics for Devices

When FortiSIEM discovers devices, it also discovers the system and application metics that can be monitored for each device, and displays these in the **Monitor Change/Performance** tab of the **Setup Wizard**. Here you can also disable the monitoring of specific metrics for devices, disable devices from being monitored, and change the polling interval for specific metrics. See Inspecting Event Pulling Methods for Devices for an explanation of the different status indicators for **System Monitor** and **Application Monitor** metrics.

1. Go to **Admin > Setup Wizard > Monitor Change/Performance**.
2. Click **Refresh** to make sure you have the latest list of devices.
   Single Line Display

   Select **Single Line Display** to view each device, along with its list of application and system monitors, on a single line.
3. To disable monitoring for a device, clear the **Enable** option for it.
4. To enable or disable monitoring of a specific metrics for a device, click on a device to select it, then click **Edit** and select **System Monitoring** or **Application Monitoring** to view the list of metrics associated with that monitor and device. You can also enable or disable the metrics for a device's monitor type by clicking on the **System Monitoring** or **Application Monitoring** section for the device.
5. To change the polling interval for a metric, in the **More** menu, select **Set Intervals**. Select the **Monitor Type** and **Device**, and then set the interval.
6. When you are done making changes, click **Apply**.

Setting Up Synthetic Transaction Monitoring Tests

A Synthetic Transaction Monitoring (STM) test lets you test whether a service is up or down, and measure the response time. An STM test can range from something as simple as pinging a service, to something as complex as sending and receiving an email or a nested Web transaction. Setting up an STM test involves defining the type of monitor, associating the monitor definition to a device and testing it, and then deploying the STM test to a Supervisor or Collector. You can view the results of STM tests in the **Synthetic Transaction Monitoring** page, either by navigating to **Summary Dashboard > Availability/Performance > Application Summary > Synthetic Transaction Monitoring**, or to **Admin > Setup Wizard > Synthetic Transaction Monitoring**, and then clicking on **Monitoring Status**. You can also report on the results of STM tests in the reports **Top Applications By Synthetic Transaction Response Time** and **Top Applications By Synthetic Transaction Response Time - Detailed view.** When an STM test fails, three system rules are triggered, and you can receive an email notification of that failure by creating a notification policy for these rules.

System Rule	Description
Service Degraded - Slow Response to STM	Detects that the response time of an end-user monitored service is greater than a defined threshold (average over 3 samples in 15 minutes is more than 5 seconds)
Service Down - No Response to STM	Detects a service suddenly went down from the up state and is no longer responding to synthetic transaction monitoring probes.
Service Staying Down - No Response to STM	Detects a service staying down, meaning that it went from up to down and did not come up, and is no longer responding to end user monitoring probes

1. Go to **Admin > Setup Wizard > Synthetic Transaction Monitoring**.
2. Click **Add**.
3. Enter a **Name** and **Description** for the test.
4. For **Frequency**, enter how often, in minutes, you want the test to run.
5. Select the **Protocol** for your test.
   See Protocol Settings for Synthetic Transaction Monitoring Tests for more information about the settings and test results for specific protocols.
6. Click **Save**.
   You now have to associate the STM test with a target host name, IP address, or IP range.
7. Click **Create and Test**.
8. For **Monitoring Definition** select one of the STM tests you have created.
9. For **Host Name** or **IP/Range**, enter the information for your STM test target.
10. For **Port**, click **+** and enter any ports to use when connecting to the target with this test.
11. Click **OK**.
    FortiSIEM will run the test and verify if it is successful. If it succeeds, it will be added to the list of tests with a yellow **Star** next to it, indicating that it has been added but is not yet running.
12. Click **Apply All** to begin executing your tests at their set frequency.
    The yellow Star will be removed from your test after it executes against the target the first time

Protocol Settings for Synthetic Transaction Monitoring Tests

This table describes the settings associated with the various protocols used for setting up Synthetic Transaction Monitoring tests.

Protocol	Description	Settings	Notes
Ping	Checks packet loss and round trip time	**Maximum Packet Loss PCT**: tolerable packet loss **Maximum Average Round Trip Time**: tolerable round trip time (seconds) from FortiSIEM to the destination and back If either of these two thresholds are exceeded, then the test is considered as failed.	
LOOP Email	This test sends an email to an outbound SMTP server and then attempts to receive the same email from a mailbox via IMAP or POP. It also records the end-to-end time.	**Timeout** : the time limit by which the end to end LOOP EMAIL test must complete. **Outgoing Settings**: these specify the outgoing SMTP server account for sending the email. • **SMTP Server**: name of the SMTP server • **User Name**: user account on the SMTP server • **Email Subject**: content of the subject line in the test email **Incoming Settings**: These specify the inbound IMAP or POP server account for fetching the email. • **Protocol Type**: choose IMAP or POP • **Server**: name of the IMAP or POP server • **User Name**: user account on the IMAP or POP server • **Email Subject**: content of the subject line in the test email	Before you set up the test you will need to have set up access credentials for an outbound SMTP account for sending email, and an inbound POP/IMAP account for receiving email

Protocol	Description	Settings	Notes
HTTP(S) - Selenium Script	This test uses a Selenium script to play back a series of website actions in FortiSIEM.	**Upload**: select the java file you exported from Selenium T**otal Timeout**: the script must complete by this time or the test will be considered failed **Step Timeout**: each step must complete by this time	**How to export**: • Make sure Selenium IDE is installed within Firefox browser • Open Firefox • Launch Tools > Selenium IDE. From now on, Selenium is recording user actions • Visit websites • Once done, stop recording • Click File > Export Test case as > Java / Junit 4 /WebDriver • Save the file as .java in your desktop. This file has to be inputted in FortiSIEM.
HTTP(S) - Simple	This test connects to a URI over HTTP(s) and checks the response time and expected results	**URI**: the URI to connect to **Authentication**: any authentication method to use when connecting to this URI **Timeout**: t his is the primary success criterion - if there is no response within the time specified here, then the test fails **Contains**: an expected string in the test results **Does Not Contain**: a string that should not be contained in the test results **Response Code**: an expected HTTP(S) response code in the test results. The default is set to **200 - 204**.	

Protocol	Description	Settings	Notes
HTTP(S) - Advanced	This test uses HTTP requests to connect to a URI over HTTP(s), and checks the response time and expected results	Click **+** to add an HTTP request to run against a URI. **URI**: the URI to run the test against **SSL**: Whether or not to use SSL when connecting to the URI, and the port to connect on **Authentication**: the type of authentication use when connecting to the URI **Timeout**: t his is the primary success criterion - if there is no response within the time specified here, then the test fails **Method Type**: the type of HTTP request to use **Send Parameters**: click **+** or the **Pencil** icon to add or edit any parameters for the request **Contains**: an expected string in the test results **Does Not Contain**: a string that should not be contained in the test results **Response Code** : an expected HTTP(S) response code in the test results. The default is set to **200 - 204** . **Store Variables as Response Data for Later Use:** click **+** or the **Pencil** icon to add or edit any variable patterns that should be used as data for later tests	
TCP	This test attempts to connect to the specified port using TCP	**Timeout**: this is the single success criterion. If there is no response within the time specified here, then the test fails.	

Protocol	Description	Settings	Notes
DNS	Checks response time and expected IP address	**Query**: the domain name that needs to be resolved **Record Type**: the type of record to test against **Result**: specify the expected IP address that should be associated with the DNS entry **Timeout**: this is the primary success criterion - if there is no response within the time specified here, then the test fails	
SSH	This test issues a command to the remote server over SSH, and checks the response time and expected results	**Remote Command**: the command to run after logging on to the system **Timeout**: this is the primary success criterion - if there is no response within the time specified here, then the test fails **Contains**: an expected string in the test results	You will need to have set up an SSH credential on the target server before setting up this test. As an example test, you could set **Raw Command** to `ls`, and then set **Contains** to the name of a file that should be returned when that command executes on the target server and directory.
LDAP	This test connects to the LDAP server, and checks the response time and expected results	**Base DN**: an LDAP base DN you want to run the test against **Filter**: any filter criteria for the Base DN **Scope**: any scope for the test **Timeout**: t his is the primary success criterion - if there is no response within the time specified here, then the test fails **Number of Rows**: the expected number of rows in the test results **Contains**: an expected string in the test results **Does Not Contain**: a string that should not be contained in the test results	You will need to have set up an access credential for the LDAP server before you can set up this test

Protocol	Description	Settings	Notes
IMAP	This tests checks connectivity to the IMAP service	**Timeout** : t his is the single success criterion - if there is no response within the time specified here, then the test fails	
POP	This test checks connectivity to the IMAP service	**Timeout** : t his is the single success criterion - if there is no response within the time specified here, then the test fails	
SMTP	This test checks connectivity to the SMTP service	**Timeout** : t his is the single success criterion - if there is no response within the time specified here, then the test fails	
JDBC	This test issues a SQL command over JDBC to a target database, and checks the response time and expected results	**JDBC Type**: the type of database to connect to **Database Name**: the name of the target database **SQL**: the SQL command to run against the target database **Timeout** : t his is the primary success criterion - if there is no response within the time specified here, then the test fails **Number of Rows**: the expected number of rows in the test results **Contains**: an expected string in the test results **Does Not Contain**: a string that should not be contained in the test results	
FTP	This test issues a FTP command to the server and checks expected results	**Anonymous Login**: choose whether to use anonymous login to connect to the FTP directory **Remote Directory**: the remote directory to connect to **Timeout** : t his is the primary success criterion - if there is no response within the time specified here, then the test fails	

Protocol	Description	Settings	Notes
TRACE ROUTE	This test issues a trace route command to the destination and parses the results to create PH_DEV_MON_TRACEROUTE events, one for each hop.	**Timeout**: If there is no response from the system within the time specified here, then the test fails. **Protocol Type**: Specifies the IP protocol over which trace route packets are send - current options are UDP, TCP and ICMP **Max TTL**: Max time to live (hop) value used in outgoing trace route probe packets. **Wait Time**: Max time in seconds to wait for a trace route probe response	For the trace route from AO to destination D via hops H1, H2, H3, FortiSIEM generates 3 hop by hop PH_DEV_MON_TRACEROUTE events. **First event:** Source AO, destination H1, Min/Max/Avg RTT, Packet Loss for this hop **Second event:** Source H1, destination H2, Min/Max/Avg RTT, Packet Loss for this hop **Third event:** Source H2, destination H3, Min/Max/Avg RTT, Packet Loss for this hop **Fourth event:** Source H3, destination D, Min/Max/Avg RTT, Packet Loss for this hop

Adding a Synthetic Monitoring Test to a Business Service

You may want to add a Synthetic Transaction Monitoring (STM) test to a Business Service as part of the monitoring infrastructure for that service. However, in order to enable reporting on that STM, you need to add it to the business service as a device that FortiSIEM can then report on. This topic explains how to create a device for an STM test and add it to your business service report.

1. Create your STM as described in Setting Up Synthetic Transaction Monitoring Tests.
2. Note the IP address that your STM resolves to in Step 9 of the setup instructions.
3. In the **CMDB** tab, select **Devices**, and then select a subcategory where you want to add the STM device. You may want to create your own group where you manage your STM devices.
4. In the summary pane for the device subcategory, click **New**.
5. Complete all relevant information for the STM device, providing the IP address/range from Step 2 in the **Access IP** field of the **Summary** page.
6. Click **Save** when you're done entering device information for the STM.
7. Follow the instructions in Creating a Report to add information about the STM device to a business service report, and then use the instructions in Adding Widgets to Dashboards to add it to your dashboard.

Related Links

- Adding Devices to the CMDB Outside of Discovery
- Creating CMDB Groups and Adding Objects to Them
- Creating a Report
- Adding Widgets to Dashboards

Creating Business/IT Services

By defining an IT or Business Service, you can create a logical grouping of devices and IT components which can be monitored together.

1. Log in to your Supervisor node.
2. Go to **CMDB > Business Services**.
3. Click **New**.
4. Enter a **Name** and **Description** for the business service.
5. Select a **Device/Application Group**, and when the list of associated devices loads into the selection pane, select a device and click **>>** to add it to the **Selected Devices/Applications** for the business service.
6. Click Save when you're done adding devices to the business service.

After you have created a business service, you can select it, and the **Show Topology** option, to view it within overall IT topology. You can also use the links in the **Analysis** menu of the Business Services summary dashboard to find out more information about incidents, device availability, device and application performance, interface and event status, and real-time and historical search for a selected business service.

Data Update Subscription Service

FortiSIEM is constantly developing support for additional IT infrastructure devices. By subscribing to the FortiSIEM Data Update Service, you can receive updates when support for new devices becomes available, rather than waiting for it to be included in a formal release. In addition to devices you can also receive new rules, reports, parser updates etc.

- Data Update Overview
- Configuring Data Update

Data Update Overview

FortiSIEM data update subscription service updates your FortiSIEM deployment with the latest device support related data as it becomes available, rather than having to wait for it to be included in a formal release.

The following items can be included in an update

- New event attribute
- New event types
- New device type
- New parsers or modifications for existing parsers
- Performance monitoring templates for new devices or modified ones for existing devices
- New rules or modifications for existing rules
- New reports or modifications for existing reports - both CMDB report and event based reports
- New groups or modifications for existing groups for Event Types, Rules, Reports, Device Groups, Application Groups
- Code to handle new devices

Configuring Data Update

Provide a brief (two to three sentence) description of the task or the context for the task.

Prerequisites

- Contact FortiSIEM support and make sure that your license includes Data Update Service.
- Make sure you have **Data Update URL -** this is typically **https://images.FortiSIEM.net/upgrade/ds**- contact FortiSIEM to make sure that this information has not changed.
- Make sure you have license credentials.

Procedure

Configure Data Update Server Setting

1. Log on to FortiSIEM Supervisor with Administrator credentials
2. Go to **Admin** > **General Settings** > **System**
3. Configure Data Update Server Setting
    1. Enter **Data Update URL** (see prerequisites)
    2. Enter **Server Username** and **Server Password** - these are the license credentials
    3. Specify **Notify Email** (optional) - you will receive email when new data updates are available
    4. Click **Save**

Check Available Data Updates

1. Log on to FortiSIEM Supervisor with Administrator credentials.
2. Go to **Admin** > **Data Update**.
3. Click **Refresh**.
    1. Available data updates are shown on left.
    2. Click a version on the left and the contents for that version is shown on the right.
4. Check the current data version from **Admin** > **Cloud Health** > **Data Update Version**. The number after 3rd decimal is the data version. For example 4.4.1.38 means data version is 38.
5. Note the data version you would like to upgrade to.

Apply Data Update on Supervisor

1. SSH to FortiSIEM Supervisor as root.
2. Go to /pbin.
3. Download the data version by running ./phdownloaddata and specify the data version you would like to upgrade to.
4. Install the data version by running ./phinstalldata.

Apply Data Update on Collectors

1. Log on to FortiSIEM Supervisor with Administrator credentials.
2. Go to **Admin** > **Collector Health**.
    1. Select a Collector.
    2. Click **Download Data Update** - this downloads the data files to the collector.
    3. Click **Install Data Update** - this installs the data files on the collector.
    4. Repeat for all collectors.

Check whether Data Update Installed Successfully

1. Log on to FortiSIEM Supervisor with Administrator credentials.
2. Check **Admin** > **Cloud Health** > **Data Update Version**.
3. Check **Admin** > **Collector Health** > **Data Update Version**.

Creating Custom Parsers and Monitors for Devices

Creating a custom parser for device logs involves writing an XML specification for the parser, and then using a test event to make sure the logs are parsed correctly. Creating a custom monitor involves defining a performance object that you want to monitor, associating that performance object to a device type, event type, and event attribute type, and then testing to make sure that the monitored metrics are correctly received by FortiSIEM. You can create custom monitors for system and application performance, command outputs, and file monitoring.

Creating Event Attributes, Event Types, and Device Types

When you create a custom parser or monitor, you must also specify the device, application, event type, and event attribute to which it applies. If these objects aren't already included in the FortiSIEM CMDB, you can create them as a preliminary step to creating your parser or monitor.

- Creating Device and Application Types
- Creating Event Attribute Types
- Creating Event Types

Creating Device and Application Types

If the device or application that you want to create a parser or monitor for isn't already listed in **Admin > Device Support > Device/App Types**, you can add it.

1. Go to **Admin > Device Support > Device/App Types**.
2. Click **New**, and then choose **New Device Type** or **New Application Type**.
3. Enter the information for the new device or application type.

Device Type	• Vendor • Model • Version • Device/App Group • Biz Service group • Description
Application Type	• Vendor • Model • Version • Device/App Group • Biz Service group • Application Package Group • Description

4. Click **Save**.

Creating Event Types

After parsing an event or log, FortiSIEM assigns a unique event type to that event/log. When you create a new custom parser for device logs, you almost always have to add a new event type to FortiSIEM so the log events can be identified.

Naming Custom Event Types

All custom event types must begin with the prefix `P H_DEV_MON_CUST_ .`

1. Go to **Admin > Device Support > Event Types**.
2. Click **New**.
3. Enter a **Name** for the new event type.
4. Select the **Device Type** to associate with the event type.
   If the device type isn't included in the menu options, you can add it to FortiSIEM.
5. Select the **Event Type Group** category for this event type.
6. Select a **Severity** to associate with the event type.
7. Enter an optional **Description**.
8. Click **Save**.

Creating Event Attribute Types

Event attributes are used to capture parsed information from events. You only have to create a new attribute if the one you want use for your custom parser or monitor is not listed in **Admin > Device Support > Event Attribute Types**.

Creating an Event Attribute Type by Cloning

You can clone an existing event attribute type to use as the basis for a new one. Select the event attribute type you want to use, click Clone, and then modify as necessary.

1. Go to **Admin > Device Support > Event Attribute Types**.
2. Click **New**.
3. Enter a **Name** and **Display Name**.
4. Select the **Value Type** to associate with the event attribute type.
5. Optionally enter a **Display Format Type** and **Description**.
6. Click **Save**.

Custom Parsers

To start creating a custom parser for device logs, you should begin by reviewing the Event Parser XML Specification. Writing the XML specification is the primary task in creating a custom parser.

- Event Parser XML Specification
- Creating a Custom Parser
- Deleting or Disabling a Parser
- Exporting a Custom Parser
- Importing a Custom Parser
- Parser Examples

Event Parser XML Specification

FortiSIEM uses an XML-based parser framework to parse events. These topics describe the parser syntax and include examples of XML parser specifications.

- Custom Parser XML Specification Template
- Parser Name Specification
- Device or Application Type Specification
- Format Recognizer Specification
- Pattern Definition Specification
- Parsing Instructions Specification

Custom Parser XML Specification Template

The basic template for a custom parser XML specification includes five sections. Click on the name of any section for more information.

Section	Description
Parser Name Specification	Name of the parser file
Device Type	The type of device or application associated with the parser
Format Recognizer Specification	Patterns that determine whether an event will be parsed by this parser
Pattern Definition Specification	Defines the parsing patterns that are iterated over by the parsing instructions
Parsing Instructions Specification	Instructions on how to parse events that match the format recognizer patterns

Custom Parser XML Specification Template

```
<eventParser name="xxx">   <deviceType> </deviceType>   <eventFormatRecognizer>
</eventFormatRecognizer>   <patternDefinitions> </patternDefinitions>   <pars-
ingInstructions> </parsingInstructions></eventParser>
```

Parser Name Specification

This section specifies the name of the parser, which is used only for readability and identifying the device type associated with the parser.

```
<eventParser name="CiscoIOSParser"></eventParser>
```

Device or Application Type Specification

This section specifies the device or the application to which this parser applies. The device and application definitions enable FortiSIEM to detect the device and application type for a host from the received events. This is called **log-based discovery** in FortiSIEM. Once a received event is successfully parsed by this file, a CMDB entry is created with the device and application set from this file. FortiSIEM discovery may further refine the device.

There are two separate subsections for device and application. In each section, vendor, model and version can be specified, but version is not typically needed.

Set Version to Any

In the examples in this topic, `<Version>` is set to `ANY` because events are generally not tied to a particular version of a device or software. You could of course set this to a specific version number if you only wanted this parser to apply to a specific version of an application or device.

Vendor and Model Must Match the FortiSIEM Version

`<Vendor>` and `<Model>` entries must match the spelling and capitalization in the CMDB.

Examples of Specifications for Types of Device and Applications

Hardware Appliances

In this case, the type of event being parsed specifies the device type, for example Cisco IOS, Cisco ASA, etc.

```
<deviceType>    <Vendor>Cisco</Vendor>    <Model>IOS</Model>    <Ver-
sion>ANY</Version></deviceType>
```

Software Operating Systems that Specify the Device Type

In this case, the type of events being parsed specifies the device type, for example Microsoft Windows etc. In this case the device type section looks like

```
<deviceType>    <Vendor>Microsoft</Vendor>    <Model>Windows</Model>    <Ver-
sion>ANY</Version></deviceType>
```

Applications that Specify Both Device Type and Application

In this case, the events being parsed specify the device and application types because Microsoft SQL Server can only run on Microsoft Windows OS.

```
<deviceType>    <Vendor>Microsoft</Vendor>    <Model>Windows</Model>    <Ver-
sion>ANY</Version></deviceType><appType>    <Vendor>Microsoft</Vendor>
<Model>SQL Server</Model>    <Version>ANY</Version>    <Name> Microsoft SQL Server-
</Name></appType>
```

Applications that Specify the Application Type but Not the Device Type

Consider the example of an Oracle database server, which can run on both Windows and Linux operating systems. In this case, the device type is set to **Generic** but the application is specific. FortiSIEM depends on discovery to identify the device type.

```
<deviceType>    <Vendor>Generic</Vendor>    <Model>Generic</Model>    <Ver-
sion>ANY</Version></deviceType><appType>    <Vendor>Oracle</Vendor>    <Model>Data-
base Server</Model>    <Version>ANY</Version>    <Name>Oracle Database
Server</Name></appType>
```

Format Recognizer Specification

In many cases, events associated with a device or application will contain a unique pattern. You can enter a regular expression in the Format Recognizer section of the parser XML file to search for this pattern, which, if found, will then parse the events according to the parser instructions. After the first match, the event source IP to parser file map is cached, and only that parser file is used for all events from that source IP. A notable exception is when events from disparate sources are received via a syslog server, but that case is handled differently.

While not a required part of the parser specification, a format recognizer can speed up event parsing, especially when one parsing pattern file among many pattern files must be chosen. Only one pattern check can determine whether the parsing file must be used or not. The other less efficient option would be to examine patterns in every file. At the same time, the format recognizer must be carefully chosen so that it is not so broad to misclassify events into wrong files, and at the same time, not so narrow that it fails at classifying the right file.

Order in Which Parsers are Used

FortiSIEM parser processes the files in the specific order listed in the file `parserOrder.csv`.

Format Recognizer Syntax

The specification for the format recognizer section is:

```
<eventFormatRecognizer><![CDATA[regexpattern]]></eventFormatRecognizer>
```

In the regexpattern block, a pattern can be directly specified using regex or a previously defined pattern (in the pattern definition section in this file or in the GeneralPatternDefinitions.xml file) can be referenced.

Example Format Recognizers

Cisco IOS

All Cisco IOS events have a `%module name` pattern.

```
<patternDefinitions>    <pattern name="patCiscoIOSMod" list="begin"><![CDATA
[FW|SEC|SEC_LOGIN|SYS|SNMP|]]></pattern>    <pattern name="patCiscoIOSMod" list-
t="continue"><![CDATA[LINK|SPANTREE|LINEPROTO|DTP|PARSER|]]></pattern>    <pattern
name="patCiscoIOSMod" list="end"><![CDATA[CDP|DHCPD|CONTROLLER|PORT_SECURITY-
SP]]></pattern></patternDefinitions><eventFormatRecognizer><![CDATA[: %<:patCis-
coIOSMod>-<:gPatInt>-<:patStrEndColon>:]]></eventFormatRecognizer>
```

Cisco ASA

All Cisco ASA events have the pattern `ASA-severity-id` pattern, for example `ASA-5-12345`.

```
<eventFormatRecognizer><![CDATA[ASA-\d-\d+]]></eventFormatRecognizer>
```

Palo Alto Networks Log Parser

In this case, there is no unique keyword, so the entire message structure from the beginning to a specific point in the log must be considered.

Event

```
<14>May 6 15:51:04 1,2010/05/06
15:51:04,0006C101167,TRAFFIC,start,1,2010/05/06
15:50:58,192.168.28.21,172.16.255.78,::172.16.255.78,172.16.255.78,rule3,,,icm
p,vsys1,untrust,untrust,ethernet1/1,ethernet1/1,syslog-
172.16.20.152,2010/05/06
15:51:04,600,2,0,0,0,0,0x40,icmp,allow,196,196,196,2,2010/05/06
15:50:58,0,any,0
```

```
<eventFormatRecognizer><![CDATA[<:gPatTime>,\w+,
(?:TRAFFIC|THREAT|CONFIG|SYSTEM)]]></eventFormatRecognizer>
```

Pattern Definition Specification

In this section of the parser XML specification, you set the regular expression patterns that that FortiSIEM will iterate through to parse the device logs.

Reusing Pattern Definitions in Multiple Parser Specifications

If there is a pattern definition that you wan to use in multiple parser specification, you need to define it in the file GeneralPatternDefintions.xml, and then refer to it from your s, then it needs to be defined in the file `GeneralPatternDefinitions.xml`. The patterns in that file are named with a g prefix, and can be referenced as shown in this example:

```
<generalPatternDefinitions>
<pattern name="gPatSyslogPRI"><![CDATA[<\d+>]]></pattern>
    <pattern name="gPatMesgBody"><![CDATA[.*]]></pattern>    <pattern name-
e="gPatMonNum"><![CDATA[\d{1,2}]]></pattern>    <pattern name="gPatDay"><![CDATA
[\d{1,2}]]></pattern>    <pattern name="gPatTime"><![CDATA[\d{1,2}:\d{1,2}:\d
{1,2}]]></pattern>    <pattern name="gPatYear"><![CDATA[\d{2,4}]]></pat-
tern></generalPatternDefinitions>
```

Each pattern has a name and the regular expression pattern within the CDATA section. This the basic syntax.

```
<pattern name="patternName"><![CDATA[pattern]]></pattern>
```

This is an example of a pattern definition.

```
<patternDefinitions>    <pattern name="patIpV4Dot"><![CDATA[\d{1,3}.\d{1,3}.\d
{1,3}.\d{1,3}]]></pattern>    <pattern name="patComm"><![CDATA[[^,]+]]></pattern>
  <pattern name="patUpDown"><![CDATA[up|down]]></pattern>    <pattern name-
e="patStrEndColon"><![CDATA[[^:]*]]></pattern></patternDefinitions>
```

You can also write a long pattern definition in multiple lines and indicate their order as shown in this example. The value of the `list` attribute should be `begin` in first line and `end` in last line. If there are more than two lines, the attribute should be set to `continue` for the other lines.

```
<pattern name="patSolarisMod" list="begin"><![CDATA[ssh-
d|login|]]></pattern><pattern name="patSolarisMod" list="continue"><![CDATA[inet-
d|lpstat|]]></pattern><pattern name="patSolarisMod" list="end"><![CDATA
[su|sudo]]></pattern>
```

Parsing Instructions Specification

This section is the heart of the parser, which attempts to recognize patterns in a log message and populate parsed event attributes.

In most cases, parsing involves applying a regular expression to the log, picking up values, and setting them to event attributes. Sometimes the processing is more involved, for example when attributes need to be stored as local variables and compared before populating the event attributes. There are three key components that are used in parsing instructions: Event attributes and variables, inbuilt functions that perform operations on event attributes and variables, and `switch` and `choose` branching constructs for logical operations. Values can be collected from both unstructured and structured strings in log messages.

- Event Attributes and Variables
- Inbuilt Functions
- Branching Constructs
- Collecting Values from Unstructured Strings
- Collecting Fields from Structured Strings

Event Attributes and Variables

The dictionary of event attributes are defined in FortiSIEM database and any member not belonging to that list is considered a local variable. For readability, local variables should begin with an _, although this is not enforced.

Setting an Event Attribute to a Constant

```
<setEventAttribute attr="eventSeverity">1</setEventAttribute>
```

Setting an Event Attribute from Another Variable

The `$` symbol is used to specify the content of a variable. In the example below, attribute `hostMACAddr` gets the value stored in the local variable `_mac`.

```
<setEventAttribute attr="hostMACAddr">$_mac</setEventAttribute>
```

Inbuilt Functions

Combining Two or More Strings to Produce a Final String

This is accomplished by using the `combineMsgId` function. Here `_evIdPrefix` is the prefix, `_evIdSuffix` is the suffix, and the output will be `string1-_evIdPrefix-_evIdSuffix`.

```
<setEventAttribute attr="eventType">combineMsgId("string1", $_evIdPrefix, "-", $_
evIdSuffix)</setEventAttribute>
```

Normalize MAC Address

This is accomplished by using the `normalizeMAC` function. The output will be six groups of two nibbles separated by a colon, for example `AA:BB:CC:DD:EE:FF`.

```
<setEventAttribute attr="hostMACAddr">normalizeMAC($_mac)</setEventAttribute>
```

Compare Interface Security Level

This is accomplished by using the `compIntfSecVal` function. This primarily applies to Cisco ASA and PIX firewalls. The results returned are:

- `LESS` if `srcIntf` has strictly lower security level than `destIntf`
- `GREATER` if `srcIntf` has strictly higher security level than `destIntf`
- `EQUAL` if `srcIntf` and destIntf have identical security levels

```
<setEventAttribute attr="_result">compIntfSecVal($srcIntf, $destInt-
f)</setEventAttribute>
```

Convert Hex Number to Decimal Number

This is accomplished by using the `convertHexStrToInt` function.

```
<setEventAttribute attr="ipConnId">convertHexStrToInt($_ipCon-
nId)</setEventAttribute>
```

Convert TCP/UDP Protocol String to Port Number

This is accomplished by using the `convertStrToIntIpPort` function.

```
<setEventAttribute attr="destIpPort">convertStrToIntIpPort($_dport)</-
setEventAttribute>
```

Convert Protocol String to Number

This is accomplished by the using the `convertStrToIntIpProto` function.

```
<setEventAttribute attr="ipProto">convertStrToIntIpProto($_proStr)</-
setEventAttribute>
```

Convert Decimal IP to String

This is accomplished by using the `converIpDecimalToStr` function.

```
<setEventAttribute attr="srcIpAddr">convertIpDecimalToStr($_srcIpAd-
dr)</setEventAttribute>
```

Convert Host Name to IP

This is accomplished by using the `convertHostNameToIp` function.

```
<setEventAttribute attr="srcIpAddr">convertHostNameToIp($_sad-
dr)</setEventAttribute>
```

Add Two Numbers

This is accomplished by using the `add` function.

```
<setEventAttribute attr="totBytes">add($sentBytes, $recvBytes)</setEventAttribute>
```

Divide Two Numbers

This is accomplished by using the `divide` function.

```
<setEventAttribute attr="memUtil">divide($_usedMem, $_totalMem)</-
setEventAttribute>
```

Scale Function

This is accomplished by using the `scale` function.

```
<setEventAttribute attr="durationMSec">scale($_durationSec, 1000)</-
setEventAttribute>
```

Extract Host from Fully Qualified Domain Name

This is accomplished by using the **extractHostFromFQDN** function. If `_fqdn`` contains a `.` , get the string before the first `.`, otherwise, get the whole string.

```
<setEventAttribute attr="hostName">extractHostFromFQDN($_fqdn)</setEventAttribute>
```

Replace a String Using a Regular Expression

This is accomplished by using the `replaceStringByRegex` function.

```
<setEventAttribute attr="eventType">replaceStringByRegex($_eventType, "\s+", "_
")</setEventAttribute>e.g. _eventType: "Event Type"; eventType: "Event_Type"
```

Replace String in String

This is accomplished by using the `replaceStrInStr` function.

```
<setEventAttribute attr="computer">replaceStrInStr($_computer, "\\", "")</-
setEventAttribute>
```

Resolve DNS Name

This is accomplished by using the `resolveDNSName` function, which converts DNS name to IP address.

```
<setEventAttribute attr="destIpAddr">resolveDNSName($destName)</setEventAttribute>
```

Convert to UNIX Time

This is accomplished by using the `toDateTime` function.

```
<setEventAttribute attr="deviceTime">toDateTime($_mon, $_day, $_year, $_time)</-
setEventAttribute><setEventAttribute attr="deviceTime">toDateTime($_mon, $_day, $_
time)</setEventAttribute>
```

Trim Attribute

This is accomplished by using the **trimAttribute** function. In the example below, it is used to trim the leading and trailing dots in `destName`.

```
<setEventAttribute attr="destName">trimAttribute($destName, ".")</-
setEventAttribute>
```

Branching Constructs

`Choose` Construct

The format is:

```
<choose>    <when test='$AttributeOrVariable1 operator Value1'>          ...
    </when>    <when test='$AttributeOrVariable2 operator Value2'>          ...
    </when>    <otherwise>          ...
    </otherwise></choose>
```

`Switch` Construct

The format is:

```
    <switch>       <case>        ...
       </case>       <case>        ...
       </case>    </switch>
```

Collecting Values from Unstructured Strings

From a string input source, a regex match is applied and variables are set. The variables can be event attributes or local variables. The input will be a local variable or the default raw message variable. The syntax is:

```
 <collectAndSetAttrByRegex src="$inputString ">     <regex><![CDATA[reg-
expattern]]></regex> </collectAndSetAttrByRegex>
```

The regexpattern is specified by a list of variables and sub-patterns embedded within a larger pattern. Each variable and sub-pattern pair are enclosed within <>.

Consider an example in which the local variable `_body` is set to `list 130 permitted eigrp 172.16.34.4(Serial1 ) > 172.16.34.3, 1 packet`. From this sting we need to set the values to local variables and event attributes.

Value	Set To	Type
`130`	`_aclName`	Local Variable
`permitted`	`_action`	Local Variable
`eigrp`	`_proto`	Local Variable
`172.16.34.4`	`srcIpAddr`	Event Attribute
`Serial1`	`srcIntfName`	Event Attribute
`172.16.34.3`	`destIpAddr`	Event Attribute
`1`	`totPkts`	Event Attribute

This is achieved by using this XML. Note that you can use both the `collectAndSetAttrByRegex` and `collectFieldsByRegex` functions to collect values from fields.

```
    <collectAndSetAttrByRegex src="$_body">         <regex><![CDATA[list <_
aclName:gPatStr> <_action:gPatWord> <_proto:gPatWord> <srcIpAddr:gPatIpV4Dot>(<:sr-
cIntfName:gPatWord>) -> <destIpAddr:gPatIpV4Dot>, <totPkts:gPatInt> <:gPatMes-
gBody>]]></regex>    </collectAndSetAttrByRegex>
```

Collecting Fields from Structured Strings

The are usually two types of structured strings in device logs:

- Key=value structured
- Value list structured

In each case, two simpler specialized parsing constructs than are provided

Key=Value Structured Data

Certain logs, such as SNMP traps, are structured as `Key1 = value1 <separator> Key2 = value2,....` These can be parsed using the `collectAndSetAttrByKeyValuePair` XML attribute tag with this syntax.

```
<collectAndSetAttrByKeyValuePair sep='separatorString' src="$inputString">
<attrKeyMap attr="variableOrEventAttribute1" key="key1"/>   <attrKeyMap attr-
r="variableOrEventAttribute2" key="key2"/></collectAndSetAttrByKeyValuePair>
```

When a `key1` match is found, then the entire string following `key1` up to the `separatorString` is parsed out and stored in the attribute `variableOrEventAttribute1`.

As an example, consider this log fragment.

```
_body =
SNMPv2-SMI::enterprises.14823.2.3.1.11.1.1.60 = Hex-STRING: 07 D8 06 0B 13 15 00
00 2D 07 00    SNMPv2-SMI::enterprises.14823.2.3.1.11.1.1.11.0 = Hex-STRING: 00 16
B6 DB 12 22  SNMPv2-SMI::enterprises.14823.2.3.1.11.1.1.12.0 = Hex-STRING: 00 21
55 4D 66 B0  SNMPv2-SMI::enterprises.14823.2.3.1.11.1.1.13.0 = INTEGER: 36
SNMPv2-SMI::enterprises.14823.2.3.1.11.1.1.1.0 = Hex-STRING: 00 1A 1E C0 60 7A
SNMPv2-SMI::enterprises.14823.2.3.1.11.1.1.56.0 = INTEGER: 2    SNMPv2-SMI::en-
terprises.14823.2.3.1.11.1.1.17.0 = STRING: "00:1a:1e:c0:60:7a"
```

The corresponding parser fragment is:

```
  <collectAndSetAttrByKeyValuePair sep='\t\\| SNMP' src="$_body">     <attrKeyMap
attr="srcMACAddr" key="SNMPv2-SMI::enterprises.14823.2.3.1.11.1.1.11.0 = Hex-
STRING: "/>     <attrKeyMap attr="_destMACAddr" key="SNMPv2-SMI::en-
terprises.14823.2.3.1.11.1.1.12.0 = Hex-STRING: "/>    <attrKeyMap attr="wlanSSID"
key="SNMPv2-SMI::enterprises.14823.2.3.1.11.1.1.6.0 = STRING: "/>     <attrKeyMap
attr="wlanRadioId" key="SNMPv2-SMI::enterprises.14823.2.3.1.11.1.1.56.0 = INTEGER:
"/>     <attrKeyMap attr="apMac" key="SNMPv2-SMI::en-
terprises.14823.2.3.1.11.1.1.17.0 = STRING: "/>  </col-
lectAndSetAttrByKeyValuePair>
```

After parsing, the attribute values are set:

Value	Attribute
00 16 B6 DB 12 22	srcMACAddr
00 21 55 4D 66 B0	destMacAddr
2	wlanRadioId
00:1a:1e:c0:60:7a	apMac

Value List Structured Data

Certain application logs, such as those from Microsoft IIS, are structured as a list of values with a separator. These can be parsed using the `collectAndSetAttrByPos` XML attribute tag following this syntax.

```
     <collectAndSetAttrByPos sep='separatorString' src="$inputString">
<attrPosMap attr="variableOrEventAttribute1" pos='offset1'/>         <attrPosMap
attr="variableOrEventAttribute2" pos='offset2'/>     </collectAndSetAttrByPos>
```

When the position `offset1` is encountered, the subsequent values up to the **`separatorString`** is stored in **`variableOrEventAttribute1`**.

As an example, consider this log fragment.

```
_body =
W3SVC1 ADS-PRI 192.168.0.10 GET /Document/ACE/index.htm - 80 - 192.168.20.55
HTTP/1.1 Mozilla/5.0+(Windows;+U;+Windows+NT+5.1;+en-US;+r-
v:1.8.1.11)+Gecko/20071127+Firefox/2.0.0.11 [http://wwwin/Document/] wwwin 200 0 0
5750 445 15
```

The parser fragment is:

```
    <collectAndSetAttrByPos src="$_body" sep='  '>       <attrPosMap attr-
r="srvInstName" pos='1'/>       <attrPosMap attr="destName" pos='2'/>
<attrPosMap attr="relayDevIpAddr" pos='2'>       <attrPosMap attr="destIpAddr" pos-
='3'/>       <attrPosMap attr="httpMethod" pos='4'/>       <attrPosMap attr-
r="uriStem" pos='5'/>       <attrPosMap attr="uriQuery" pos='6'/>
<attrPosMap attr="destIpPort" pos='7'/>       <attrPosMap attr="user" pos='8'/>
   <attrPosMap attr="srcIpAddr" pos='9'/>       <attrPosMap attr="httpVersion" pos-
='10'/>       <attrPosMap attr="httpUserAgent" pos='11'/>       <attrPosMap attr-
r="httpReferrer" pos='13'/>       <attrPosMap attr="httpStatusCode" pos='15'/>
  <attrPosMap attr="httpSubStatusCode" pos='16'/>       <attrPosMap attr-
r="httpWin32Status" pos='17'/>       <attrPosMap attr="recvBytes" pos='18'/>
<attrPosMap attr="sentBytes" pos='19'/>       <attrPosMap attr="durationMSec" pos-
='20'/>    </collectAndSetAttrByPos>
```

For structured strings, techniques in this section are more efficient than in the previous section since, the expression is simpler and ONE tag can be used to parse regardless of the order in which the keys or values appear in the string.

Creating a Custom Parser

Cloning New Parsers

You can clone an existing parser and then use it as the basis for creating a new one. Select the parser you want to clone, and then click **Clone**. Modify the parser as necessary, and then make sure you use the **Up** and **Down** buttons to place it in the list of parsers at the point at which is should be applied.

Creating Custom Parsers for Multi-Tenant Deployments

Custom parsers can only be created from the Super/Global account in Service Provider FortiSIEM deployments.

- Prerequisites
- Procedure

Prerequisites

- You should have examples of the logs that you want to parse
- You should have created any new device/application types, event attribute types, or event types that you want to use in your XML specification
- You should already have written the XML specification for your parser
- You should have prepared a test event that you can use to validate the parser

Parsers Applied in Order

Parsers are applied in the order they are listed in **Admin > Device Support > Parsers**, so it is important to add your custom parser to the list in relation to any other parsers that may be applied to your device logs. If you click **Fix Order**, this will arrange the parsers with system-defined parsers at the top of the list in their original order, and user-defined parsers at the bottom. By sure to click **Apply** to make sure the change in order is picked up by the back-end module.

Procedure

1. Go to **Admin > Device Support > Parsers**.
2. Select a parser that is above the location in the list where you want to add your parser, and then click **New**.
3. Enter a **Name** for the parser.
4. Select a **Device Type** to which the parser should apply.
   If the device type doesn't appear in the menu, you should create a new device type
5. Enter a **Test Event** containing an example of an event that you want to use to validate the parser.
6. Enter the **Parser XML**.
7. Click **Validate**.
   This will validate the XML.
8. Click **Test**.
   This will send the test event to the parser to make sure it is parsed correctly, and will also test the parsers above and below yours in the list to make sure they continue to parse logs correctly.
9. If the XML for your parser validates and the test event is correctly parsed, select **Enable**.
   If you need to continue working on your parser, you can **Save** it without selecting Enable.
10. Click **Save**.

11. Click **Apply** to have the backend module pick up your parser and begin applying it to device logs.
    You should now validate that events are being parsed by creating some activity that will cause a log to be generated, and then run a query against the new device IP address and validate the parsed results.

Deleting or Disabling a Parser

Deleting User-Defined Parsers

You can only delete user-defined parsers, but both system and user-defined parsers can be disabled.

1. Go to **Admin > Device Support > Parsers**.
2. Select the parser you want to delete or disable.
3. Click **Delete** or **Disable**.
4. Click **Yes** to confirm that you want to delete or disable the parser.

Exporting a Custom Parser

To export a parser, you must also export XML files for the device/app types, event attribute types, event types, and then the parser specification file used by your parser.

1. Go to **Admin > Device Support > Device/App Types**.
2. Select the device/application types used in your parser, and then click **Export**.
3. Go to **Admin > Device Support > Event Attribute Types**.
4. Select the event attribute types used in your parser, and then click **Export**.
5. Go to **Admin > Device Support > Event Types**.
6. Select the event types used in your parser, and then click **Export**.
7. Go to **Admin > Device Support > Parsers**.
8. Select the parser specification for your parser, and then click **Export**.

Importing a Custom Parser

Importing a custom parser involves importing four XML files: the XML files containing any device/app types, event attribute types, or event types that you have created for this parser, followed by the parser specification XML file.

1. For each device/app type, event attribute type, or event type XML file that is required for your parser, go to the appropriate tab in **Admin > Device Support**, and then click **Import**.
2. **Browse** to the location of your XML file, and then click **Upload**.
3. Go to **Admin > Device Support > Parsers**, and then click **Import**.
4. **Browse** to the location of your parser specification XML file, and then click **Upload**.
5. Follow the instruction in Creating a Custom Parser to validate your XML and test the parser, and to make sure it appears in the correct position in the list of parsers.

Parser Examples

- Cisco IOS Syslog Parser

Cisco IOS Syslog Parser

The objective is to parse this syslog message:

```
<190>91809: Jan  9 02:38:47.872: %SEC-6-IPACCESSLOGP: list testlog permitted tcp
192.168.20.33(3438) -> 69.147.86.184(80), 1 packet
```

Creating the appropriate parser requires these steps.

- Add Device Type
- Create the Parser Specification and Add Local Patterns
- Define the Format Recognizer
- Parse the Syslog Header
- Parse the Syslog Body
- Final Parser
- Parsed Output

Add Device Type

Create a file `CiscoIOSParser.xml` with this content.

```
<eventParser name="CiscoIOSParser">   <deviceType>      <Vendor>Cisco</Vendor>
 <Model>IOS</Model>      <Version>ANY</Version>   </deviceType></eventParser>
```

Create the Parser Specification and Add Local Patterns

Create the parser XML file with this content, and add the pattern definition `patCiscoIOSMod` for detecting IOS modules such as SEC.

```
<eventParser name="CiscoIOSParser">   <deviceType>       <Vendor>Cisco</Vendor>
  <Model>IOS</Model>      <Version>ANY</Version>   </deviceType>   <pat-
ternDefinitions>      <pattern name="patCiscoIOSMod" list="begin">   <![CDATA
[FW|SEC|SEC_LOGIN|SYS|SNMP|]]></pattern>     <pattern name="patCiscoIOSMod" list-
t="continue"> <![CDATA[LINK|SPANTREE|LINEPROTO|DTP|PARSER|]]></pattern>     <pat-
tern name="patCiscoIOSMod" list="end"><![CDATA[CDP|DHCPD|CONTROLLER|PORT_SECURITY-
SP]]></pattern>     <pattern name="patStrEndColon"><![CDATA[[^:]*]]></pattern>
<pattern name="patComm"><![CDATA[[^,]+]]></pattern>  </pat-
ternDefinitions></eventParser>
```

Define the Format Recognizer

Add this format recognizer for detecting `%SEC-6-IPACCESSLOGP,` which is a signature of Cisco IOS syslog messages.

```
<eventParser name="CiscoIOSParser">   <deviceType>       <Vendor>Cisco</Vendor>
  <Model>IOS</Model>      <Version>ANY</Version>   </deviceType>   <pat-
ternDefinitions>      <pattern name="patCiscoIOSMod" list="begin">   <![CDATA
[FW|SEC|SEC_LOGIN|SYS|SNMP|]]></pattern>     <pattern name="patCiscoIOSMod" list-
t="continue"> <![CDATA[LINK|SPANTREE|LINEPROTO|DTP|PARSER|]]></pattern>
```

```
<pattern name="patCiscoIOSMod" list="end"><![CDATA[CDP|DHCPD|CONTROLLER|PORT_
SECURITY-SP]]></pattern>     <pattern name="patStrEndColon"><![CDATA[[^:]*]]></pat-
tern>     <pattern name="patComm"><![CDATA[[^,]+]]></pattern>   </pat-
ternDefinitions>   <eventFormatRecognizer>     <![CDATA[: %<:patCiscoIOSMod>-
<:gPatInt>-<:patStrEndColon>:]]>   </eventFormatRecognizer></eventParser>
```

Parse the Syslog Header

A syslog message consists of a syslog header, and a body. For better organization, we first parse the syslog header and event type. Subsequent code will include event type specific parsing, which is why event type is extracted in this step. In this example, the header is in boldface.

<190>91809: Jan 9 02:38:47.872: %SEC-6-IPACCESSLOGP: `list testlog permitted tcp 192.168.20.33(3438) -> 69.147.86.184(80), 1 packet`

The XML code for parsing the header does the following:

1. Matches the pattern `<190>91809: Jan 9 02:38:47.872: %SEC-6-IPACCESSLOGP:`
2. Sets the `eventType` attribute to `IOS-SEC- IPACCESSLOGP`.
3. Sets `deviceTime`.
4. Sets event severity (1-7 scale in Cisco IOS, 1=> most severe, to normalized 1-10 scale in FortiSIEM where 10=>most severe)
5. Saves the event `list testlog permitted tcp 192.168.20.33(3438) -> 69.147.86.184 (80), 1 packet` in a temporary variable `_body`.

Note that the patterns `gPatSyslogPRI`, `gPatMon`, `gPatDay`, `gPatTime`, `gPatInt`, `gPatmesgBody` are global patterns that are defined in the GeneralPatternDefinitions.xml file:

```
<generalPatternDefinitions>    <pattern name="gPatSyslogPRI"><![CDATA
[<\d+>]]></pattern>   <pattern name="gPatMesgBody"><![CDATA[.*]]></pattern>  <pat-
tern name="gPatMon"> <![CDATA[Jan|Feb|Mar|Apr|May|Jun|Jul|Aug|Sep|Oct|Nov|Dec|\d
{1,2}]]></pattern>  <pattern name="gPatDay"><![CDATA[\d{1,2}]]></pattern>  <pat-
tern name="gPatTime"><![CDATA[\d{1,2}:\d{1,2}:\d{1,2}]]></pattern> <pattern name-
e="gPatInt"><![CDATA[\d+]]></pattern></generalPatternDefinitions>
```

This parser file XML fragment for parsing the example syslog message looks like this:

```
<parsingInstructions>    <collectFieldsByRegex src="$_rawmsg">      <regex><!
[CDATA[<:gPatSyslogPRI>?<:gPatMon>\s+<:gPatDay>\s+<:gPatTime> %<_evIdPre-
fix:patCiscoIOSMod>-<_severity:gPatInt>-<_evIdSuffix:patStrEndColon>: <_
body:gPatMesgBody>]]></regex>    </collectFieldsByRegex>    <setEventAttribute
attr="eventType">combineMsgId("IOS-", $_evIdPrefix, "-", $_evIdSuffix)</-
setEventAttribute>    <choose>        <when test='$_severity IN "6, 7"'>
 <setEventAttribute attr="eventSeverity">1</setEventAttribute>        </when>
  <when test='$_severity = "1"'>           <setEventAttribute attr-
r="eventSeverity">10</setEventAttribute>        </when>        <when test='$_sever-
ity = "2"'>           <setEventAttribute
attr="eventSeverity">8</setEventAttribute>        </when>        <when test='$_
severity IN "3, 4"'>          <setEventAttribute attr-
r="eventSeverity">5</setEventAttribute>      </when>      <when test='$_severity =
"5"'>          <setEventAttribute attr="eventSeverity">2</setEventAttribute>
</when>    </choose><parsingInstructions>
```

Parse the Syslog Body

The parsing is done on an eventType by eventType basis, since the formats are eventType specific. Parsing the syslog body involves three steps:

1. Parsing the action string. Based on the action staring value (`permit` or `denied`), modify the `eventType` by appending the action string value at the end, and also modify the `eventSeverity` values.
2. Parsing the protocol, source and destination IP, port, and totalPackets.
3. Converting the protocol string to a protocol integer.

```
<choose>   <when test='$eventType IN "IOS-SEC-IPACCESSLOGP, IOS-SEC-IPACCESSLOGDP,
IOS-SEC-IPACCESSLOGRP"'>       <collectAndSetAttrByRegex src="$_body">
<regex><![CDATA[list <_aclName:gPatStr>\s+<_action:gPatWord>\s+<_pro-
to:gPatWord>\s+<srcIpAddr:gPatIpV4Dot>\(<srcIpPort:gPatInt>\)<:gPatMesgBody>-
>\s+<destIpAddr:gPatIpV4Dot>\(<destIpPort:gPatInt>\),\s+<totPkts:gPatInt>
<:gPatMesgBody>]]>                  </regex>       </collectAndSetAttrByRegex>
  <choose>          <when test='$_action = "permitted"'>              <setEventAt-
tribute attr="eventType">combineMsgId("IOS-", $_evIdPrefix, "-", $_evIdSuffix, "-
PERMITTED")</setEventAttribute>        <setEventAttribute attr-
r="eventSeverity">1</setEventAttribute>          </when>          <when test='$_
action = "denied"'>               <setEventAttribute attr="eventType">combineMsgId
("IOS-", $_evIdPrefix, "-", $_evIdSuffix, "-DENIED")</setEventAttribute>
    <setEventAttribute attr="eventSeverity">3</setEventAttribute>          </when>
      </choose>       <setEventAttribute attr="ipProto">convertStrToIntIpProto($_
proto)</setEventAttribute>   </when></choose>
```

Final Parser

```
<eventParser name="CiscoIOSParser">    <deviceType>      <Vendor>Cisco</Vendor>
  <Model>IOS</Model>      <Version>ANY</Version>    </deviceType>    <pat-
ternDefinitions>        <pattern name="patCiscoIOSMod" list="begin">   <![CDATA
[FW|SEC|SEC_LOGIN|SYS|SNMP|]]></pattern>        <pattern name="patCiscoIOSMod"
list="continue"> <![CDATA[LINK|SPANTREE|LINEPROTO|DTP|PARSER|]]></pattern>
<pattern name="patCiscoIOSMod" list="end"><![CDATA[CDP|DHCPD|CONTROLLER|PORT_
SECURITY-SP]]></pattern>        <pattern name="patStrEndColon"><![CDATA
[[^:]*]]></pattern>        <pattern name="patComm"><![CDATA[[^,]+]]></pattern>
</patternDefinitions>    <parsingInstructions>    <!—parse header -->    <col-
lectFieldsByRegex src="$_rawmsg">      <regex><![CDATA[<:gPatSys-
logPRI>?<:gPatMon>\s+<:gPatDay>\s+<:gPatTime> %<_evIdPrefix:patCiscoIOSMod>-<_
severity:gPatInt>-<_evIdSuffix:patStrEndColon>: <_body:gPatMesgBody>]]></regex>
</collectFieldsByRegex>    <setEventAttribute attr="eventType">combineMsgId("IOS-
", $_evIdPrefix, "-", $_evIdSuffix)</setEventAttribute>    <choose>        <when
test='$_severity IN "6, 7"'>            <setEventAttribute attr-
r="eventSeverity">1</setEventAttribute>        </when>        <when test='$_sever-
ity = "1"'>           <setEventAttribute
attr="eventSeverity">10</setEventAttribute>        </when>        <when test='$_
severity = "2"'>           <setEventAttribute attr-
r="eventSeverity">8</setEventAttribute>        </when>        <when test='$_sever-
ity IN "3, 4"'>          <setEventAttribute
attr="eventSeverity">5</setEventAttribute>      </when>      <when test='$_sever-
ity = "5"'>          <setEventAttribute attr="eventSeverity">2</setEventAttribute>
     </when>    </choose>    <!—parse body -->    <choose>      <when
```

```
test='$eventType IN "IOS-SEC-IPACCESSLOGP, IOS-SEC-IPACCESSLOGDP, IOS-SEC-
IPACCESSLOGRP"'>          <collectAndSetAttrByRegex src="$_body">         <regex><!
[CDATA[list <_aclName:gPatStr>\s+<_action:gPatWord>\s+<_pro-
to:gPatWord>\s+<srcIpAddr:gPatIpV4Dot>\(<srcIpPort:gPatInt>\)<:gPatMesgBody>-
>\s+<destIpAddr:gPatIpV4Dot>\(<destIpPort:gPatInt>\),\s+<totPkts:gPatInt>
<:gPatMesgBody>]]>                  </regex>          </collectAndSetAttrByRegex>
        <choose>              <when test='$_action = "permitted"'>
<setEventAttribute attr="eventType">combineMsgId("IOS-", $_evIdPrefix, "-", $_
evIdSuffix, "-PERMITTED")</setEventAttribute>            <setEventAttribute attr-
r="eventSeverity">1</setEventAttribute>             </when>             <when
test='$_action = "denied"'>                 <setEventAttribute attr-
r="eventType">combineMsgId("IOS-", $_evIdPrefix, "-", $_evIdSuffix, "-DENIED")</-
setEventAttribute>                  <setEventAttribute
attr="eventSeverity">3</setEventAttribute>              </when>          </choose>
         <setEventAttribute attr="ipProto">convertStrToIntIpProto($_pro-
to)</setEventAttribute>      </when>    </choose><parsingInstructions>
```

Parsed Output

Input syslog:

```
<190>91809: Jan 9 02:38:47.872: %SEC-6-IPACCESSLOGP: list testlog permitted
tcp 192.168.20.33(3438) -> 69.147.86.184(80), 1 packet
```

Parsed fields:

1. **phRecvTime**: the time at which the event was received by FortiSIEM
2. **phDeviceTime**: Jan 9 02:38:47 2010
3. **eventType**: SEC-IPACCESSLOGP-PERMITTED
4. **eventSeverity**: 3
5. **eventSeverityCategory**: LOW
6. **aclName**: testlog
7. **ipProto**: 6
8. **srcIpAddr**: 192.168.20.33
9. **destIpAddr**: 69.147.86.184
10. **srcIpPort**: 3438
11. **destIpPort**: 80
12. **totPkts**: 1

The master list of event attributes supported by FortiSIEM is here

Custom Performance Monitors

Creating a custom performance monitor involves creating a performance object that specifies the monitoring access protocol to use, maps event attributes available for that protocol to FortiSIEM event attribute types, and then associates those attributes to an event type. You can use system or user-defined device types, event attribute types, and event types when creating the performance object.

- Creating a Custom Performance Monitor
- Monitoring Protocol Configuration Settings
- Mapping Monitoring Protocol Objects to Event Attributes
- Exporting a Custom Performance Monitor
- Importing a Custom Performance Monitor
- Examples of Custom Performance Monitors

Creating a Custom Performance Monitor

You create custom performance monitors by defining the performance object that you want to monitor, including the relationship between the performance object and FortiSIEM events and event attributes, and then associating the performance object to a device type.

Creating Custom Performance Monitors for Enterprise and Multi-Tenant Deployments

In Service Provider FortiSIEM deployments, custom performance performance have to be created by the Super/Global account, and apply to all organizations. In enterprise deployments, custom performance monitors can be created by any user who has access to the **Admin** tab.

- Prerequisites
- Procedure

Prerequisites

- You should review the configuration settings for the monitoring protocols that you will use in your monitor, and be ready to provide the appropriate OIDs, classes, or database table attributes for the access protocol.
- You should have created any new device/application types, event attribute types, or event types that you want to use in your performance monitor
- You should have the IP address and access credentials for a device that you can use to test the monitor

Procedure

Creating the Performance Object and Applying it to a Device

1. Go to **Admin > Device Support > Performance Monitoring**.
2. Click **New**.
3. Enter a **Name** for the performance monitor.
4. For **Type**, select either **System** or **Application**.
5. For **Method**, select the monitoring protocol for the performance monitor.
   See the topics under Monitoring Protocol Configuration Settings for more information about the configuration settings for each type of monitoring protocol.
6. Click **New** next to **List of Attributes**, and create the mapping between the performance object and FortiSIEM event attributes.
   Note that the Method you select will determine the name of this mapping and the configuration options that are available. See Mapping Monitoring Protocol Objects to Event Attributes for more information.
7. Select the **Event Type** that will be monitored.
8. Enter the **Polling Frequency** for the monitor.
9. Enter a **Description**.
10. Click **Save**.
11. In **Admin > Device Support > Performance Monitoring**, under **Enter Device Type to Performance Object Mapping**, click **New**.
12. Enter a **Name** for the mapping.
13. In the top pane of the dialog, select the **Device Type** to which you want to apply the monitor.
    Whenever a device belonging to the selected device type is discovered, FortiSIEM will attempt to apply the performance monitor to it.

14. In the bottom pane of the dialog, select the custom performance monitor.
15. Click **Save**.

Testing the Performance Monitor

1. Go to **Admin > Device Support > Performance Monitoring.**
2. Select the performance monitor.
3. Click **Test**.
4. For **IP**, enter the IP address of the device that you want to use to test the monitor.

   Testing for Multi-Tenant Deployments

   If you have a Service Provider FortiSIEM, select the Supervisor or Collector where the device is monitored.

5. Click **Test**. If the test succeeds, click **Close**, and then click **Apply** to register the new monitor with the backend module.

After you have successfully tested and applied the performance monitor, you should initiate discovery of the device that it will monitor, and then make sure that the new monitor is enabled as described in Managing Monitoring of System and Application Metrics for Devices.

Monitoring Protocol Configuration Settings

These topics describe the configuration settings for monitoring protocols such as SNMP, WMI, and JDBC that are used for creating custom performance monitors.

- JDBC Configuration Settings
- JMX Configuration Settings
- SNMP Configuration Settings for Custom Performance Monitors
- WMI Configuration Settings for Custom Performance Monitors

JDBC Configuration Settings

When configuring JDBC as the access protocol for a custom performance monitor, use these settings. You may also want to review the topic Custom JDBC Performance Monitor for a Custom Table as example of how to set up a custom performance monitor using JDBC.

Field	Setting/Notes
Method	JDBC
Database Type	Select the type of database to connect to
SQL Query	The SQL Query to execute when connecting
List of Columns	This creates the mapping between columns in the database and FortiSIEM event attributes. See Mapping Monitoring Protocol Objects to Event Attributes for more information.
Where Clauses	This indicates whether the database table being queried has a fixed set of rows, or whether it is growing over time. An example of this would be a table containing logs, in which case FortiSIEM would keep track of the last entry and only pull the new ones. There are three options here: 1. There is a fixed set of rows and all rows are needed. Leave all options cleared. 2. There is a fixed set of rows and a fixed number of rows are needed. Select **Fixed Records** and enter the number of required rows. 3. The table is growing and only new values are needed. Select **Retrieve all new values since last retrieve time of column**, and enter the name of the column that represents time in the database. FortiSIEM will keep track of the largest value in this column and only pull entries greater than that value during the next polling interval.

JMX Configuration Settings

When configuring JMX as the monitoring protocol for a custom performance monitor, use these settings. You may also want to review the topic Custom JMX Monitor for IBM Websphere as an example of creating a custom JMX performance monitor.

Field	Setting/Notes
Method	JMX
MBean	Enter the MBean interface that you want to monitor, or click the downward arrow to browse the JMX tree and select it. Note that the option you select here will determine the objects that are available when you select an **Object Attribute** for the **List of Attributes**. See the next section in this topic for information on how to find

Identifying MBean Names and Attributes for Custom Applications

This section discusses how to get MBean names and attributes for custom J2EE based applications.

1. Launch JConsole on your workstation and connect to the application.
2. Select the **MBeans** tab.
3. Browse to the application you want to monitor, and select it.
4. In the right pane you will see the `MBeanInfo`. Note the `ObjectName`, while the attributes for the application will be listed in the tree view.

SNMP Configuration Settings for Custom Performance Monitors

When configuring SNMP as the access protocol for a custom performance monitor, use these settings. You may also want to review the topics Custom SNMP Monitor for D-Link Interface Network Statistics and Custom SNMP Monitor for D-Link HostName and SysUpTime as example of how to set up a custom performance monitor using SNMP.

Field	Settings/Notes
Method	SNMP
Parent OID	The parent Object Identifier (OID) is used to optimize the number of SNMP GETs required for pulling the various individual OIDs. You can enter this directly, or click the downward arrow to select it from an MIB file. Several different MIB files are available to select from, s ee Importing OID Definitions from a MIB File for more information.
Parent ID is table	Select **is table** if the OIDs you want to monitor are in a table with at least one row. An example would be interface metrics, such as `ifInOctets` and `ifOutOctets`, since there is an interface metric for each interface.
List of OIDs	The OIDs you want to monitor mapped to FortiSIEM event attributes. The selection you make for **Parent OID** determines the options available in the **OID** menu when you select **New**.

Importing OID Definitions from a MIB File

Many devices include MIB files that you can then use to create a custom performance monitor for the device. This involves creating a configuration file based on information in the MIB file, using that file as input for the `mib2xml` executable, and then placing the resulting output file in the `/data/mibXml` directory of your Supervisor. Once placed in this directory, you can select the file from the **MIB File List** menu to select the parent OID, which will then also affect which OIDs you can select for the OID to event attribute mapping.

Procedure

1. Collect the device OID files you want to use and place them in a directory where the mib2XML
2. Create the input config file with these fields, and name it with the `.cfg` file designation.
   See the attached alcatel.cfg file for an example.

Field	Description
group	This is the number of MIB file group. MIB files need to be analyzed as a group because of cross-references within them. The group attribute specifies an ID for each group and needs to be unique for every group.
mibFile	The name of the MIB file being analyzed. There can be multiple entries. Be sure to specify the path to the MIB files.
vendor	The name of the device vendor for the MIB file
model	The model name or number for the device
evtPrefix	As SNMP trap notification definitions in the MIB file are parsed, an event file is generated for each SNMP trap. This field specifies the event type prefix.
enterpriseId	The enterprise ID number for this vendor, which is used for generating the SNMP trap parser

3. Run `mib2XML <filename>.cfg`.
4. Move the resulting `.mib.xml` file to the `/data/mibXml` directory of your Supervisor.

Example

In this example, a set of MIB files from an Alcatel 7x50 device are used to generate the XML output file.

1. Sample MIB files:
   TIMETRA-CHASSIS-MIB.mib
   TIMETRA-GLOBAL-MIB.mib
   TIMETRA-SYSTEM-MIB.mib
   TIMETRA-TC-MIB.mib
2. Information in these files, and the paths to them, are then used to create this config file.
   alcatel.cfg
3. Running `mib2xml alcatel.cfg` generates both an output and an mib2XML file.
   alcatel.out
   TIMETRA-TC-MIB.mib.xml

WMI Configuration Settings for Custom Performance Monitors

When configuring WMI as the monitoring protocol for a custom performance monitor, use these settings. You may also want to review the topic Custom WMI Monitor for Windows Domain and Physical Registry as example of how to set up a custom performance monitor using WMI.

Field	Settings
Method	WMI
Parent Class	WMI metrics are defined in the form of a parent class having multiple attributes. For example, the parent class `Win32_ComputerSystem` has the attributes `Domain` and `TotalPhysicalMemory`.
Is Table	If the parent WMI class is a table with one or more rows, select this option.

Mapping Monitoring Protocol Objects to Event Attributes

When you select a monitoring protocol for your custom performance monitor, you must also establish the relationship between the objects used by that protocol and event attributes in FortiSIEM. For example, creating a performance monitor that uses SNMP to monitor a device requires that you create a mapping between the SNMP OIDs that you want to monitor, and set of event attributes. This topic describes the configuration settings that you will use to create these object-to-event attribute relationships.

Procedure

1. When creating your custom performance monitor, after you have selected the **Method**, click **New** next to **List of Attributes**.
   Depending on the monitoring protocol that you select, this table may be named **List of OIDs** (SNMP), or **List of Columns** (JDBC).
2. In the first field, enter or select the monitoring protocol object that you want to map to FortiSIEM event attribute. Your options depend on the monitoring protocol you selected for Method.

Monitoring Protocol	Field Name	Settings/Notes
SNMP	**OID**	Select an MIB file from the **MIB File List**, and then select the OID that you want to create the mapping for.
WMI	**Attribute**	Enter an attribute of the WMI class you entered for **Parent Class**.
JMX	**Object Attribute**	The **MBean** you select determines the attributes you can select. You will also have to enter a **Name** and **Private Key** for the MBean attribute.
JDBC	**Column Name**	Enter the name of the column in the **SQL Query** that you are using to monitor the database.

3. Select the **Format** for the object attribute.
   Your options will depend on the monitoring protocol you selected for Method.
4. For **Type**, select **Raw Value** or **Counter**.
5. For **Event Attribute**, select the FortiSIEM event attribute that the monitoring protocol object should map to. If you need to create a new event attribute, see Creating Event Attribute Types.
6. Create any **Transforms** of the values returned for the monitoring protocol object.
   See the next section for more information how to configure transforms.
7. Click **Save** when you are done creating the mappings, and then complete the configuration of your custom performance monitor.

Creating Transforms

You can use a transform to convert the value returned for your monitoring project object into a more physically meaningful or usable metric. You an create multiple transforms, and they will be evaluated in the order shown in the table. Multiple transforms can be selected – they are evaluated in sequential order as shown in the display table

1. Next to **Transforms**, click **New**.
2. For **Type**, select **System** or **Custom**.
3. For **Formula**, either select a system-defined transformation formula from the menu if you selected **System** for **Type**, or enter a formula if you selected **Custom**.
4. Click **Save**.

Exporting a Custom Performance Monitor

To export a parser, you must also export XML files for the device/app types, event attribute types, event types, and then the monitor.

1. Go to **Admin > Device Support > Device/App Types**.
2. Select the device/application types used in your monitor, and then click **Export**.
3. Go to **Admin > Device Support > Event Attribute Types**.
4. Select the event attribute types used in your monitor, and then click **Export**.
5. Go to **Admin > Device Support > Event Types**.
6. Select the event types used in your monitor, and then click **Export**.
7. Go to **Admin > Device Support > Performance Monitoring**.
8. Select the monitor, and then click **Export**.

Importing a Custom Performance Monitor

Importing a custom performance monitor involves importing four XML files: the XML files containing any device/app types, event attribute types, or event types that you have created for this parser, followed by the custom performance monitor file.

1. For each device/app type, event attribute type, or event type XML file that is required for your monitor, go to the appropriate tab in **Admin > Device Support**, and then click **Import**.
2. **Browse** to the location of your XML file, and then click **Upload**.
3. Go to **Admin > Device Support > Performance Monitors**, and then click **Import**.
4. **Browse** to the location of your performance monitor file, and then click **Upload**.
5. Follow the instructions in Creating a Custom Performance Monitor to test and apply your performance monitor.

Examples of Custom Performance Monitors

- Custom JDBC Performance Monitor for a Custom Table
- Custom JMX Monitor for IBM Websphere
- Custom SNMP Monitor for D-Link HostName and SysUpTime
- Custom SNMP Monitor for D-Link Interface Network Statistics
- Custom WMI Monitor for Windows Domain and Physical Registry

Custom JDBC Performance Monitor for a Custom Table

Planning

Examining the Table Structure

For this example, consider two custom Oracle tables that you want to monitor.

1. A table called `HEALTH_STATIC_DEMO` that does not have time stamp as a column. The table does not grow with time, and the `HEALTH` column is updated by the application.

```
create table HEALTH_STATIC_DEMO
{
   ID        VARCHAR2 (200) not null,
   HOST_NAME NVARCHAR2 (200) not null,
   HEALTH    NVARCHAR2 (50)
}
```

2. A table called `HEALTH_DYNAMIC_DEMO` that has a time-stamp in the column `create_time`. Only records with a more recent time-stamp than previous ones have to be pulled in, and every time a new record is written, it includes a time stamp.

```
create table HEALTH_DYNAMIC_DEMO
{
   ID          VARCHAR2 (200) not null,
   HOST_NAME   NVARCHAR2 (200) not null,
   HEALTH      NVARCHAR2 (50),
   CREATE_TIME DATE not null
}
```

Creating New Device Types, Event Attribute Types, and Event Types

In this case, you only need to create two new event types to handle the contents of the two tables.

Naming Custom Event Types

All custom event types must begin with the prefix `P H_DEV_MON_CUST_ .`

Event Types

Name	Device Type	
`PH_DEV_MON_CUST_JDBC_PERFORMANCE_STATIC`	Generic	Low
`PH_DEV_MON_CUST_JDBC_PERFORMANCE_DYNAMIC`	Generic	Low

Adding New JDBC Performance Objects

Each table requires its own performance object for monitoring.

Performance Object Configuration for Static Table HEALTH_STATIC_DEMO

Field	Setting
Name	`jdbc_static_perfObj`
Type	Application
Method	JDBC
Database Type	Oracle Database Server
SQL Query	`select * from health_static_demo`
List of Columns	**Column Name** / **Name** / **Format** / **Event Attribute** host_name / STRING / hostName health / STRING / health
Where Clauses	Not applicable, since the table doesn't grow over time
Event Type	`PH_DEV_MON_CUST_JDBC_PERFORMANCE_STATIC`
Polling Frequency	180 seconds

Performance Object Configuration for Dynamic Table HEALTH_DYNAMIC_DEMO

Field	Setting
Name	`jdbc_dynamic_perfObj`
Type	Application
Method	JDBC
Database Type	Oracle Database Server
SQL Query	`select * from health_dynamic_demo`
List of Columns	**Column Name** / **Name** / **Format** / **Event Attribute** host_name / STRING / hostName cpu_util / DOUBLE / cpuUtil mem_util / DOUBLE / memUtil create_time / STRING / createTime
Where Clauses	retrieve all new values since last retrieve time of column create_time
Event Type	`PH_DEV_MON_CUST_JDBC_PERFORMANCE_STATIC`
Polling Frequency	180 seconds

Associating Device Types to Performance Objects

In this example, the Oracle database runs on Microsoft Windows, so you would need to associate Microsoft Windows device types to the two performance objects. Because the discovered device type has to exactly match one of device types in this association in order for the discovery module to initiate monitoring, you would need to add other device types, such as Linux, if you also wanted to monitor Oracle databases over JDBC on those devices.

Edit Device to Performance Object

Field	Settings
Name	windows_oracle_perf_association
Device Types	• Microsoft Windows • Microsoft Windows 7 • Microsoft Windows 98 • Microsoft Windows ME • Microsoft Windows NT • Microsoft Windows Server 2000 • Microsoft Windows Server 2003 • Microsoft Windows Server 2008 • Microsoft Windows Vista • Microsoft Windows XP
Perf Objects	• jdbc_static_perfObj(JDBC) - Default Interval:3mins • jdbc_dynamic_perfObj(JDBC) - Default Interval:3mins

Testing the Performance Monitor

Before testing the monitor, make sure you have defined the access credentials for the database server, created the IP address to credentials mapping, and tested connectivity to the server.

1. Go to **Admin > Device Support > Performance Monitoring**.
2. Select one of the performance monitors you created, and then click **Test**.
3. For **IP**, enter the address of the Oracle database server, and select either the Supervisor or Collector node that will retrieve the information for this monitor.
4. Click **Test** .
   You should see `succeed` under **Result** , and a parsed event attributes in the test result pane.
5. When the test succeeds, click **Close**, and then click **Apply** to register the new monitor with the backend module.

Enabling the Performance Monitor

1. Discover or re-discover the device you want to monitor.
2. Once the device is successfully discovered, make sure that the monitor is enabled and pulling metrics.

Writing Queries for the Performance Metrics

You can now use a simple query to make sure that that the metrics are pulled correctly. The search results should display the metrics for the event attributes you defined.

1. Create a structured historical search, and in the **Filter Criteria**, enter `Event Type = "PH_DEV_MON_CUST_JDBC_PERFORMANCE_STATIC"; Group by:` [None]
   This should show the entries in the `HEALTH_STATIC_DEMO` table

2. Create a structured historical search, and in the **Filter Criteria**, enter `Event Type = "PH_DEV_MON_CUST_JDBC_PERFORMANCE_SDynamic"; Group by:` [None]
   This should show the entries in the `HEALTH_DYNAMIC_DEMO` table .

Custom SNMP Monitor for D-Link Interface Network Statistics

This example shows how to create a custom performance monitor for network interface statistics for D-link switches. In this case, the result is a table, with one set of metrics for each interface.

Planning

Matching SNMP OIDs to FortiSIEM Event Attribute Types

If you run the command `snmpwalk -v 1 -c <community> <ip> .1.3.6.1.2.1.2.2.1` against the D-Link switch, you should see an output similar to this:

```
IF-MIB::ifIndex.1 = INTEGER: 1
IF-MIB::ifIndex.2 = INTEGER: 2
IF-MIB::ifIndex.3 = INTEGER: 3
IF-MIB::ifIndex.4 = INTEGER: 4
IF-MIB::ifIndex.5 = INTEGER: 5
...
```

To get the interface index, you would run `snmpwalk -v 1 -c <community> <ip> .1.3.6.1.2.1.2.2.1.1`:

```
IF-MIB::ifIndex.1 = INTEGER: 1
IF-MIB::ifIndex.2 = INTEGER: 2
IF-MIB::ifIndex.3 = INTEGER: 3
IF-MIB::ifIndex.4 = INTEGER: 4
IF-MIB::ifIndex.5 = INTEGER: 5
...
```

To get interface queue length (the `outQLen` event attribute in FortiSIEM), you would run `snmpwalk -v 1 -c <community> <ip> .1.3.6.1.2.1.2.2.1.21`:

```
IF-MIB::ifOutQLen.1 = Gauge32: 0
IF-MIB::ifOutQLen.2 = Gauge32: 0
IF-MIB::ifOutQLen.3 = Gauge32: 0
IF-MIB::ifOutQLen.4 = Gauge32: 0
IF-MIB::ifOutQLen.5 = Gauge32: 0
...
```

To get interface speed, you would run `snmpwalk -v 1 -c <community> <ip> .1.3.6.1.2.1.2.2.1.5`:

```
IF-MIB::ifSpeed.1 = Gauge32: 1000000000
IF-MIB::ifSpeed.2 = Gauge32: 1000000000
IF-MIB::ifSpeed.3 = Gauge32: 1000000000
IF-MIB::ifSpeed.4 = Gauge32: 1000000000
```

```
IF-MIB::ifSpeed.5 = Gauge32: 1000000000
...
```

To get received bytes (the `recvBitsPerSec` event attribute in FortiSIEM), you would run `snmpwalk -v 1 -c <community> <ip> .1.3.6.1.2.1.2.2.1.10`:

```
IF-MIB::ifInOctets.1 = Counter32: 0
IF-MIB::ifInOctets.2 = Counter32: 1247940872
IF-MIB::ifInOctets.3 = Counter32: 0
IF-MIB::ifInOctets.4 = Counter32: 0
IF-MIB::ifInOctets.5 = Counter32: 0
...
```

Finall,y to get sent bytes (the `sentBitsPerSec` event attribute in FortiSIEM), you would run `snmpwalk -v 1 -c <community> <ip> .1.3.6.1.2.1.2.2.1.16`:

```
IF-MIB::ifOutOctets.1 = Counter32: 0
IF-MIB::ifOutOctets.2 = Counter32: 1271371281
IF-MIB::ifOutOctets.3 = Counter32: 0
IF-MIB::ifOutOctets.4 = Counter32: 0
IF-MIB::ifOutOctets.5 = Counter32: 0
...
```

From these outputs you can see that if you want to create a performance monitor for D-Link switch uptime, you need to:

1. Create a new device type, since D-Link switches are not supported in this release.
2. Create an event type, `PH_DEV_MON_CUST_DLINK_INTF_STAT`, that will contain the event attribute types `outQLen` , `recvBitsPerSec`, and `sentBitsPerSec`, which are already part of the FortiSIEM event attribute library, and `hostNameSnmpIndx` and `intfSpeed`, which you need to create.
3. Create the mapping between the SNMP OIDs and the event attributes:
    1. OID `.1.3.6.1.2.1.2.2.1.1` and `hostNameSnmpIndx`
    2. OID `.1.3.6.1.2.1.2.2.1.5` and `intfSpeed`
    3. OID `.1.3.6.1.2.1.2.2.1.21` and `outQLen`
    4. OID `.1.3.6.1.2.1.2.2.1.10` and `recvBitsPerSec`
    5. OID .1.3.6.1.2.1.2.2.1.16 and `sentBitsPerSec`

Creating New Device Types, Event Attributes, and Event Types

Device Type

Create a new device type with these attributes:

Field	Setting
Vendor	D-Link
Model	DGS
Version	Any

Device/App Group	**Devices > Network Devices > Router Switch**
Biz Service Group	<no selection>
Description	D-Link Switch

Event Attribute Types

Create these event attribute types:

Name	Display Name	Value Type	Display Format Type
hostSnmpIndex	Host Interface SNMP Index	INT64	<left blank>
intfSpeed	Interface Speed in bits/sec	INT64	<left blank>

Event Types

Naming Custom Event Types

All custom event types must begin with the prefix `P H_DEV_MON_CUST_` .

Create this event type:

Name	Device Type	Severity
`PH_DEV_MON_CUST_INTF_STAT`	D-Link DGS	Low

Adding the D-Link SNMP Performance Object

In this case, you will create one performance object that will map the SNMP OIDs to the FortiSIEM event attribute types, and then associate them with the `PH_DEV_MON_CUST_INTF_STAT` event type. When you create the `recvBitsPerSec` and `sentBitsPerSec` mapping you will also add a sequential transform to convert the cumulative metric to a rate, and then convert bytes per second to bits per second. .

Performance Object Configuration for Event Type PH_DEV_MON_CUST_INTF_STAT

Field	Setting
Name	D-LinkIntStat
Type	System
Method	SNMP
Parent OID	.1.3.6.1.2.1.2.2.1
Parent OID is Table	Selected
List of OIDs	**Object Attribute** / **Name** / **Format** / **Type** / **Event Attribute** .1.3.6.1.1.2.1.2.2.1.1 / IntfIndex / INTEGER / RawValue / `hostSnmpIndex` .1.3.6.1.1.2.1.1.2.1.5 / intfSpeed / Gauge32 / RawValue / `intfSpeed` .1.3.6.1.1.2.1.1.2.1.10 / recvBitsPerSec / Counter32 / Counter / recvBitsPerSec .1.3.6.1.1.2.1.1.2.1.16 / sentBitsPerSect / Counter32 / Counter / sentBitsPerSect .1.3.6.1.1.2.1.1.2.1.21 / outInftQ / Gauge32 / RawValue / OutQLen
Event Type	PH_DEV_MON_CUST_INTF_STAT
Polling Frequency	60 seconds

Transform Formula for recvBitsPerSec and sentBitsPerSec Event Attributes

Type	Formula
system	toRate
system	BytesPerSecToBitsPerSec

Associating Device Types to Performance Objects

In this case you would only need to make one association with the D-Link DGS device you created.

Field	Settings
Name	`D-LinkPerfObj`
Device Types	• D-Link DGS
Perf Objects	• D-LinkIntfStat(SNMP) - Default Interval:1mins

Testing the Performance Monitor

Before testing the monitor, make sure you have defined the access credentials for the D-Link device, created the IP address to credentials mapping, and tested connectivity.

1. Go to **Admin > Device Support > Performance Monitoring**.
2. Select the performance monitor you created, and then click **Test**.
3. For **IP**, enter the address of the device, and select either the Supervisor or Collector node that will retrieve the information for this monitor.
4. Click **Test**.
   You should see `succeed` under **Result**, and the parsed event attributes in the test result pane.
5. When the test succeeds, click **Close**, and then click **Apply** to register the new monitor with the backend module.

Enabling the Performance Monitor

1. Discover or re-discover the device you want to monitor.
2. Once the device is successfully discovered, make sure that the monitor is enabled and pulling metrics.

Writing Queries for the Performance Metrics

You can now use a simple query to make sure that that the metrics are pulled correctly. The search results should display the metrics for the event attributes you defined.

Create a structured historical search with these settings:

Filter Criteria	Display Columns	Time	For

			Organizations
Structured `Reporting IP IN <IP Range> AND Event Type =" PH_DEV_MON_CUST_INTF_ STAT"; Group by: Host Name, Host Interface`	Host Name,Host Interface SNMP Index,MAX(Out Intf Queue), AVG(Intf Speed), AVG(Sent Bit Rate), AVG (Received Bit Rate)	Last 10 Minutes	All

Custom JMX Monitor for IBM Websphere

This example illustrates how to write a custom performance monitor for retrieving IBM Websphere thread, heap memory, and non-heap memory metrics.

Planning

Creating New Device Types, Event Attribute Types, and Event Types

In this case, the IBM Websphere device type is already supported by FortiSIEM, but you need to create new event attributes and event types for the metrics you want to retrieve.

Event Attribute Types

Name	Display Name	Value Type	Display Format Type
websphere_heapPCT	WebSphere HeapPct	INT64	
websphere_numThreads	WebSphere NumThreads	INT64	
websphere_maxThreads	WebSphere MaxThreads	INT64	
websphere_threadPct	WebSphere ThreadPct	INT64	
websphere_numClass	WebSphere NumClass	INT64	
websphere_heapUsed	WebSphere HeapUsed	INT64	Bytes
websphere_heapMax	WebSphere HeapMax	INT64	Bytes
websphere_heapCommitted	WebSphere HeapCommitted	INT64	Bytes
websphere_nonHeapUsed	WebSphere NonHeapUsed	INT64	Bytes
websphere_nonHeapMax	WebSphere NonHeapMax	INT64	Bytes
websphere_nonHeapCommitted	WebSphere NonHeapCommitted	INT64	Bytes

Event Types

Naming Custom Event Types

All custom event types must begin with the prefix `P H_DEV_MON_CUST_` .

Name	Device Type	Severity
PH_DEV_MON_CUST_WEBSPHERE_HEAPMEMORY	IBM WebSphere App Server	Low
PH_DEV_MON_CUST_WEBSPHERE_NON_HEAPMEMORY	IBM WebSphere App Server	Low
PH_DEV_MON_CUST_WEBSPHERE_THREAD	IBM WebSphere App Server	Low

Adding New IBM WebSphere Performance Objects

Each of the event types requires creating a performance object for monitoring.

Performance Object Configuration for Event Type PH_DEV_MON_CUST_WEBSPHERE_HEAPMEMORY

<table>
<tr><th>Field</th><th>Setting</th></tr>
<tr><td>Name</td><td>websphere_heapMemory_perfObj</td></tr>
<tr><td>Type</td><td>Application</td></tr>
<tr><td>Method</td><td>JMX</td></tr>
<tr><td>MBean</td><td><code>java.lang:type=Memory</code></td></tr>
<tr><td>List of Attributes</td><td>
<table>
<tr><th>Object Attribute</th><th>Private Key</th><th>Name</th><th>Format</th><th>Event Attribute</th></tr>
<tr><td>HeapMemoryUsage</td><td>committed</td><td>committed</td><td>Long</td><td><code>websphere_heapCommitted</code></td></tr>
<tr><td>HeapMemoryUsage</td><td>used</td><td>used</td><td>Long</td><td><code>websphere_heapUsed</code></td></tr>
<tr><td>HeapMemoryUsage</td><td>max</td><td>max</td><td>Long</td><td><code>websphere_heapMax</code></td></tr>
<tr><td></td><td></td><td></td><td>Long</td><td><code>websphere_heapPCT</code></td></tr>
</table>
</td></tr>
<tr><td>Event Type</td><td><code>PH_DEV_MON_CUST_WEBSPHERE_HEAPMEMORY</code></td></tr>
<tr><td>Polling Frequency</td><td>180 seconds</td></tr>
</table>

Performance Object Configuration for Event Type PH_DEV_MON_CUST_WEBSPHERE_THREAD

For the `webSphere_threadPct` **Event Attribute**, you will enter a transform as shown in the second table.

<table>
<tr><th>Field</th><th>Setting</th></tr>
<tr><td>Name</td><td><code>websphere_thread_perfObj</code></td></tr>
<tr><td>Type</td><td>Application</td></tr>
<tr><td>Method</td><td>JMX</td></tr>
<tr><td>MBean</td><td><code>java.lang:type=Threading</code></td></tr>
<tr><td>List of Attributes</td><td>
<table>
<tr><th>Object Attribute</th><th>Private Key</th><th>Name</th><th>Format</th><th>Event Attribute</th></tr>
<tr><td>ThreadCount</td><td></td><td>ThreadCount</td><td>Long</td><td><code>websphere_numThreads</code></td></tr>
<tr><td>PeakThreadCount</td><td></td><td>PeakThreadCount</td><td>Long</td><td><code>websphere_maxThreads</code></td></tr>
<tr><td></td><td></td><td></td><td>Long</td><td><code>websphere_threadPCT</code></td></tr>
</table>
</td></tr>
<tr><td>Event Type</td><td><code>PH_DEV_MON_CUST_WEBSPHERE_THREAD</code></td></tr>
<tr><td>Polling Frequency</td><td>180 seconds</td></tr>
</table>

Transform Formula for websphere_threadPCT Event Attribute

Click **New** next to **Transforms** in the dialog to enter the formula.

Field	Settings
Object Attribute	<blank>
Name	<blank>
Private Key	<blank>
Format	Long
Event Attribute	websphere_threadPct
Transforms	**Type** **Formula** custom ThreadCount*100/PeakThreadcount

Performance Object Configuration for Event Type PH_DEV_MON_CUST_WEBSPHERE_NON_HEAPMEMORY

Field	Setting				
Name	`websphere_nonHeapMemory_perfObj`				
Type	Application				
Method	JMX				
MBean	`java.lang:type=Memory`				
List of Attributes	**Object Attribute**	**Private Key**	**Name**	**Format**	**Event Attribute**
	NonHeapMemoryUsage	used		Long	`websphere_nonHeapUsed`
	NonHeapMemoryUsage	committed		Long	`websphere_nonHeapCommitted`
	NonHeapMemoryUsage	max		Long	`websphere_nonHeapMax`
Event Type	`PH_DEV_MON_CUST_WEBSPHERE_NON_HEAPMEMORY`				
Polling Frequency	180 seconds				

Associating Device Types to Performance Objects

In this example, IBM WebSphere runs on Microsoft Windows, so you would need to associate Microsoft Windows device types to the three performance objects. Because the discovered device type has to exactly match one of device types in this association in order for the discovery module to initiate these monitors, you would need to add other device types, such as Linux, if you also wanted to monitor IBM Websphere over JMX on those devices.

Edit Device to Performance Object

Field	Settings
Name	windows_oracle_perf_association
Device Types	• Microsoft Windows • Microsoft Windows 7 • Microsoft Windows 98 • Microsoft Windows ME • Microsoft Windows NT • Microsoft Windows Server 2000 • Microsoft Windows Server 2003 • Microsoft Windows Server 2008 • Microsoft Windows Vista • Microsoft Windows XP
Perf Objects	• websphere_thread_perfObj(JMX) - Default Interval:3mins • websphere_thread_perfObj(JMX) - Default Interval:3mins • websphere_nonHeapMemory_perfObj (JMX) - Default Interval:3mins

Testing the Performance Monitor

Before testing the monitor, make sure you have defined the access credentials for the server, created the IP address to credentials mapping, and tested connectivity.

1. Go to **Admin > Device Support > Performance Monitoring**.
2. Select one of the performance monitors you created, and then click **Test**.
3. For **IP**, enter the address of the Oracle database server, and select either the Supervisor or Collector node that will retrieve the information for this monitor.
4. Click **Test** .
   You should see `succeed` under **Result** , and the parsed event attributes in the test result pane.
5. When the test succeeds, click **Close**, and then click **Apply** to register the new monitor with the backend module.

Enabling the Performance Monitor

1. Discover or re-discover the device you want to monitor.
2. Once the device is successfully discovered, make sure that the monitor is enabled and pulling metrics.

Writing Queries for the Performance Metrics

You can now use a simple query to make sure that that the metrics are pulled correctly. The search results should display the metrics for the event attributes you defined.

Create a structured historical search with these settings:

Filter Criteria	Display Columns	Time	For Organizations
Structured `Reporting IP IN <IP Range> AND Event Type CONTAIN "ph_dev_mon_cust_web"; Group by: [None]`	Event Receive Time,Reporting IP, Event	Last 60 Minutes	All

Custom SNMP Monitor for D-Link HostName and SysUpTime

Although D-link switches and routers are not supported in this release of FortiSIEM, you can still use the custom monitor feature to create a system uptime event that will collect basic performance metrics like `hostName` and `SysUpTime`.

Planning

Mapping SNMP OIDs to FortiSIEM Event Attribute Types

If you run the command `snmpwalk -v 1 -c <community> <ip> .1.3.6.1.2.1.1` against the D-Link switch, you should see an output similar to this:

```
SNMPv2-MIB::sysDescr.0 = STRING: DGS-1210-48          2.00.011
SNMPv2-MIB::sysObjectID.0 = OID: SNMPv2-SMI::enterprises.171.10.76.11
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (157556100) 18 days, 5:39:21.00
SNMPv2-MIB::sysContact.0 = STRING:
SNMPv2-MIB::sysName.0 = STRING: SJ-Test-Lab-D-Link
SNMPv2-MIB::sysLocation.0 = STRING: San Jose
SNMPv2-MIB::sysServices.0 = INTEGER: 72
SNMPv2-MIB::sysORLastChange.0 = Timeticks: (157555949) 18 days, 5:39:19.49
```

To get sysU`ptime`, you would run `snmpwalk -v 1 -c <community> <ip> .1.3.6.1.2.1.1.3`:

```
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (157577770) 18 days, 5:42:57.70
```

To get `hostname`, you run `snmpwalk -v 1 -c <community> <ip> .1.3.6.1.2.1.1.5`:

```
SNMPv2-MIB::sysName.0 = STRING: SJ-Test-Lab-D-Link
```

From these outputs you can see that if you want to create a performance monitor for D-Link switch uptime, you need to:

1. Create a new device type, since D-Link switches are not supported in this release
2. Create an event type, `PH_DEV_MON_CUST_DLINK_UPTIME`, that will contain the event attribute types `hostName` and `SysUpTime`, which are already part of the FortiSIEM event attribute type library.
3. Create the mapping between the SNMP OIDs and the event attributes:
   - OID `.1.3.6.1.2.1.1.5` and `hostName`.
   - OID `.1.3.6.1.2.1.1.5` and `SysUpTime`.

Creating New Device Types, Event Attribute Types, and Event Types

Device Type

Create a new device type with these attributes:

Field	Setting
Vendor	D-Link
Model	DGS
Version	Any
Device/App Group	**Devices > Network Devices > Router Switch**
Biz Service Group	<no selection>
Description	D-Link Switch

Event Attribute Types and Event Types

Both `sysUptime` and `hostName` are included in the **Event Attribute Types**, so you only need to create a new event type, `PH_DEV_MON_CUST_DLINK_UPTIME`, that will contain them.

Naming Custom Event Types

All custom event types must begin with the prefix `P H_DEV_MON_CUST_ .`

Name	Device Type	Severity	Description
`PH_DEV_MON_CUST_DLINK_UPTIME`	D-Link DGS	0 - Low	D-Link Uptime

Adding the D-Link SNMP Performance Object

In this case, you will create one performance object that will map the SNMP OIDs to the FortiSIEM event attribute types `hostName` and `SysUptime`, and then associate them with the `PH_DEV_MON_CUST_DLINK_UPTIME` event type. When you create the `SysUpTime` mapping you will also add a transform to convert system time to centiseconds to seconds as shown in the second table.

Performance Object Configuration for Event Type PH_DEV_MON_CUST_DLINK_UPTIME

Field	Setting				
Name	D-LinkUptime				
Type	System				
Method	SNMP				
Parent OID	.1.3.6.1.1.2.1.1				
Parent OID is Table	<left cleared>				
	Object Attribute	**Name**	**Format**	**Type**	**Event Attribute**
List of OIDs	.1.3.6.1.1.2.1.1.5	Host Name	String	RawValue	`hostName`
	.1.3.6.1.1.2.1.1.3	Uptime	Timeticks	RawValue	`SysUpTime`
Event Type	`PH_DEV_MON_CUST_DLINK_UPTIME`				
Polling Frequency	10 seconds				

Transform Formula for SysUptime Event Attribute

Type	Formula
custom	uptime/100

Associating Device Types to Performance Objects

In this case you would only need to make one association with the D-Link DGS device you created.

Field	Settings
Name	`D-LinkPerfObj`
Device Types	• D-Link DGS
Perf Objects	• D-LinkUptime(SNMP) - Default Interval:0.17mins

Testing the Performance Monitor

Before testing the monitor, make sure you have defined the access credentials for the D-Link device, created the IP address to credentials mapping, and tested connectivity.

1. Go to **Admin > Device Support > Performance Monitoring**.
2. Select the performance monitor you created, and then click **Test**.
3. For **IP**, enter the address of the device, and select either the Supervisor or Collector node that will retrieve the information for this monitor.
4. Click **Test**.
   You should see `succeed` under **Result**, and the parsed event attributes in the test result pane.
5. When the test succeeds, click **Close**, and then click **Apply** to register the new monitor with the backend module.

Enabling the Performance Monitor

1. Discover or re-discover the device you want to monitor.
2. Once the device is successfully discovered, make sure that the monitor is enabled and pulling metrics.

Writing Queries for the Performance Metrics

You can now use a simple query to make sure that that the metrics are pulled correctly. The search results should display the metrics for the event attributes you defined.

Create a structured historical search with these settings:

Filter Criteria	Display Columns	Time	For Organizations
Structured `Reporting IP IN <IP Range> AND Event Type = "PH_DEV_MON_CUST_DLINK_UPTIME"; Group by: [None]`	Event	Last 10 Minutes	All

Custom WMI Monitor for Windows Domain and Physical Registry

Planning

Mapping Windows WMI Classes to FortiSIEM Event Attribute Types

If you run the command `wmic -U <domain>/<user>%<pwd> //<ip> "select * from Win32_ComputerSystem` against a Windows server, you will see an output similar to this:

```
CLASS: Win32_ComputerSystem
AdminPass-
wordStatus::SEP::Auto-
maticManagedPagefile::SEP::AutomaticResetBootOption::SEP::AutomaticResetCapability::SEP::B

1::SEP::True::SEP::True::SEP::True::SEP::3::SEP::3::SEP::True::SEP::Normal
boot::SEP::WIN2008-ADS::SEP::3::SEP::Win32_ComputerSystem::SEP::-
420::SEP::True::SEP::AT/AT COMPATIBLE::SEP::WIN2008-
ADS::SEP::FortiSIEM.net::SEP::5::SEP::True::SEP::3::SEP::False::SEP::NULL::SEP::
(null)::SEP::3::SEP::(null)::SEP::VMware, Inc.::SEP::VMware Virtual Plat-
form::SEP::WIN2008-ADS::SEP::(null)::SEP::True::SEP::1::SEP::1::SEP::NULL::SEP::
([MS_VM_CERT/SHA1/27d66596a61c48dd3dc7216fd715126e33f59ae7],Welcome to the Virtual
Machine)::SEP::True::SEP::3932100000::SEP::0::SEP::NULL::SEP::False::SEP::0::SEP::-
::0::SEP::3::SEP::(null)::SEP::Windows User::SEP::1::SEP::-1::SEP::-1::SEP::(LM_
Workstation,LM_Server,Primary_Domain_Con-
troller-
,Timesource,NT,DFS)::SEP::OK::SEP::NULL::SEP::0::SEP::NULL::SEP::0::SEP::X86-based
PC::SEP::3::SEP::4293496832::SEP::FortiSIEM\Administrator::SEP::6::SEP::(null)
```

From this output you can see that the `Win32_ComputerSystem` WMI class has two attributes:

1. `Domain`
2. `TotalPhysicalMemory`

From these outputs you can see that if you want to create a performance monitor for Windows Domain and Physical Registry, you need to

1. Create an event type, `PH_DEV_MON_CUST_WIN_MEM`, that will contain the event attribute types `Domain` and `memTotalMB`, both of which are already contained in the FortiSIEM event attribute types library.
2. Create the mapping between the WMI class attributes and the FortiSIEM event attribute types:
   - WMI class attribute `Domain` and `Domain`.
   - WMI class attribute `TotalPhysicalMemory` (Bytes) and `memTotalMB` (type INT64). Because `TotalPhysicalMemory` returns in bytes, and `memTotalMB` is in `INT64`, a transform will be required to convert the metrics.

Creating New Device Types, Event Attributes, and Event Types

Device Type

Since Microsoft Windows is supported by FortiSIEM, you don't need to create a new device type.

Event Attribute Types and Event Types

Both `Domain` and `memTotalMB` are included in the FortiSIEM event attribute type library, so you only need to create a new event type, `PH_DEV_MON_CUST_WIN_MEM`, that will contain them.

Naming Custom Event Types

All custom event types must begin with the prefix `P H_DEV_MON_CUST_` .

Name	Device Type	Severity	Description
`PH_DEV_MON_CUST_WIN_MEM`	Microsoft Windows	0 - Low	Windows Domain and Memory

Adding the Microsoft Windows WMI Performance Object

In this case, you will create one performance object that will map the WMI Class attributes to the FortiSIEM event attribute types `Domain` and `memTotalMB`, and then associate them with the `PH_DEV_MON_CUST_WIN_MEM` event type. When you create the `memTotalMB` mapping you will also add a transform to convert bytes to INT64 as shown in the second table.

Performance Object Configuration for Event Type PH_DEV_MON_CUST_DLINK_UPTIME

Field	Setting
Name	WinMem
Type	System
Method	WMI
Parent Class	Win32_ComputerSystem
Parent Class is Table	<left cleared>
List of Attributes	**Attribute** / **Format** / **Type** / **Event Attribute** Domain / String / RawValue / `domain` TotalPhysicalMemory / Integer / RawValue / `memTotalMB`
Event Type	`PH_DEV_MON_CUST_WIN_MEM`
Polling Frequency	20 seconds

Transform Formula for TotalPhysicalMemory Event Attribute Type

Type	Formula
custom	TotalPhysicalMemory/1024/1024

Associating Device Types to Performance Objects

In this example, you would need to associate Microsoft Windows device types to the performance object.

Edit Device to Performance Object

Field	Settings
Name	WinMisc
Device Types	• Microsoft Windows • Microsoft Windows NT • Microsoft Windows Server 2000 • Microsoft Windows Server 2003 • Microsoft Windows Server 2008 • Microsoft Windows Vista • Microsoft Windows XP
Perf Objects	• WinMem(WMI) - DefaultInterval:0.33mins

Testing the Performance Monitor

Before testing the monitor, make sure you have defined the access credentials for the server, created the IP address to credentials mapping, and tested connectivity.

1. Go to **Admin > Device Support > Performance Monitoring**.
2. Select one of the performance monitors you created, and then click **Test**.
3. For **IP**, enter the address of the Microsoft Windows server, and select either the Supervisor or Collector node that will retrieve the information for this monitor.
4. Click **Test**.
   You should see `succeed` under **Result**, and the parsed event attributes in the test result pane.
5. When the test succeeds, click **Close**, and then click **Apply** to register the new monitor with the backend module.

Enabling the Performance Monitor

1. Discover or re-discover the device you want to monitor.
2. Once the device is successfully discovered, make sure that the monitor is enabled and pulling metrics.

Writing Queries for the Performance Metrics

You can now use a simple query to make sure that that the metrics are pulled correctly. The search results should display the metrics for the event attributes you defined.

Create a structured historical search with these settings:

Filter Criteria	Display Columns	Time	For Organizations
`Host IP = <IP> AND Event Type = " PH_DEV_MON_CUST_WIN_MEM";` **`Group by:`** `[None]`	Event Receive Time,Reporting IP,Domain,Total Memory (MB)	Last 10 Minutes	All

Custom Command Output Monitor

You may already have commands or scripts for your devices that collect important metrics or perform some useful function. By creating a custom command output monitor, you can import the output of those commands into the FortiSIEM event database, where it can be used to create reports, write rules to alert against anomolies, or trigger the execution of scripts. Creating a custom command output monitor involves collecting a sample output from the command, and then creating a performance object that uses regex to parse the command output, maps the output event attributes to FortiSIEM event attribute types, and then associates those to an event type.

- Creating a Custom SSH Command Output Monitor
- Creating a Custom Multi-Line SSH Command Output Monitor
- Creating a Custom WINEXE Command Output Monitor

Device Types Supported for Custom SSH Command Output Monitors

- Linux variants
- Unix variants - IBM AIX, HP UX
- Microsoft Windows (with Cygwin tools installed that allows SSH)
- Cisco IOS, NX-OS, ASA, CatOS
- Juniper JunOS, SSG, ISG
- PaloAlto PANOS
- Fortinet FortiGate
- HP Procurve, H3C
- Extreme Ntwork XOS
- Foundry BigIron
- Avaya ERS

Device Types Supported for Custom WINEXE Command Output Monitors

- Microsoft Windows

Creating a Custom SSH Command Output Monitor

In this example, the regular expression is used to parse a single line of the command output.

Planning

Mapping SSH Command Outputs to FortiSIEM Event Attribute Types

In this example, you want to monitor the output of the `iostat` command. On a Linux machine, the output would look similar to this:

```
[root@centos6-yu ~]# iostat
Linux 2.6.32-358.23.2.el6.x86_64 (centos6-yu.accelops.net.cn)   03/28/2014

avg-cpu:  %user   %nice %system %iowait  %steal   %idle
           0.34    0.00    0.42    0.01    0.00   99.24

Device:            tps   Blk_read/s   Blk_wrtn/s   Blk_read   Blk_wrtn
sda               1.49         5.68        48.77   17172938  147406124
sda1              0.00         0.00         0.00       4726         28
sda2              1.49         5.68        48.77   17166980  147406096
dm-0              6.18         5.68        48.77   17162762  147405360
dm-1              0.00         0.00         0.00       2376          0
dm-2              0.00         0.00         0.00       1098        736
```

This is the event log that you will want to produce in FortiSIEM:

```
[PH_DEV_MON_CUST_CMD]:[hostIpAddr]=10.1.20.52,[hostName]=centos6-yu.FortiSIEM.net,
[readBytes]=17292116.00,
[readRate]=5.71,[tps]=1.49,[writtenBytes]=147500688.00,[writtenRate]=48.73,
[diskName]="sda2"
```

From this example, you can see that to create a monitor for the iostat command output, you would need to:

1. Create the event attribute types `readBytes`, `readRate`, `tps`, `writtenBytes`, `writtenRate`, and `diskName`, to correspond to `Blk_read`, `Blk_read/s`, `tps`, `Blk_wrtn`, `Blk_wrtn/s`, and `Device` from the command output.
2. Create an event type, `PH_DEV_MON_CUST_CMD`, that will contain the event attribute types `readBytes`, `readRate`, `tps`, `writtenBytes`, `writtenRate`, and `diskName`,
3. Create a performance object containing the regular expression that will parse the command output and match value positions to event attribute types, and then associate those event attribute types and values to `PH_DEV_MON_CUST_CMD`.

Creating New Event Attribute Types and Event Types

Event Attributes

Create these event attribute types:

Name	Display Name	Value Type	Display Format Type
diskName	Disk Name	Rawvalue	STRING
tps	Transactions/s	Rawvalue	DOUBLE
readRate	Read Rate	Rawvalue	DOUBLE
readBytes	Read Bytes	Rawvalue	INTEGER
writtenBytes	Written Bytes	Rawvalue	INTEGER
writtenRate	Written Rate	Rawvalue	DOUBLE

Event Types

Naming Custom Event Types

All custom event types must begin with the prefix `P H_DEV_MON_CUST_ .`

Create this event type:

Name	Device Type	Severity
PH_DEV_MON_CUST_CMD	Centos IOS	Low

Adding the iostat Command Output Performance Object

In this case, you will create one performance object that will use a regular expression to parse the command output, match value positions in the command output against FortiSIEM event attributes, and then associate those with the event type `PH_DEV_MON_CUST_CMD`.

Performance Object Configuration for Event Type PH_DEV_MON_CUST_CMD

Field	Setting
Name	cmd-iostat
Type	Application
Method	Login
Used For	Command Output Monitoring
Command	`iostat`
Regular Expression	`(^[^]+)\s+([0-9]+\.?[0-9]+|\d+)\s+([0-9]+\.?[0-9]+|\d+)\s+([0-9]+\.?[0-9]+|\d+)\s+([0-9]+\?[0-9]+|\d+)\s+([0-9]+\.?[0-9]+|\d+)`
Matched Attribute Count	6

List of Attributes

Matched Position	Format	Type	Event Attribute
1	STRING	RawValue	`diskName`
2	DOUBLE	RawValue	`tps`
3	DOUBLE	RawValue	`readRate`
5	INTEGER	RawValue	`readBytes`
6	INTEGER	RawValue	`writtenBytes`
4	DOUBLE	RawValue	`writtenRate`

Event Type	`PH_DEV_MON_CUST_CMD`
Polling Frequency	60 seconds

Associating Device Types to Performance Objects

Field	Settings
Name	`cmd-iostat`
Device Types	• Centos Linux
Perf Objects	• cmd-iostat(SSH)- Default Interval:1mins

Testing the Performance Monitor

Before testing the monitor, make sure you have defined the access credentials for the D-Link device, created the IP address to credentials mapping, and tested connectivity.

1. Go to **Admin > Device Support > Performance Monitoring**.
2. Select the performance monitor you created, and then click **Test**.
3. For **IP**, enter the address of the device, and select either the Supervisor or Collector node that will retrieve the information for this monitor.
4. Click **Test**.
   You should see `succeed` under **Result**, and the parsed event attributes in the test result pane.
5. When the test succeeds, click **Close**, and then click **Apply** to register the new monitor with the backend module.

Enabling the Performance Monitor

1. Discover or re-discover the device you want to monitor.
2. Once the device is successfully discovered, make sure that the monitor is enabled and pulling metrics.

Writing Queries for the Performance Metrics

You can now use a simple query to make sure that that the metrics are pulled correctly. The search results should display the metrics for the event attributes you defined.

Create a structured historical search with these settings:

Filter Criteria	Display Columns	Time	For Organizations
Structured `Reporting IP IN <IP Range> AND Event Type ="PH_DEV_MON_CUST_CM"; Group by:[None]`	Disk Name,Transactions/s, Read Rate, Read Bytes, Written Bytes, Written Rate	Last 10 Minutes	All

Creating a Custom WINEXE Command Output Monitor

There may be times when you want the output of a PowerShell command from a Microsoft server as an input for FortiSIEM. Because PowerShell commands can't be sent via SSH, you need to configure a WINEXE performance object to send the command, parse the output, and associate values to FortiSIEM event attribute types.

Often there is a need to have powershell command output from Microsoft servers into FortiSIEM. These commands cannot be run on Windows systems via SSH. The equivalent way of remotely running a command on Windows systems is Winexe. FortiSIEM will run the Winexe command on Windows systems, collect the output and parse the output into fields for use in FortiSIEM analytics.

Planning

For this example, assume you want to monitor disabled users in Microsoft Active Directory. You would use this command:

```
./winexe -U '<user>%<pwd>' //10.1.2.11 'powershell -NonInteractive -InputFormat
none -OutputFormat text -Command "& {Import-Module ActiveDirectory;Get-ADUser -
LDAPFilter {(useraccountcontrol:1.2.840.113556.1.4.803:=2)}}"'
```

which would have an output similar to this:

```
DistinguishedName : CN=Guest,CN=Users,DC=sh-FortiSIEM,DC=com
Enabled           : False
GivenName         :
Name              : Guest
ObjectClass       : user
ObjectGUID        : dfe5eb21-557f-4550-9cea-5db4ee74317f
SamAccountName    : Guest
SID               : S-1-5-21-2731518400-1375262604-1712717995-501
Surname           :
UserPrincipalName :

DistinguishedName : CN=krbtgt,CN=Users,DC=sh-FortiSIEM,DC=com
Enabled           : False
GivenName         :
Name              : krbtgt
ObjectClass       : user
ObjectGUID        : 13a3703b-185c-4208-98b5-0e65ff638593
SamAccountName    : krbtgt
SID               : S-1-5-21-2731518400-1375262604-1712717995-502
Surname           :
UserPrincipalName :

DistinguishedName : CN=sshd,CN=Users,DC=sh-FortiSIEM,DC=com
Enabled           : False
GivenName         :
Name              : sshd
ObjectClass       : user
ObjectGUID        : aec0ab40-647a-48ae-ba61-9cf31c08794d
SamAccountName    : sshd
SID               : S-1-5-21-2731518400-1375262604-1712717995-1225
Surname           :
```

```
UserPrincipalName :

DistinguishedName : CN=ywang12,DC=sh-FortiSIEM,DC=com
Enabled           : False
GivenName         :
Name              : ywang12
ObjectClass       : user
ObjectGUID        : df694fd2-cf44-49d1-9ecc-42d8a87102c3
SamAccountName    : ywang12
SID               : S-1-5-21-2731518400-1375262604-1712717995-1253
Surname           :
UserPrincipalName :
```

From this example, you can see that to create a monitor for the iostat command output, you would need to:

1. Create an event type, `PH_DEV_MON_CUST_DISABLED_USERS`, that will contain the event attribute types `distName, samAccount`, and `sid,` all of which are already contained in the FortiSIEM event attribute types library, and which match to `DistinguishedName`, `SamAccountName`, and `SID` in the command output.
2. Create a performance object containing the regular expression that will parse the command output and match values against the event attribute types, and then associate those event attribute types and values to `PH_DEV_MON_CUST_CMD`.

After enabling the WIINEXE output monitor, you should see an event similar to this in FortiSIEM:

```
 [PH_DEV_MON_CUST_DISABLED_USERS]:[hostIpAddr]=10.1.2.11,[hostName]=WIN2K8-SHRPNT,
[distName]="CN=krbtgt,CN=Users,DC=sh-FortiSIEM,DC=com",[samAccount]="krbtgt",[sid]-
]="S-1-5-21-2731518400-1375262604-1712717995-502"
```

Creating New Event Attribute Types and Event Types

Event Types

Naming Custom Event Types

All custom event types must begin with the prefix `P H_DEV_MON_CUST_ .`

Create this event type:

Name	Device Type	Severity
`PH_DEV_MON_CUST_DISABLED_USERS`	Cisco IOS	Low

Adding the show interfaces Command Output Performance Object

In this case, you will create one performance object that will use a regular expression to parse the command output, match value positions in the command output against FortiSIEM event attributes, and then associate those with the event type `PH_DEV_MON_CUST_DISABLED_USERS`.

Performance Object Configuration for Event Type PH_DEV_MON_CUST_DISABLED_USERS

Name	WINEXE-AD-Disabled-Users-Output			
Type	System			
Method	WINEXE			
Used For	Command Output Monitoring			
Command	Import-Module ActiveDirectory:Get-ADUser -LDAPFilter {(useraccountcontrol:1.2.840.113556.1.4.803:2)}			
Regular Expression	\nDistinguishedName\s+:\s+(.*?)\n.*?SamAccountName\s+:\s+(.*?)\nSID\s+(.*?)\n			
Matched Attribute Count	3			
List of Attributes	**Matched Position**	**Format**	**Type**	**Event Attribute**
	1	STRING	RawValue	disName
	2	STRING	RawValue	samAccount
	3	STRING	RawValue	sid
Event Type	PH_DEV_MON_CUST_DISABLED_USERS			
Polling Frequency	60 seconds			

Associating Device Types to Performance Objects

Field	Settings
Name	`DiscoverDisabledUsers`
Device Types	• MIcrosoft Windows Server 2008 • MIcrosoft Windows Server 2008 R2 • MIcrosoft Windows Server 2012 • MIcrosoft Windows Server 2012 R2
Perf Objects	• WINEXE-AD-Disabled-Users-Output(WINEXE)-Default Interval:1mins

Testing the Performance Monitor

Before testing the monitor, make sure you have defined the access credentials for the D-Link device, created the IP address to credentials mapping, and tested connectivity.

1. Go to **Admin > Device Support > Performance Monitoring**.
2. Select the performance monitor you created, and then click **Test**.
3. For **IP**, enter the address of the device, and select either the Supervisor or Collector node that will retrieve the information for this monitor.
4. Click **Test**.
   You should see `succeed` under **Result**, and the parsed event attributes in the test result pane.
5. When the test succeeds, click **Close**, and then click **Apply** to register the new monitor with the backend module.

Enabling the Performance Monitor

1. Discover or re-discover the device you want to monitor.
2. Once the device is successfully discovered, make sure that the monitor is enabled and pulling metrics.

Writing Queries for the Performance Metrics

You can now use a simple query to make sure that that the metrics are pulled correctly. The search results should display the metrics for the event attributes you defined.

Create a structured historical search with these settings:

Filter Criteria	Display Columns	Time	For Organizations
Structured `Event Type = PH_DEV_MON_CUST_DISABLED_USERS; Group by:[None]`	Event Receive	Last 10 Minutes	All

Creating a Custom Multi-Line SSH Command Output Monitor

In some cases, the output from a command may run over several lines. An example, as shown in the code block below, is the `show interfaces` command for Cisco IOS routers. Here the information for each interface, such as Vlan1, Vlan2, etc., needs to be consolidated into a single FortiSIEM event. This topic will show you how to configure a performance object for multi-line SSH command outputs, including an example of the regular expression you would use to parse the example output.

- Planning
- Adding the show interfaces Command Output Performance Object
- Associating Device Types to Performance Objects
- Testing the Performance Monitor
- Enabling the Performance Monitor
- Writing Queries for the Performance Metrics

Planning

Mapping a Multi-Line SSH Command Output to FortiSIEM Event Attribute Types

In this example, you want to monitor the output of the '`show interfaces`' command, which would look similar to this for a Cisco IOS router:

```
Vlan1 is up, line protocol is up
  Hardware is EtherSVI, address is 00d0.055b.5000 (bia 00d0.055b.5000)
  Description: DevNet
  Internet address is 192.168.20.1/22
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive not supported
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:00, output 00:00:00, output hang never
  Last clearing of "show interface" counters never
  Input queue: 1/75/12681/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 3583000 bits/sec, 1726 packets/sec
  5 minute output rate 3118000 bits/sec, 1064 packets/sec
  L2 Switched: ucast: 2060202231 pkt, 586057481378 bytes - mcast: 62824587 pkt,
9271104426 bytes
  L3 in Switched: ucast: 43940778993 pkt, 16358818361299 bytes - mcast: 0 pkt, 0
bytes mcast
  L3 out Switched: ucast: 37329069590 pkt, 18769383194932 bytes mcast: 0 pkt, 0
bytes
     44460046444 packets input, 16420615020121 bytes, 0 no buffer
     Received 52655932 broadcasts (0 IP multicasts)
     0 runts, 0 giants, 146 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     37746681819 packets output, 18872504999045 bytes, 0 underruns
     0 output errors, 0 interface resets
```

```
     0 output buffer failures, 0 output buffers swapped out
Vlan2 is up, line protocol is up
  Hardware is EtherSVI, address is 00d0.055b.5000 (bia 00d0.055b.5000)
  Description: ServerNet
  Internet address is 192.168.0.1/24
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive not supported
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:00, output 00:00:01, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/16/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 1652000 bits/sec, 367 packets/sec
  5 minute output rate 258000 bits/sec, 177 packets/sec
  L2 Switched: ucast: 3422947811 pkt, 2275729058787 bytes - mcast: 4291290 pkt,
528654887 bytes
  L3 in Switched: ucast: 17926721335 pkt, 14810495462969 bytes - mcast: 0 pkt, 0
bytes mcast
  L3 out Switched: ucast: 13822525718 pkt, 7788778830975 bytes mcast: 0 pkt, 0
bytes
     19067733427 packets input, 15044884652941 bytes, 0 no buffer
     Received 4283101 broadcasts (0 IP multicasts)
     0 runts, 0 giants, 2 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     13850959642 packets output, 7791605865261 bytes, 0 underruns
     0 output errors, 0 interface resets
     0 output buffer failures, 0 output buffers swapped out
Vlan3 is up, line protocol is up
  Hardware is EtherSVI, address is 00d0.055b.5000 (bia 00d0.055b.5000)
  Description: newbuildnet
  Internet address is 192.168.24.1/24
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive not supported
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:04, output 00:00:01, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 23000 bits/sec, 1 packets/sec
  5 minute output rate 1000 bits/sec, 1 packets/sec
  L2 Switched: ucast: 319623039 pkt, 321540971691 bytes - mcast: 6427637 pkt,
563598014 bytes
  L3 in Switched: ucast: 9237477530 pkt, 10166398798345 bytes - mcast: 0 pkt, 0
bytes mcast
  L3 out Switched: ucast: 5881512921 pkt, 4457997315264 bytes mcast: 0 pkt, 0
bytes
```

```
9289735817 packets input, 10171188457635 bytes, 0 no buffer
Received 6427548 broadcasts (0 IP multicasts)
0 runts, 0 giants, 0 throttles
0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
5939896982 packets output, 4471143181770 bytes, 0 underruns
0 output errors, 0 interface resets
0 output buffer failures, 0 output buffers swapped out
```

From this example, you can see that to create a monitor for the 'show interfaces' command output, you would need to:

1. Create an event type, `PH_DEV_MON_CUST_SHOW_INTF`, that will contain the event attribute types `intfName`, `recvBitsPerSec`, `recvPacketsPerSec`, `sentBitsPerSec`, and `sentPacketsPerSec`, all of which are already contained in the FortiSIEM event attribute types library.
2. Create a performance object containing the regular expression that will parse the command output and match values against the event attribute types, and then associate those event attribute types and values to `PH_DEV_MON_CUST_CMD`.

Creating New Event Attribute Types and Event Types

Event Types

Naming Custom Event Types

All custom event types must begin with the prefix `P H_DEV_MON_CUST_ .`

Create this event type:

Name	Device Type	Severity
`PH_DEV_MON_CUST_SHOW_INTF`	Cisco IOS	Low

Adding the show interfaces Command Output Performance Object

In this case, you will create one performance object that will use a regular expression to parse the command output, match value positions in the command output against FortiSIEM event attributes, and then associate those with the event type `PH_DEV_MON_CUST_SHOW_INTF`.

Performance Object Configuration for Event Type PH_DEV_MON_CUST_SHOW_INTF

<table>
<tr><th>Field</th><th>Setting</th></tr>
<tr><td>Name</td><td>ssh-multiline-CiscoIOS</td></tr>
<tr><td>Type</td><td>System</td></tr>
<tr><td>Method</td><td>Login</td></tr>
<tr><td>Used For</td><td>Command Output Monitoring</td></tr>
<tr><td>Command</td><td>show interfaces</td></tr>
<tr><td>Regular Expression</td><td>\n(\S*?) is [administratively down|up|down]
(?!\n\S.)*5 minute input rate\s+
(\d+)\s+bits\/sec.*?5 minute output rate\s+
(\d+)\s+bits\/sec,\s+(\d+)\s+packets\/sec</td></tr>
<tr><td>Matched Attribute Count</td><td>5</td></tr>
<tr><td>List of Attributes</td><td>
<table>
<tr><th>Matched Position</th><th>Format</th><th>Type</th><th>Event Attribute</th></tr>
<tr><td>1</td><td>STRING</td><td>RawValue</td><td>intfName</td></tr>
<tr><td>2</td><td>INTEGER</td><td>RawValue</td><td>recvBitsPerSec</td></tr>
<tr><td>3</td><td>INTEGER</td><td>RawValue</td><td>recvPacketsPerSec</td></tr>
<tr><td>4</td><td>INTEGER</td><td>RawValue</td><td>sentBitsPerSec</td></tr>
<tr><td>5</td><td>INTEGER</td><td>RawValue</td><td>sentPacketsPerSec</td></tr>
</table>
</td></tr>
<tr><td>Event Type</td><td>PH_DEV_MON_CUST_SHOW_INTF</td></tr>
<tr><td>Polling Frequency</td><td>60 seconds</td></tr>
</table>

Associating Device Types to Performance Objects

Field	Settings
Name	`ssh-Cisco-Intf-Status`
Device Types	• Cisco IOS
Perf Objects	• ssh-multiline-CiscoIOS(SSH)-Default Interval:1mins

Testing the Performance Monitor

Before testing the monitor, make sure you have defined the access credentials for the Cisco IOS device, created the IP address to credentials mapping, and tested connectivity.

1. Go to **Admin > Device Support > Performance Monitoring**.
2. Select the performance monitor you created, and then click **Test**.
3. For **IP**, enter the address of the device, and select either the Supervisor or Collector node that will retrieve the information for this monitor.
4. Click **Test**.
   You should see `succeed` under **Result**, and the parsed event attributes in the test result pane.
5. When the test succeeds, click **Close**, and then click **Apply** to register the new monitor with the backend module.

Enabling the Performance Monitor

1. Discover or re-discover the device you want to monitor.
2. Once the device is successfully discovered, make sure that the monitor is enabled and pulling metrics.

Writing Queries for the Performance Metrics

You can now use a simple query to make sure that that the metrics are pulled correctly. The search results should display the metrics for the event attributes you defined.

Create a structured historical search with these settings:

Filter Criteria	Display Columns	Time	For Organizations
Structured `Event Type = " P H_DEV_MON_CUST_SHOW_ INTF"; Group by:[None]`	Event Receive	Last 10 Minutes	All

Custom File Monitor

You can create custom file monitors to monitor changes to directories and specific files, and also to trigger incidents when the content of a monitored file is changed from a target gold file.

- Agent-less File-Integrity Monitoring
- Agent-less Target File Monitoring

Agent-less File-Integrity Monitoring

You can use file integrity monitoring to make sure that critical files and directories on servers are not modified. When you enable a file integrity monitor for a specific file or directory, the monitor will:

1. Log in to the system using SSH.
2. Compute the checksums of the files or a directory, including all files in the directory.
3. Periodically verify the computed checksums.
4. Create an event when a change to the checksums is detected.

Supported Servers

File and directory integrity monitoring is supported for these servers:

- Linux variants
- Unix variants
- Windows (with Unix tools installed that allow SSH)

Example Events

These are examples of events that are generated by FortiSIEM when a file or directory is modified, deleted, or has its permissions changed.

File Monitors and Event Types

Unlike other custom monitors, you don't need to set the event type to associate with the monitor. When you select **File Monitor** for the **Used For** option, this automatically associates the event types with the file or directory you specify for monitoring. These examples include the event type associated with each monitoring event.

A Directory is Modified by Adding a File

Event Type: `PH_DEV_MON_CUST_FILE_CREATE`

```
Thu Mar 27 16:33:27 2014 CO228SP222 FortiSIEM-FimLog
file="/home/admin/DirectoryMon/file4.txt"
hash="d41d8cd98f00b204e9800998ecf8427e" user="root" group="root"
access="rw-r--r--" size="0"
type="create" hostIp="192.168.64.228"
```

A Specific File is Modified

Event Type: PH_DEV_MON_CUST_FILE_CHANGE_CONTENT

```
Thu Mar 27 16:37:50 2014 CO228SP222 FortiSIEM-FimLog
file="/home/admin/TargetFileMon/tartget1.txt" prehash="3c9d4de73e30f41eabaef892d507894c"
hash="3c9d4de73e30f41eabaef892d507894c" user="root" group="root" access="rw-r--r--"
size="26" targetfilehash="cc3d5ed5fda53dfa81ea6aa951d7e1fe" type="modify"
hostIp="192.168.64.228"
```

A Specific File is Deleted

Event Type: PH_DEV_MON_CUST_FILE_DELETE

```
Thu Mar 27 16:33:52 2014 CO228SP222 FortiSIEM-FimLog  file-
e="/home/admin/DirectoryMon/file3.txt"
hash="cc3d5ed5fda53dfa81ea6aa951d7e1fe" user="root" group="root" access="rw-r--r--
" size="18"
type="delete" hostIp="192.168.64.228"
```

Permissions or Ownership of a Specific File or Any File in a Directory is Changed

Event Type: PH_DEV_MON_CUST_FILE_CHANGE_ATTRIB.

For permissions changes, look for the preaccess and access attributes.

For ownership changes, look for the preuser, user, pregroup, and group attributes.

```
Thu Mar 27 16:33:10 2014 CO228SP222 FortiSIEM-FimLog  file-
e="/home/admin/FileMon/file1.txt"
hash="cc3d5ed5fda53dfa81ea6aa951d7e1fe" preuser="root" user="admin" pre-
group="root" group="admin"
preaccess="rw-r--r--" access="rwxrwxrwx" size="18" type="change" hostIp-
p="192.168.64.228"
```

File Scan Event

Event Type: PH_DEV_MON_CUST_FILE_SCAN

When FortiSIEM scans a file or a directory, this event is generated and can be reported against.

```
Thu Mar 27 13:59:26 2014 CO228SP222 FortiSIEM-FimLog  file-
e="/home/admin/TargetFileMon/tartget1.txt"
hash="cc3d5ed5fda53dfa81ea6aa951d7e1fe" user="root" group="root" access="rw-r--r--
" size="18"
targetfilehash="cc3d5ed5fda53dfa81ea6aa951d7e1fe" type="scan" hostIp-
p="192.168.64.228"
```

Adding the File Integrity Monitoring Performance Object

In Service Provider deployments, the performance object should be created by the Super/Global account, and will apply to all organizations. For both Service Provider and enterprise deployments, the performance object can be created for an organization by any user who has access to the **Admin** tab.

In this case, you will create one performance object for each file or directory you want to monitor. You don't need to create a new event type or event attribute type, as these are automatically associated with the performance object when you select **File Monitoring** for the **Used For** field.

Performance Object Configuration for File Integrity Monitoring

Field	Setting
Name	LinuxFileMon
Type	Application
Method	Login
Used For	File Monitor
File Path	`home/admin/FileMon/file.txt`
Polling Frequency	30 seconds

Performance Object Configuration for Directory Integrity Monitoring

Field	Setting
Name	LinuxDirMon
Type	Application
Method	Login
Used For	File Monitor
File Path	`home/admin/DirectoryMon`
Polling Frequency	30 seconds

Associating Device Types to Performance Objects

You should associate the performance object to the Linux, Unix, or SSH-capable Windows device type that contains the file or directory path you want to monitor.

Testing the Performance Monitor

Before testing the monitor, make sure you have defined the access credentials for the device, created the IP address to credentials mapping, and tested connectivity.

1. Go to **Admin > Device Support > Performance Monitoring**.
2. Select the performance monitor you created, and then click **Test**.

3. For **IP**, enter the address of the device, and select either the Supervisor or Collector node that will retrieve the information for this monitor.
4. Click **Test**.
   You should see `succeed` under **Result**, and the parsed event attributes in the test result pane.
5. When the test succeeds, click **Close**, and then click **Apply** to register the new monitor with the backend module.

Enabling the Performance Monitor

1. Discover or re-discover the device you want to monitor.
2. Once the device is successfully discovered, make sure that the monitor is enabled and pulling metrics.

Writing Queries for the Performance Metrics

You can now use a simple query to make sure that that the metrics are pulled correctly.

Change: Audited File Added/Deleted

Create a structured historical search with these settings:

Filter Criteria	Display Columns	Time	For Organizations
Structured `Event Type IN ("PH_DEV_MON_CUST_FILE_CREATE","PH_DEV_MON_CUST_FILE_DELETE") Group by:[None]`	Event Receive Time	Last 1 Day	All

Change: Audited File Content Modifications

Create a structured historical search with these settings:

Filter Criteria	Display Columns	Time	For Organizations
Structured `Event Type ="PH_DEV_MON_CUST_FILE_DELTA" Group by:[None]`	Event Receive Time, Host	Last 1 Day	All

Change: Audited File Attribute Modifications

Create a structured historical search with these settings:

Filter Criteria	Display Columns	Time	For Organizations
Structured `Event Type =" PH_DEV_MON_CUST_ FILE_CHANGE_ATTRIB" Group by: [None]`	Event Receive Time, Host	Last 1 Day	All

Agent-less Target File Monitoring

You can use target file monitoring to make sure that a specific file, for example a device configuration file, is always identical in content to a gold standard target file that you import into FortiSIEM. When you enable a target file monitor, it will:

1. Pre-compute the checksum of the gold standard target file imported into FortiSIEM.
2. Periodically, log in to the system using SSH and compute the checksum of the file.
3. Create an event when the content of the monitored file is different than the gold standard target file.

Supported Servers

Target file monitoring is supported for these servers:

- Linux variants
- Unix variants
- Windows (with Unix tools installed that allow SSH)

Example Events

Two events that are generated by FortiSIEM when the target file is modified.

File Monitors and Event Types

Unlike other custom monitors, you don't need to set the event type to associate with the monitor. When you select **File Monitor** for the **Used For** option, this automatically associates the event types with the file or directory you specify for monitoring. These examples include the event type associated with each monitoring event.

Event Type: `PH_DEV_MON_CUST_TARGET_FILE_CHANGE`

This indicates that content of the target file has changed. You can see that the values for `prehash` and `hash` are different.

```
Thu Mar 27 16:37:50 2014 CO228SP222 FortiSIEM-FimLog
file="/home/admin/TargetFileMon/tartget1.txt"
prehash="3c9d4de73e30f41eabaef892d507894c" hash="3c9d4de73e30f41eabaef892d507894c"
user="root" group="root" access="rw-r--r--" size="26"
targetfilehash="cc3d5ed5fda53dfa81ea6aa951d7e1fe" type="modify"
hostIp="192.168.64.228"
```

Event Type: `PH_DEV_MON_CUST_TARGET_FILE_DELTA`

This indicates what was changed,as you can see with the`addedItem`, `deletedItem`, `oldSVNVersion`, and `newSVNVersion` attributes.

```
<14>Mar 27 14:02:28 VA223_TestaThon phPerfMonitor[3740]: [PH_DEV_MON_
CUST_TARGET_FILE_DELTA]:[eventSeverity]=PHL_INFO,
```

```
[procName]=phPerfMonitor,[fileName]=phSvnUpdate.cpp,[lineNumber]=205,
[phCustId]=1,[hostName]=CO228SP222,
[hostIpAddr]=192.168.64.228,
[fileName]=/home/admin/TargetFileMon/tartget1.txt,[oldSVNVersion]=15,
[newSVNVersion]=20,
[deletedItem]=(none),[addedItem]=newline;,[phLogDetail]=
```

Adding the File Integrity Monitoring Performance Object

In Service Provider deployments, the performance object should be created by the Super/Global account, and will apply to all organizations. For both Service Provider and enterprise deployments, the performance object can be created for an organization by any user who has access to the **Admin** tab.

In this case, you will create one performance object in which you will upload the gold target file and enter the path to the file you want to monitor. You don't need to create a new event type or event attribute type, as these are automatically associated with the performance object when you select **File Monitoring** for the **Used For** field.

Performance Object Configuration for File Integrity Monitoring

Field	Setting
Name	LinuxTargetFileMon
Type	Application
Method	Login
Used For	File Monitor
File Path	`home/admin/FileMon/file.txt`
Target File	Click **Upload** and browse to the location of the file that you want to use as the gold target

Associating Device Types to Performance Objects

You should associate the performance object to the Linux, Unix, or SSH-capable Windows device type that contains the file or directory path you want to monitor.

Testing the Performance Monitor

Before testing the monitor, make sure you have defined the access credentials for the device, created the IP address to credentials mapping, and tested connectivity.

1. Go to **Admin > Device Support > Performance Monitoring**.
2. Select the performance monitor you created, and then click **Test**.
3. For **IP**, enter the address of the device, and select either the Supervisor or Collector node that will retrieve the information for this monitor.
4. Click **Test**.

You should see `succeed` under **Result**, and the parsed event attributes in the test result pane.

5. When the test succeeds, click **Close**, and then click **Apply** to register the new monitor with the backend module.

Enabling the Performance Monitor

1. Discover or re-discover the device you want to monitor.
2. Once the device is successfully discovered, make sure that the monitor is enabled and pulling metrics.

Checking the Difference between Versions of Monitored Files

When the monitor detects a difference between the files, it will trigger the rule `Audited target file content modified`, and the rule will continue to trigger and generate incidents until the checksums of the files match. You can compare the original monitored file against the new version in the CMDB.

1. Go to **CMDB > Devices**.
2. Select the device where the monitored filed is located
3. Click the **Configuration** tab.
   In the left pane you will see a list of all the files, and their versions, on the device.
4. To compare files, select one, CNTRL/select the other, and then click **Diff**.

Custom Configuration Change Monitoring

This features provides a way for collecting configuration files for any device and monitoring changes.

Define a new vendor, model (Optional)

If the device vendor and model is not yet defined in FortiSIEM, then the new definition needs to be added.

To check whether you device is already defined

1. Go to **Admin > Device Support > Device/App Types**
2. In the **Search** area, type in the vendor name and see if it exists.

To add a new device type

1. Go to **Admin > Device Support > Device/App Types**
2. Click **New**
3. Fill in the following information:
   - **Vendor**: Type in the name of the Vendor (e.g. Fortinet or Cisco)
   - **Model**: Type in the model - be very generic - preferable software model e.g. FortiOS, IOS - do not enter hardware model for appliances
   - **Version**: Most of the time ANY
   - **Device/App Group**: Select the CMDB Group to which the new device will belong
   - **Business Service Group**: Define the Business Service Group to which the new device will belong
   - **Description**: Add description
4. Click **Save**

Create a valid access method

1. Go to **Admin > Setup > Credentials (Step 1)**.
2. Click **Add**.
3. Create an SSH credential
   a. **Device Type** - Select your device.
   b. **Access Protocol** - Set to SSH.
   c. Define **User Name** and **Password**.
4. Click **Save**.
5. Go to **Admin > Setup > Step 2: IP Range to Credentials**.
6. Click **Add**.
7. Enter the following information for IP Range to Credential Mapping:
   1. **IP/Range** - the access IP of the device.
   2. **Credentials** - pick the credential in Step 3.
   3. Click **OK**.
8. Select the entry and Click **Test Connectivity** or **Test Connectivity without Ping**.
9. Make sure **Test Connectivity** succeeds.

Create a Performance Object

1. Go to **Admin > Device Support > Performance Monitoring**.
2. Under **Enter Performance Object** are, Click **New**.
3. Enter the following information to create a new **Performance Object**:
    1. **Name** - enter a name for reference
    2. **Type** - set to System
    3. **Method** - set to LOGIN
    4. **Used For** - set to Configuration Monitoring
    5. **Expect Script** - Click **Upload** to store a configuring pulling expect script in FortiSIEM
    6. **Polling Frequency** - determines how often configuration will be pulled - recommended 30 minutes
4. Click **Save**

Create Device Type to Performance Object association

1. Go to **Admin > Device Support > Performance Monitoring**.
2. Under **Enter Device Type to Performance Object** Association, Click **New**.
3. Enter the following information to create an association:
    1. **Name** - enter a name for reference
    2. **Device Types** - select the relevant device type for custom configuration polling
    3. **Perf Objects** - Select the performance object created in previous step
4. Click **Save**

Discover the device

1. Go to **Admin > Setup > Discovery**.
2. Click **Add**.
3. In **Include Range,** enter the **IP address** of the device.
4. **Click OK**.
5. Select the entry and then click **Discover**.

Validation Check

The expect script will be executed and configuration will be discovered.

1. Go to **Admin > Setup > Monitor Change/Performance**. Search for the device and check the configuration monitoring task under **System Monitor**
2. Go to **CMDB.** Search for the device and check for the configuration under **Configuration** tab for the selected device.

Configuring Event Handling

This section describes certain event handling operations that happen at the moment events are received in FortiSIEM.

- Event Dropping
- Event Forwarding
- Event Organization Mapping
- Multi-line Syslog Handling

Event Dropping

Some devices and applications generate a significant number of logs, which may be very verbose, contain little valuable information, and consume storage resources. You can configure Event Dropping rules that will drop events just after they have been received by FortiSIEM, preventing these event logs from being collected and processed. I mplementing these rules may require some thought to accurately set the event type, reporting device type, and event regular expression match, for example. However, dropped events do not count towards licensed Events per Second (EPS), and are not stored in the Event database. Dropped event also do not appear in reports, and do not trigger rules. You can also specify that events should be dropped but stored, so event information will be available for searches and reports, but will not trigger rules. And example of an event type that you might want to store but not have trigger any rules would be an IPS event that is a false positive.

1. Log in to your Supervisor node.
   For Service Provider deployments you should log in to the Super/Global account if you want to set a system-wide event dropping rule. If you want to set an event-dropping rule for a specific organization, either log in as an administrator for that organization, or or log in using the Super/Global Account and then select the organization to which the rule should apply when you are creating it.
2. Go to **Admin > General Settings > Event Handling**.
3. Under **Event Dropping Rule**, click **Add**.
4. Next to **Reporting Device**, click **Edit**, and use the CMDB Browser to find device group or individual device that you want to create the rule for.
5. Next to **Event Type**, click **Edit**, and use the Event Type Browser to find the group of event types, or a specific event type, that you want to create the rule for.
6. If the event type you select has an **Source IP** or **Destination IP** attribute, you can enter specific IP addresses to which the rule should apply.
7. For **Regex Filter**, enter any regular expressions you want to use to filter the log files.
   If any matches are made against your regular expression, then the event will be dropped.
8. For Service Provider deployments, select the **Organization** to which the rule should apply.
9. Select the **Action** that should be taken when the event dropping rule is triggered.
10. Enter any **Description** for the rule.
11. Click **Save**.

Implementation Notes

1. All matching rules are implemented by FortiSIEM, and inter-rule order is not important. If you create a duplicate of an event dropping rule, the first rule is in effect.
2. If you leave a rule definition field blank, then that field is not evaluated. For example, leaving **Event Type** left blank is the same as selecting **All Event Types**.
3. FortiSIEM drops the event at the first entry point. If your deployment uses Collectors, events are dropped by the Collectors. If your deployment doesn't use Collectors, then the event will be dropped by the Worker or Supervisor where the event is received.
4. You can use the report System Event Processing Statistics to view the statistics for dropped events. When you run the report, select AVG(Policy Dropped Event Rate(/sec) as one of the dimensions for Chart For to see events that have been dropped to this policy.

Event Forwarding

In systems management, many servers may need access to forward logs, traps and Netflows from network devices and servers, but it is often resource intensive for network devices and servers to forward logs, traps and Netflows to multiple destinations. For example, most Cisco routers can forward Netflow to two locations at most. However, FortiSIEM can forward/relay specific logs, traps and Netflows to one or more destinations. A Super, Worker or Collector can forward events - the one which receives and parses the event forwards it. If you want to send a log to multiple destinations, you can send it to FortiSIEM, which will use an event forwarding rule to send it to the desired locations.

1. Log in to your Supervisor node.
2. Go to **Admin > General Settings > Event Handling**.
3. Under **Event Forwarding Rule**, for Service Provider deployments, select the organization for which the rule will apply.
4. Click **Add**.
5. For **Sender IP**, enter the IP address of the device that will be sending the logs.
6. For **Severity**, select an operator and enter a severity level that must match for the log to be forwarded.
7. Select the **Traffic Type** to which the rule should apply.
   The **Forward To > Port** field will be populated based on your selection here.
8. For **Forward to > IP**, enter the IP address to which the event should be forwarded.
9. Click **OK**.

Multiple Destinations from the Same Sender IP

If you want the same sender IP to forward events to multiple destinations, create a rule for each destination.

Duplicate Rules Create Duplicate Logs

FortiSIEM will implement all rules that you create and enable, so if you create a duplicate of an event forwarding rule, two copies of the same log will be sent to the destination IP.

Event Organization Mapping

FortiSIEM can handle reporting devices that are themselves Service Provider and hence have organization names in events that they send. This section describes how you can map organization names in external events to those on FortiSIEM so that those events have the correct FortiSIEM organizations.

Adding Organization Mapping Rules

1. Go to **Admin** > **General Settings** > **Event Handling** > **Event Organization Handling**
2. Click **Add** to add a rule
3. Select **Enabled** if this rule is to be enforced
4. Select the **Device Type** of the sender. This has to be a device that FortiSIEM understands and able to parse events.
5. Select the **Event Attribute** that contains the external organization name. FortiSIEM will map the value in this field to FortiSIEM organization.
6. Select the **Collectors** that are going to receive the events. By default any collectors would be able to do this but it is possible to scope down if needed. This field is optional.
7. Specify the **Reporting IP/Range** of the Service Provider devices that are sending events. Format of this field is a comma separated list of IP addresses intermixed with IP ranges, e.g. 10.1.1.1,10.1.1.2,10.10.1.1-10.10.1.250.
8. Specify the **Org Mapping.**
    a. Click **Edit**
    b. Select the **System (FortiSIEM) organization** on the left column
    c. Click the **Event Organization** and enter the external Organization name corresponding to the System Organization on the left column
9. Click **OK** to Save.

Implementation Notes

Do not define overlapping rules - make sure no overlaps in (Collector, Reporting IP/Range,Event Attribute) between multiple rules.

Multi-line Syslog Handling

Often applications generate a single syslog in multiple lines. For analysis purposes, the multiple lines need to put together into a single log. This feature enables you to do that.

User can write multiple multi-line syslog combining rules based on reporting IP and begin and ending patterns. All matching syslog within the begin and ending pattern are combined into a single log.

To create a multi-line syslog rule,

1. Go to **Admin** > **General Settings** > **Event Handling**
2. Scroll down to **Multiline syslog** section
3. Click **Add**
4. Enter the following information
    a. **Enabled** - check this if the rule needs to be effective
    b. **Sender IP** - the source of the syslog - format is a single IP, IP range, CIDR and a combination of the above separated by comma
    c. **Protocol** - TCP or UDP since syslog can come via either of these protocols
    d. **Organization** - syslog from devices belonging to this organization will be combined into one line
    e. **Begin Pattern** - combining syslog starts when the regular expression specified here is encountered
    f. **End Pattern** - combining syslog stops when the regular expression specified here is encountered
5. Click **Save**

Example 1 - Syslog over UDP

In this case, **Begin Pattern** is required and **End Pattern** is optional.

- If a packet matches the **Begin Pattern**, FortiSIEM will hold it in memory and wait for the next packet.
- If the 2nd packet also matches the **Begin Pattern**, continue waiting.
- If the 3rd packet doesn't match the **Begin Pattern**, flush out the 2 events (1+2 and 3).
- If any packet matches the **End Pattern**, flush out.
- The **Begin Pattern** is in each packet of a multiline syslog. Remove them except the 1st packet.

For example, the receiver gets these packets:

<syslog header> I come to

<syslog header> work

<syslog header> every day

If you set the **Begin Pattern** to a regular expression to match the <syslog header> and leave the **End Pattern** to be empty, then the three syslogs are combined into a single syslog

<syslog header> I come to work every day

If you set the **Begin Pattern** to a regular expression to match the <syslog header> and leave the **End Pattern** to match work, then the first two syslogs are combined into a single syslog, while the third one is left alone.

<syslog header> I come to work

<syslog header> work

Example 2 - Syslog over TCP - octet counting

Octet counting means that there is a header that specifies the length of the syslog. In this case, syslog is not combined. There is no need to combine, since the source can send large syslog messages.

Example 3 - syslog over TCP - non-transparent framing

In non-transparent framing, two syslogs sent over a TCP stream is delineated by the "\n" character. In this case, either **Begin Pattern** or **End Pattern** is required. Both can be present as well.

- If the **Begin Pattern** is matched in the TCP stream, a multi-line syslog combination begins
- If the **End Pattern** is matched in the TCP stream, multi-line syslog combination ends
- If the **Begin Pattern** is again matched in the TCP stream, the previous multi-line syslog combination ends

TCP syslog stream: id=0,name=<1>name=a,id=1<2>name=b,id=2<3>

Begin pattern is <\d+> and end pattern is id=\d+. This results in 3 syslogs

id=0,name=

<1>name=a,id=1

<2>name=b,id=2

And <3> will be held for next packet.

If the Begin pattern is <\d+> and end pattern is empty, this also results in 3 syslogs as before.

Managing FortiSIEM

This chapter describes the following:

General System Administration

Topics in this section contain information on monitoring the health of your FortiSIEM deployment, general system settings such as language, date format, and system logos, and how to add devices to a maintenance calendar.

FortiSIEM Backend Processes

This topic provides a brief description of FortiSIEM backend system processes, and the nodes (Supervisor, Collector, Worker) that use them.

Process	Function	Used by Supervisor	Used by Worker	Used by Collector
phMonitor	Monitoring other processes	X	X	X
phDiscover	Pulling basic data from target	X		X
phPerfMonitor	Execute performance job	X	X	X
phAgentManager	Execute event pulling job	X	X	X
phCheckpoint	Execute checkpoint monitoring	X	X	X
phEventPackage	Uploading event/SVN file to Supervisor/Worker			X
phParser	Parsing event to shared store (SS)	X	X	X
phDataManager	Save event from SS to Event DB	X	X	
phRuleMaster	Determines if a rule should trigger	X		
phRuleWorker	Aggregates data for rules	X	X	
phQueryMaster	Merges data from QueryWorker	X		
phQueryWorker	Executes a query task	X	X	
phReportMaster	Merge data from ReportWorker	X		
phReportWorker	Aggregates data for reports	X	X	
phIPIdentityMaster	Merges IP identity information	X		
phIdentityWorker	Collects IP identity information	X	X	
Apache	Receives event/SVN files from the Collector	X	X	

Administrator Tools

This topic describes administration tools and scripts that are included with your FortiSIEM deployment, along with information on where to find and how to use them.

Tool	Description	How to Use It

phTools	phTools is a simple tool for starting and stopping backend processes, and for getting change log information. When you upgrade your deployment, for example, you would use phTools to stop all backend processes.	Log in to the FortiSIEM host machine as `root`.
		Usage
		`[root@FortiSIEM]# phtools`
		Commands: `--change-log`, `--start`, `--stop`, `--stats`
		Target: ALL
		`--change-log` also supports `ERROR`, `TRACE`, `INFO`, `DEBUG`, `CRITICAL`

TestSegmentReader	Test Segment Reader is used to quickly read data segments in the eventdb through the command line. You can use this to manually inspect data integrity and parsed event attributes.	Log into the FortiSIEM host machine as `root`. **Usage** `[root@FortiSIEM]# TestSegmentReader <segmentDir>`
phExportEvent	Used to export event information to a CSV file	See Exporting Events to Files
TestDBPurger	A script to selectively delete event data per org and time interval **Use Only to Delete Data for a Single Date**: You should only use this script to delete data for a single date and organization. If you try to delete data for multiple dates, the script will fail.	You can find the script at `/opt/phoenix/bin/TestDBPurger`. Run it in terminal mode and follow the instructions.

Managing User Activity

In the User Activity page you can view the users who are logged into your system, user query activity, and locked out users. You can also log users out of the system, stop active user queries, and lock or unlock users from being able to log in. Click the **User Activity** icon in the upper-right corner of the FortiSIEM web interface to access user activity information.

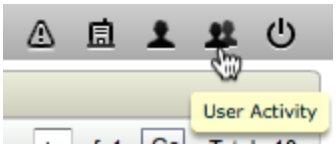

- Managing Logged In Users
- Managing Locked Out Users
- Managing Active User Queries

Managing Logged In Users

In the **Logged In Users** tab of the **User Activity** page you can see the users who are currently logged in to your system. You can also log users out of the system, with an option to lock them out as well.

1. Log in to your Supervisor node.
2. In the upper-right corner of the FortiSIEM web interface, click the **User Activity** icon.

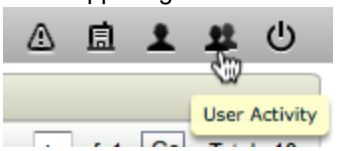

3. Click the **Logged In Users** tab.
   You will see a list of all the users who are currently in your system.
4. If you want to log a user out of the system, select the user and click **Log Out**.
5. If you want to lock a user out of the system, select the user and click **Log Out and Lock Out**.

Managing Locked Out Users

In the **Locked Users** tab of the **User Activity** page you can see the users who are currently locked out of your system, and also unlock them.

1. Log in to your Supervisor node.
2. In the upper-right corner of the FortiSIEM web interface, click the **User Activity** icon.

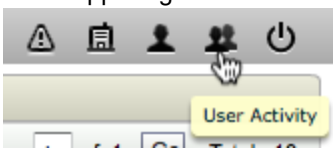

3. Click the **Locked Users** tab.
   You will see a list of all users who are locked out of the system.
4. To unlock a user, select the user and then click **Unlock**.

Managing Active User Queries

In the **User Queries** tab of the **User Activity** page you can see the user queries that are running in your system, and also stop queries.

1. In the upper-right corner of the FortiSIEM web interface, click the **User Activity** icon.

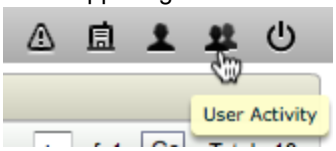

2. Click the **User Queries** tab.
   You will see a list of all the queries that are currently running in your system.
3. To stop a query, select it and then click **Stop Query**.

Creating Maintenance Window for Devices

You can add a device to a maintenance window. During this period, the device is not monitored, and alerts for the device are not triggered. If you have FortiSIEM Service Provider deployment and you log in as a Super/Global user, you can schedule maintenance events for single organizations, the Super/Global organization, or add devices from multiple organizations to the same maintenance event.

1. Log in to your Supervisor node.
2. Go to **Admin > Maintenance Calendar**.
3. Click **Add**.
4. Enter a **Name** and **Description** for the maintenance event.
5. Set the **Time Range** and **Date Range** for the maintenance event.
   Recurring Maintenance Events: Select **From start date on** to set up recurring maintenance events.
6. Under **Groups and Devices**, click **Edit**.
7. If you have FortiSIEM Service Provider deployment, select the **Organization** that has the devices you want to add to the maintenance calendar.
   Multiple Organizations, One Maintenance Event: If you are the Super/Global user, it is possible to add devices from different organizations to the same maintenance event.
8. Add **Folders** or **Items** to the maintenance event by selecting them, and then using the **Folder >>** and **Item >>** buttons to move them into the selection pane.
9. Click **OK** when you're done selecting Folders and Items.
10. Select **Generate incidents for devices under active maintenance** if you want incidents for devices that are part of this maintenance event to be triggered.
11. Click **OK**.
12. You will now see your maintenance event listed on the calendar. Mouse over any calendar entry to view details of the maintenance event.

Creating Maintenance Window for Synthetic Transaction Monitoring jobs

You can add a Synthetic Transaction Monitoring (STM) job to a maintenance event. During the maintenance event, the STM job is not executed and hence related alerts do not trigger.

If you have FortiSIEM Service Provider deployment and you log in as a Super/Global user, you can schedule maintenance events for single organizations, the Super/Global organization, or add devices from multiple organizations to the same maintenance event.

1. Log in to your Supervisor node.
2. Go to **Admin > Maintenance Calendar**.
3. Click **Add**.
4. Enter a **Name** and **Description** for the maintenance event.
5. Set the **Time Range** and **Date Range** for the maintenance event.
   Recurring Maintenance Events: Select **From start date on** to set up recurring maintenance events.
6. Under **Groups and Devices**, click **Edit**.
7. If you have FortiSIEM Service Provider deployment, select the **Organization** that has the devices you want to add to the maintenance calendar.
   Multiple Organizations, One Maintenance Event: If you are the Super/Global user, it is possible to add devices from different organizations to the same maintenance event.
8. Click **Synthetic Transaction Monitor** (STM) to see all the STM jobs under **Items** in the windows below.
9. Select the Items from the bottom left and then click **Item >>** to move them into the selection pane.
10. Click **OK** to Save the configuration.
11. Select **Generate incidents for devices under active maintenance** if you want incidents for devices that are part of this maintenance event to be triggered.
12. Click **OK**.
13. You will now see your maintenance event listed on the calendar. Mouse over any calendar entry to view details of the maintenance event.

Creating Reverse SSH Tunnels to Debug Collector Issues

Using SSH Tunnels to Connect to Managed Endpoints

When you want to quickly debug an issue, you often need to connect to a managed endpoint directly from a browser using protocols such as Telnet/SSH, RDP, or VNC to HTTP(S), depending on the operating system of the endpoint. However, in a Service Provider deployment, the managed endpoint could be behind a firewall and across the Internet. To further complicate matters, the firewall may not permit an inbound connection for management protocols for security reasons, and also may not allow quick policy changes.

The FortiSIEM solution to this situation is to build a reverse SSH tunnel between the Collector and the Supervisor. The firewall already allows HTTP(S) sessions from Collector to Supervisor. After also being configured to also allow SSH connections from Collector to Supervisor, FortiSIEM builds an on-demand reverse SSH Tunnel initiated by the Collector. You can then use the tunnel to open a remote management session from your browser to the remote managed endpoint. This blog post on The Geek Stuff describes the process for setting up reverse SSH tunnels on Linux, and provides some additional technical details.

If the managed endpoint is directly accessible from your browser, FortiSIEM can open a direct session. The devices have to be discovered first, and based on this information, FortiSIEM can determine whether to launch a direct or Collector-based session.

- If the device is discovered by the Supervisor, then it opens a direct session
- If the device is discovered by a Collector, then it opens a reverse SSH tunnel from the collector, and then initiates a session over this tunnel

FortiSIEM has several features for managing SSH tunnels, including:

- You can define the port of the reverse SSH tunnel. By default it is set to **19999,** but it can be changed to any port.
- FortiSIEM automatically times out each tunnel after a day, although you can manually delete a tunnel at any time
- FortiSIEM provides full tunnel management auditing, such as a reporting on who creates and deletes a tunnel
- FortiSIEM supports a broad group of connectivity protocols protocols. You can can launch any connectivity application by specifying the port, and FortiSIEM will create the tunnel.
- RBAC is supported at the Collector level - if the user can visit the Collector health page, then the user can open a remote collector tunnel.

Browser Plugins and Connectivity Protocol Support

Since FortiSIEM runs from a browser, some integrations are possible if certain browser plugins are installed. The best use case is:

- Using the Firefox browser to connect to FortiSIEM
- The FireSSH browser plugin is already installed in Firefox
- You launch a remote session to the managed endpoint over SSH

- FortiSIEM launches the FireSSH browser plugin and passes the managed endpoint IP
- You type in your user name and password, and if the authentication succeeds, then the shell appears

This table lists the browsers, and the protocols supported by their plugins, that you can use to connect to the managed endpoint.

Note: Always type the end host/device credentials for direct connections over a reverse tunnel even though the displayed IP/port belongs to the Supervisor.

Web Browser	Connectivity Protocol	Supported Browser Plugin	Integration
Firefox	SSH	FireSSH	The plugin launches. You need to provide your user name and password for the end host/device
	Telnet	None	A dialog shows the Supervisor's port/tunnel endpoint to connect to. Use your favorite external telnet client to telnet to <Supervisor-IP> and the port.
	HTTP(S)	None required	Another tab opens. You will need to provide your user name and password if the endpoint device requires it.
	RDP	None	A dialog shows the Supervisor's port/tunnel endpoint to connect to. Use your favorite external remote desktop client to connect to <Supervisor-IP> and the port.
	VNC	None	A dialog shows the Supervisor's port/tunnel endpoint to connect to. Use your favorite external VNC client to connect to <Supervisor-IP> and the port.
	Other	None	A dialog shows the Supervisor's port/tunnel endpoint to connect to. Use your favorite external application client to connect to <Supervisor-IP> and the port.

Web Browser	Connectivity Protocol	Supported Browser Plugin	Integration
Chrome	SSH	FireSSH	The plugin launches. You need to provide your user name and password for the end host/device.
	Telnet	None	A dialog shows the Supervisor's port/tunnel endpoint to connect to. Use your favorite external telnet client to telnet to <Supervisor-IP> and the port.
	RDP	Chrome RDP	A dialog opens for the Chrome RDP plugin. Make sure your popup blocker is disabled, or that you allow popups from this site. Click **Launch App** to launch the plugin in a new tab. A dialog shows the Supervisor's port/tunnel endpoint to connect to. Enter <Supervisor-IP>:<Supervisor Port> to connect. Alternatively, you can use your favorite RDP client.
	HTTP(S)	None required	Another tab opens. You will need to provide your user name and password if the endpoint device requires it.
	VNC	None	A dialog shows the Supervisor's port/tunnel endpoint to connect to. Use your favorite external VNC client to connect to <Supervisor-IP> and the port.
	Other	None	A dialog shows the Supervisor's port/tunnel endpoint to connect to. Use your favorite external application client to connect to <Supervisor-IP> and the port.

Web Browser	Connectivity Protocol	Supported Browser Plugin	Integration
Safari (on OSX only)	SSH	Mac Terminal	A new terminal window launches and connects via SSH to <Supervisor-IP> and <Supervisor-port>. Enter your user name and password for the end host/device.
	Telnet	Mac Terminal	A new terminal window launches and connects via telnet to <Supervisor-IP> and <Supervisor-port>. Enter your user name and password for the end host/device.
	RDP	None	A dialog opens for the Chrome RDP plugin. Make sure your popup blocker is disabled, or that you allow popups from this site. Click **Launch App** to launch the plugin in a new tab. A dialog shows the Supervisor's port/tunnel endpoint to connect to. Enter <Supervisor-IP>:<Supervisor Port> to connect. Alternatively, you can use your favorite RDP client.
	HTTP(S)	None required	Another tab opens. You will need to provide your user name and password if the endpoint device requires it.
	VNC	None	A dialog shows the Supervisor's port/tunnel endpoint to connect to. Use your favorite external VNC client to connect to <Supervisor-IP> and the port.
	Other	None	A dialog shows the Supervisor's port/tunnel endpoint to connect to. Use your favorite external application client to connect to <Supervisor-IP> and the port.
Internet Explorer	SSH, Telnet, RDP, HTTP (S), VNC, Other	No plugin integration	Create the tunnel and then connect to the <Supervisor-Port> that is displayed using an external application.

Firewall Configuration

If there is a firewall between the Collector and the Supervisor, the firewall needs to allow SSH from the Collector to the Supervisor. The default setting uses a non-standard port, **19999**, so make sure you configure the firewall between the Collector and the Supervisor to allow outbound TCP connections on port 19999.

Using Role-Based Access Control to Limit Access to Tunnel Creation, Viewing, and Closing

For security and management reasons, you may want to limit the ability of users to create tunnels. The easiest way to do this is through user roles that have defined access capabilities. For example

- To prevent the creation of **any** tunnels for a role, disallow access to the **CMDB** tab for that role, or disallow access to the particular device or device group. This second option lets you create fine-grained controls for tunnel creation, for example:

- Admins who are able to view Network devices can only open tunnels to Network devices
- Admins who are able to view Servers can only open tunnels to Servers
- Admins who are able to view a custom-created device group can only open tunnel to that specific custom group

- To prevent viewing and closing existing tunnels, disallow access to the **Admin > Collector Health** page.

Related Links

- Setting Up User Roles

Auditing the Creation and Deletion of SSH Tunnels

FortiSIEM includes a system-defined report that shows the SSH tunnel open/close history for the time range that you specify.

1. Log in to your Supervisor node.
2. Go to **Analytics > Reports > System Audit**.
3. Select the **SSH Tunnel Open/Close History** report.
4. Run the report as described in Running System and User-Defined Reports and Baseline Reports.

Creating a Remote Tunnel to a Device Monitored by a Collector

Prerequisites

- You should review the browsers and plugins that are supported for the connectivity protocol you want to use to connect to the device.

Procedure

1. Log in to your Supervisor node.
2. Go to **CMDB > Devices**.
3. Search for or browse to the device you want to establish the connection to.
4. In the **IP Address** column for that device, click on the IP address associated with it to open the **Options** menu.
5. In the Options menu, select **Connect To...** .
6. Enter the **Protocol** and **Port** you want to use to connect to the device.
   For **SSH** this is **Port 22**.
7. Select **Create Tunnel**.
   A tunnel will be established between the Supervisor and the Collector that is monitoring the device.
8. Use your browser and plugins to establish remote connectivity to the device as described in Creating Reverse SSH Tunnels to Debug Collector Issues.

Managing Remote Tunnels to Collector Devices

After you have created tunnels to collector devices, you can view and manage those tunnels in the **Collector Health** page.

1. Log in to your Supervisor node.
2. Go to **Admin > Collector Health**.
3. Click **Tunnels**. The existing tunnels will be displayed in a table with these columns:

Column Name	Description
Host IP	The IP address of the managed endpoint
Super Port	Sessions are opened on this port on the Supervisor to connect to the managed endpoint. This ensures that the Supervisor will use the correct tunnel to reach the managed endpoint.
Protocol	The protocol used to establish the connection to the endpoint
Collector	The Collector that monitors the endpoint
PID	The process ID of the tunnel. If you kill this process, it will kill the tunnel
Opened Time	The time when the tunnel was opened

4. You can close a tunnel by selecting it and then clicking **Close**, or you can close all tunnels at the same time by clicking **Close All**.

Managing System Date Format and Logos

The **UI** page under **Admin > General Setting**s contains fields that you can use to change the date format for your FortiSIEM user interface, and to upload logos to be used within the user interface and on PDF reports.

1. Log in to your Supervisor node.
2. Go to **Admin > General Settings > UI**.
3. Select the **Date Format** you want to use to display dates in the user interface, and then click **Change**.
4. Click **Change** to choose a **UI Logo** that will be displayed alongside the main application tabs for your FortiSIEM deployment.
   The logo file must be in in PNG format, and should not be more than 200 pixels wide or 60 pixels high (54 pixels is the ideal height).
5. Click **Change** to choose a **Report Logo** that will be used in the header of reports you export to PDF.
   The logo file must be in SVG format, 160 pixels wide and 40 pixels high, or other dimensions with a 4:1 width/height ratio.

For Service Provider installs, UI Logos can also be set on a per organization basis.

1. SSH to Supervisor via root
2. Change user to admin 'su admin'
3. Change directory by running 'cd /opt/glassfish3/glassfish/domains/domain1/applications/phoenix/phoenix-web-1.0_war/resources/header'
4. Create a logo per organization
   a. mkdir org
   b. cd org
   c. Create Organizations IDs as directories. Eg: ‘mkdir 2001’ (To find Org ids, Goto Admin > Setup Wizard > Organizations > ID)
5. Copy PNG files to respected Organizations as logo.png. For example: /opt/glassfish3/glassfish/domains/domain1/applications/phoenix/phoenix-web-1.0_war/resources/header/org/2001/logo.png
6. Logon to Organization e.g: Org1 (id: 2001) and make sure that UI logo is updated

Steps to convert JPEG, GIF and PNGs to SVG format

Note - FortiSIEM only accepts SVG formatted logo. All other formats must be converted to SVG formats first.

1. Upload 160 x 40 JPEG/GIF/PNG logo to http://vectormagic.com
2. Download SVG formatted logo from converter
3. Upload converted SVG formatted logo on FortiSIEM UI (Admin > General Settings > UI)

Viewing Cloud Health and System Information

The **Admin > Cloud Health** page shows you the status of the nodes in your deployment, as well as the processes running on them.

1. Go to **Admin > Cloud Health**.
2. Click on any node to view its **Process Details**.
   See FortiSIEM Backend Processes for more information about the system role played by each process.
3. You can access other information about your FortiSIEM deployment by clicking the **Alert** icon in the upper-right corner of the user interface, which will show you **Alerts** and **Tasks** for the system within the last 24 hours.

Viewing System Errors

You can view system errors from any page in the FortiSIEM user interface by clicking on the **System Errors** link directly under the URI address window in your browser.

Viewing Collector Health

If your FortiSIEM deployment includes Collectors, you can monitor the status of the Collectors in the **Admin > CollectorHealth** page. You can also upgrade Collectors from this page, as described in Setting Up the Image Server for Collector Upgrades.

1. Log in to your Supervisor node.
2. Go to **Admin > Collector Health**.
3. Select a Collector and click **Show Processes** to see the processes running on that Collector.
   See FortiSIEM Backend Processes for more information about the processes that run on Collectors.
4. You can also **Stop** or **Start** a Collector by selecting it and clicking the appropriate button.

Properties associated with Collector Health include:

Collector Property	Description
Org Name	Name of the organization to which the Collector belongs
Collector Name	The name of the Collector
IP Address	The IP address of the Collector
Status	The status of the Collector as either **Up** or **Down**
Health	Displays the health of the Collector based on the health of the modules running on it. If Health is **Critical**, it means that one of the modules is not running on the Collector.
Up Time	Total time that the Collector has been up
Last Performance Data	The time when the collector last reported its performance status to the cloud
Last Status Update	The time when the collector last reported its status to the cloud
Last Event Data	T he time when the collector last reported events to the cloud
CPU Utilization	Overall CPU utilization of the Collector
Memory Utilization	Overall memory utilization of the Collector
Version	Which version of FortiSIEM the Collector is running on
Build Date	The date on which the version of FortiSIEM the Collector is running on was built
Upgrade Version	If the Collector has been upgraded, the version it was upgraded to
Install Status	If you upgrade the Collector, the status of the upgrade is shown here as either **Success** or **Failed**
Download Status	If an image was downloaded to the Collector as described in Setting Up the Image Server for Collector Upgrades, the status of the download is shown here as **Success** or **Failed**
Allocated EPS	The number of events per second (EPS) dynamically allocated by the system to this collector. See Dynamic Distribution of Events per Second (EPS) across Collectors for more information about how EPS is allocated across Collectors.
Incoming EPS	The EPS that the Collector is currently seeing

Viewing License Information and Adding Nodes to a License

The **License Management** page in the **Admin** tab shows information associated with your current FortiSIEM license, and allows you to add virtual appliances and Report Servers to your deployment as your license allows.

1. Log in to your Supervisor node.
2. Go to **Admin > License Management**.
3. Under **License Information** you will see detailed information about both **Allowed** and **Current Usage** for the number of virtual appliances, EPS, number of devices, and other attributes associated with you FortiSIEM license.
4. Under **VA Information** you will see the name and IP address of the virtual appliances, and their roles, in your FortiSIEM deployment. Click **Add**, and then enter an IP address for other nodes that you want to add to your license.
5. Under **Report Server Information** you will see the IP address of any Report Servers in your deployment. Click **Add**, and then enter an IP address for other Report Servers that you want to add to your license.

Calculations for License Usage Statistics

Statistic	Calculation	Notes
EPS		FortiSIEM calculates the EPS for your system using a counter that records the total number of received events in a three minute time interval. Every second, a thread wakes up and checks the counter value. If the counter is less than 110% of the license limit (using the calculation 1.1 x EPS License x 180) , then FortiSIEM will continue to collect events. If you exceed 110% of your licensed EPS, events are dropped for the remainder of the three minute window, and an email notification is triggered. At the end of the three minute window the counter resets and resumes receiving events.
Number of Devices		Each entry in **CMDB > Devices** counts as one device. Exceptions to this are: • Workstations • Mobile Devices • VoIP Phones These devices are not counted against the number of devices that are licensed for your deployment.

FortiSIEM Event Categories and Handling

This topic provides a brief description of various types of event categories in FortiSIEM.

Event Categories

System Event Category	Description	Counted in EPS License	phstatus -a outout	Stored in DB?
0	External events and not flow events (e.g. syslog, SNMP Trap, Event pulling)	Yes	EPS	Yes
1	Incidents (events that begin with PH_RULE)	No	EPS INTERNAL	Yes
2	FortiSIEM Audit Events (events that begin with PH_AUDIT)	No	EPS INTERNAL	Yes
3	FortiSIEM Internal system logs, free format	No	EPS INTERNAL	Yes
4	External flow events (Netflow, Sflow)	Yes	EPS	Yes
5	FortiSIEM Internal health events for summary dashboards	No	EPS INTERNAL	Yes
6	FortiSIEM Performance Monitoring events (events that begin with PH_DEV_MON)	Yes	EPS PERF	Yes
7	AO Beaconing events	No	EPS INTERNAL	Yes
8	FortiSIEM Real Time Performance Probe Events	No	EPS INTERNAL	No
99	FortiSIEM Internal Rule Engine	No	EPS INTERNAL	No

Event handling at various nodes

Running "phstatus -a" command at various nodes provides the events handled by that node.The output shows the statistics at 3min, 15min and 30 min averages.

```
EPS: 3 Min: 26.19 15 Min: 30.36 30 Min: 28.85
```

```
EPS INTERNAL: 3 Min: 0.35 15 Min: 0.38 30 Min: 0.35
```

```
EPS PERF: 3 Min: 0.00 15 Min: 0.00 30 Min: 0.00
```

- If you run "phstatus -a" at a Collector, the output shows the events handled by that collector
- If you run "phstatus -a" at a Worker, the output shows the events handled by that Worker - includes events sent by devices directly to that Worker or events sent by Collectors
- If you run "phstatus -a" at a Supervisor, you get the aggregated view across all nodes

Reported EPS by events

The following events report eps which includes EPS (EXTERNAL) and EPS PERF - to be measured against license

1. PH_SYSTEM_EVENTS_PER_SEC: this reports eps at a organization level
2. PH_SYSTEM_PERF_EVENTS_PER_SEC: this reports performance monitoring related eps (counted against license)
3. PH_SYSTEM_INTERNAL_EVENTS_PER_SEC: this reports internal eps (not counted against license)
4. PH_SYSTEM_IP_EVENTS_PER_SEC: this reports eps reported by a device level
5. PH_SYSTEM_DEVAPP_EVENTS_PER_SEC: his reports eps reported by a device level but also has vendor, model info

Changing Dashboard Theme

The **UI** page under **Admin > General Setting**s contains fields that you can use to change the theme for widget dashboards.

- My Dashboard
- Availability/Performance > Avail/Perf Widgets
- Biz Svc Dashboard
- Dashboards By Function

To do this:

1. Log in to your Supervisor node.
2. Go to **Admin > General Settings > UI**.
3. Select the **Dashboard Theme** you want to use, and then click **Change**.
4. Refresh the browser.

Global Setting

Currently the dark theme setting is a global setting - so all users would have the same theme.

Installing OS Security Patches

You may want to install OS level security patches to fix some recently found vulnerabilities.

First check whether the CVEs you are interested in have already been patched by the current FortiSIEM version. You can do this by running the following command.

```
rpm --changelog -q httpd
```

To upgrade OS packages on Collectors, run the following command as root

```
/opt/phoenix/bin/phUpdateSystem.sh
```

To upgrade OS packages on Super/Workers, run the following command as root

```
yum -y --exclude=google-chrome-stable update
```

We use a headless chrome browser for STM but chrome is not supported by Google on CentOS6 or 7 platforms. To upgrade that package to the latest version, we use a third party system.

Run the following commands as root on Super/Worker/Collector

```
sed --in-place -e 's/\(.*phsetosaccelops\)/#\1/' /var/spool/cron/root
yum reinstall -y centos-release
wget http://chrome.richardlloyd.org.uk/install_chrome.sh
chmod u+x install_chrome.sh
./install_chrome.sh
sed --in-place -e 's/^#\(.*phsetosaccelops\)/\1/' /var/spool/cron/root
```

Working with the Configuration Management Database (CMDB)

The Configuration Management Database (CMDB) contains:

- Discovered information about your IT infrastructure such as devices, networks, applications, and users
- Information derived from your discovered infrastructure, including network topology and inter-device relationships such as the relationship of WLAN Access Points to Controller, and Virtual Machines to ESX Hosts.
- Information about system objects such as rules, reports, business services, event types, networks, and ports/protocols

You can find and manage all this information under the **CMDB** tab.

CMDB Categorization of Devices and Applications

Categorization of Devices and Applications

FortiSIEM uses four methods to identify and categorize devices and applications in the CMDB.

From Discovery - Network Devices

When FortiSIEM discovers a device, it looks for keywords in the SNMP `sysDescr` attribute and also probes for the SNMP `sysObjectID` attribute. Internal tables are then used to map a discovered device to one or more CMDB device groups based on these attributes.

- Keywords from the `sysDescr` attribute are matched against the system table `Device Vendor and Model`
- Keywords from the `sysObjectID` attribute are matched against the system table `Device Vendor and Model`
- Matches from the Device Vendor and Model table are then matched against the `ApprovedDeviceVendor.csv` table that is used to create the categories in the CMDB Devices/Applications.

From Discovery - Applications

FortiSIEM discovers applications by discovering the processes that are running on a server. The table `AppMapping.csv` maps process names to Applications, Application Groups, and application folders in the CMDB.

From Logs

FortiSIEM includes a large number of log parsers, each of of which is associated with a Device Vendor/Model and Application Vendor/Model. When the log is parsed by FortiSIEM, the Device/Application/Vendor information is matched against the table `ApprovedDeviceVendor.csv`, which then assigns the application or device to the appropriate CMDB Device/Application folder.

Special Cases

There are some special cases that cannot be categorized using discovery or log information. An example is Microsoft Active Directory. It is an application, but there is no explicit process for i.t as it is part of the kernel or big system service. In this case, specific logs are used: Windows Security logs 672, 673 to detect Microsoft Domain Controller 2000, 2003, and Windows Security logs 4768, 4769 to detect Microsoft Windows Domain Controller 2008, 2012.

Examples

Categorizing a Cisco IOS Router/Switch

This is an example of categorizing a device using discovery. In this case, the `Cisco IOS` substring in the SNMP `sysDescr` attribute is used to detect a Cisco IOS device,

```
[desktop]$ snmpwalk -v 2c -c public 192.168.20.1 sysDescr
SNMPv2-MIB::sysDescr.0 = STRING: Cisco IOS Software, s72033_rp Software (s72033_
rp-ADVIPSERVICESK9_WAN-M), Version 12.2(33)SXI1, RELEASE SOFTWARE (fc3)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2009 by Cisco Systems, Inc.
Compiled Sat 28-Mar-09 10:29 by pr
```

Then this entry in `ApprovedDeviceVendor.csv` maps the Device Vendor/Model `Cisco IOS` to the **Router/Switch** category in the CMDB. `PH_SYS_DEVICE_ROUTER_SWITCH` is the internal ID of the category.

```
#id,Vendor,ModelOS,Version,Type,CMDB Folder Id,,Biz Service,Access Pro-
tocol,Parsed,Priority
301,Cisco,IOS,ANY,Appliance,PH_SYS_DEVICE_ROUTER_SWITCH,,,"TELNET,SSH",1,10
```

Categorizing Fortinet Firewalls

This is also an example of categorizing a device by discovery. In this case, the `SNMPv2-SMI::enterprises.12356` substring in the SNMP `sysObjectId` attribute is used to detect a Fortinet Firewall device.

```
[desktop]$ snmpwalk -v 2c -c public 172.16.255.82 sysObjectID
SNMPv2-MIB::sysObjectID.0 = OID: SNMPv2-SMI::enterprises.12356.101.1.502
```

Then this entry in the `ApprovedDeviceVendor.csv` table maps the Device Vendor/Model `Fortinet FortiOS` to the **Firewall** and **Network IOS** categories in the CMDB, since Fortinet is a UTM device. `PH_SYS_DEVICE_FIREWALL` and `PH_SYS_DEVICE_NETWORK_IPS` are the internal IDs of the categories .

```
#id,Vendor,ModelOS,Version,Type,CMDB Folder Id,,Biz Service,Access Pro-
tocol,Parsed,Priority
21,Fortinet,FortiOS,ANY,Appliance,"PH_SYS_DEVICE_FIREWALL,PH_SYS_DEVICE_NETWORK_
IPS,PH_SYS_DEVICE_SEC_GW",,PH_SYS_BizSrvc_FW,"TELNET,SSH",1,10
```

Categorizing Microsoft IIS

This is an example of categorizing an application based on a running process. In this case, SNMP discovers a process `svchost.exe` with the path `-k iissvcs`.

```
[desktop]$ snmpwalk -v 2c -c public 192.168.0.10 | grep 1148
HOST-RESOURCES-MIB::hrSWRunIndex.1148 = INTEGER: 1148
HOST-RESOURCES-MIB::hrSWRunName.1148 = STRING: "svchost.exe"HOST-RESOURCES-
MIB::hrSWRunParameters.1148 = STRING: "-k iissvcs"
```

This entry in the `AppMapping.csv` table is then used to map the process name `svchost.exe` with the path name `-k iissvcs` to a Microsoft IIS application.

```
#Application group name,package signature,process name,process parameter,process
Description,Priority,Ports,Group
Microsoft IIS,,svchost.exe,"-k iissvcs",Microsoft IIS,10,"http,https",PH_SYS_APP_
WEB_SERVER,
```

Categorizing Cisco ASA

This is an example of categorizing a device based on logs. The Cisco ASA parser has has a Device Vendor/Model associated with it, and when a log from the Cisco ASA device is parsed by FortiSIEM, this entry in `ApprovedDeviceVendor.csv` maps the Device Vendor/Model `Cisco ASA` to the **Firewall** and **VPN**

Gateway categories in the CMDB. `PH_SYS_DEVICE_FIREWALL` and `PH_SYS_DEVICE_VPN_GATEWAY` are the internal IDs of these categories.

```
#id,Vendor,ModelOS,Version,Type,CMDB Folder Id,,Biz Service,Access Pro-
tocol,Parsed,Priority
11,Cisco,ASA,ANY,Appliance,"PH_SYS_DEVICE_FIREWALL,PH_SYS_DEVICE_VPN_
GATEWAY",,"PH_SYS_BizSrvc_FW,PH_SYS_BizSrvc_VPN","TELNET,SSH",1,10
```

Categorizing Microsoft IIS

This is an example of categorizing an application based on logs. The Microsoft IIS (via Snare) parser has a Device Vendor/Model associated with it, and when a log from Microsoft IIS is processed by FortiSIEM, this entry in `ApprovedDeviceVendor.csv` maps the Device Vendor/Model `Microsoft` to the **Windows Server** and **Web Server** categories in the CMDB. `PH_SYS_DEVICE_WINDOWS_SERVER` and `PH_SYS_APP_WEB_SERVER` are the internal IDs of these categories.

the following entry in

```
#id,Vendor,ModelOS,Version,Type,CMDB Folder Id,,Biz Service,Access Pro-
tocol,Parsed,Priority
901,Microsoft,IIS,ANY,Application,"PH_SYS_DEVICE_WINDOWS_SERVER,PH_SYS_APP_WEB_
SERVER",Microsoft IIS,,None,1,10
```

Overview of the CMDB User Interface

While the Summary and Widget dashboard views of your IT infrastructure provide real-time monitoring and reporting on your IT infrastructure, the **CMDB** view provides more in-depth detail about devices, applications, users, and other IT infrastructure components as they are listed in the CMDB, as well as the ability to manage these objects.

- Tab Overview
- Inventory Management and Edit Details Controls
- User Interface Controls for Device View
- Data Collection Status

Tab Overview

This screenshot shows the **Device** view of the **CMDB** tab with **Devices** selected in the **Device View** of the IT infrastructure hierarchy. For any type of object you select in the hierarchy, the CMDB will load a **Summary** view of the objects in the top pane, and **Details** for any individual object you select from the summary in the bottom pane. While the available details will change depending on the type of object you select, all objects in the CMDB view will have **Inventory Management** controls in the summary pane, and an **Edit Details** control in the Details pane.

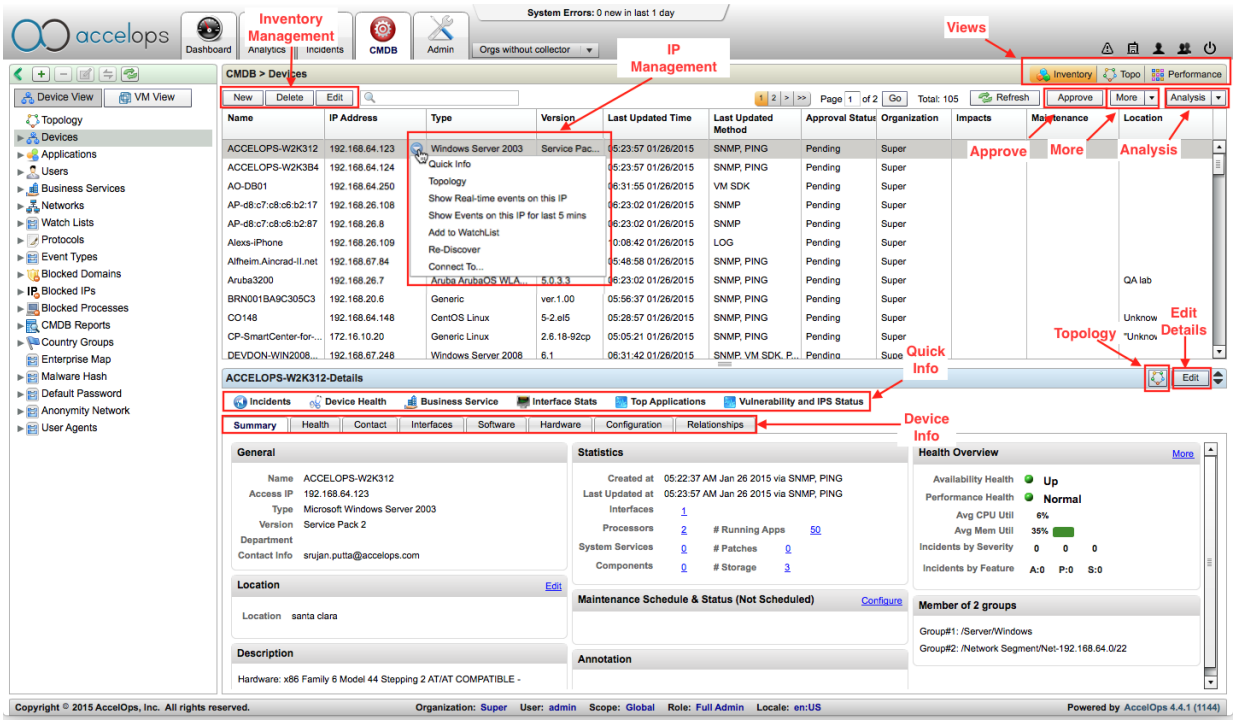

Inventory Management and Edit Details Controls

UI Control	Description
New	Add a new object to the CMDB. **Manually Adding Devices to the CMDB**: In most cases you will want to add devices to the CMDB through the device discovery process, but there are some situations in which you may want to add them manually, as described in Adding Devices to the CMDB Outside of Discovery and Adding a Synthetic Monitoring Test to a Business Service.
Delete	Delete a selected object from the CMDB.
Edit	Edit details about the selected object. You can also use the **Edit Details** button in the **Details** pane for the same purpose. You can also set device-specific properties to use in defining per-device thresholds.

User Interface Controls for Device View

The view of devices in the CMDB provides you with a number of ways to access information about a device. Some of the device user interface controls in the CMDB view you can also find in the dashboard summary view of devices, such as the **Analysis** menu and the **Quick Info** view of a device.

UI Control	Description
Views	• **Inventory** A summary of all devices of that type in the CMDB • **Topo** Shows all devices of the selected type in a topology view • **Performance** Shows a Performance Summary dashboard for all devices of that type
IP Management	Hover your mouse cursor over the IP address associated with a device to open the IP Management menu • **Quick Info** Loads the **Quick Info** for the device, which you can also see by selecting **Quick Info** in the **Analysis** menu • **Topology** Shows the device's location in the network topology, which you can also see by clicking the **Topology** button in the device **Details** pane • **Show Real-Time events on this IP** Loads a Real Time Search with the selected IP address in the search criteria • **Show Events on this IP for the Past 5 Minutes** Loads and Historial search with the selected IP address in the search criteria and the **Time** filter set to **Last 5 Mins** • **Add to WatchList** Add that IP address to a WatchList
More	• **Location** Displays any location information associated with the device • **ChangeOrg** For multi-tenant deployments, change the organization associated with the device • **Impacted Org** Shows organizations that device is associated with • **Maintenance** Displays the maintenance schedule for the device • **Export General Info** Exports a summary view of selected devices, or a detailed view of information for a specific device, in PDF or CSV format
Approve	Approve any newly-discovered devices
Analysis	The Analysis menu contains a number of options for component analytics, depending on the component selected. See Using the Analysis Menu for more information. You can also access the Analysis menu for a component by hovering your mouse over the component's **Device IP** menu until the blue **Quick Info** icon appears, and then clicking the icon.

UI Control	Description
Quick Info	The **Quick Info** view of a device, which you can also access through the **Analysis** menu or hovering your mouse cursor over the **Device IP** column, displays **General and Health information** for the device, and when appropriate, **Identity and Location** information. It also contains links to additional information about the device: • **Incidents** An exportable summary of incidents associated with the device • **Health** **Availability**, **Performance**, and **Security** health information for the device**.** You can also access this information by clicking the **Device Health** user interface control, or by selecting **Device Health** in the **Analysis** menu. • **BizService** Any business services impacted by the device. You can also access this information by selecting **Impacted Business Services** in the **Analysis** menu. • **Applications** Displays a report on the top 10 applications associated with the device by **Average CPU Utilization** over the past hour • **Vulnerability and IP Status** (Not used in the **Dashboard** view) Displays the vulnerability status reports that are also available by selecting **Vulnerability and IPS Status** in the **Analysis** menu • **Hardware Health** (Used only for the **CMDB/Storage** view) Displays health information for the hardware being used for storage • **Interfaces** Displays a report on the top 10 interfaces associated with the device by average throughput • **Topology** Shows the device's location in the network topology**.** You can also access this information by selecting **Topology** in the **Analysis** menu. The Quick Info view also contains two links, **Goto Config Item**, which links to the device entry in the CMDB, and **Goto Identity**, which links to **Analytics > Identity and Location Report**, where you can edit this information for the device.

UI Control	Description
Device Info	Each tab contains information about a specific aspect of the device, as well as an **Edit** button to change information: • **Summary** General organizational and operational information about the device • **Health** **Availability** , **Performance** , and **Security** health reports for the device. You can also access this information by selecting a device in the Summary dashboard, and then click **Health,** or by going to **Quick Info > Health** after selecting the device . If any **Incidents** are displayed, click the number to view the **Incident Summary** . Depending on the reported metric, you can zoom in for a closer look at graphs and reports by clicking the **Magnifying Glass** icon that appears when you hover your mouse cursor over them. • **Monitor** Shows Event Receive Status and Performance Monitor Status - when data was last collected and status • **Contact** Contact information for the device • **Interfaces** Interfaces connected to the device • **Software** Software running on the device. Categories include **Installed Software**, **Running Applications**, **Windows Services**, and **Installed Patches**. In the **Installed Software** category you can use the **Diff... button** to compare different versions of software you've installed. • **Hardware** Information about the hardware associated with the device.Categories include **Processors**, **Storage**, **SAN Storage**, **System BIOS**, **Components**, **SAN Ports**, **RAID Groups**, **LUNs**, and **Storage Groups**. • **Configuration** Configuration files associated with the device. You can compare configuration files by selecting two or more, and then clicking **Diff...** • **Relationships** Other devices that this device interacts with
Topology	Shows the selected device in the Topology view
Edit Details	Click to edit the **Summary**, **Contact Info**, **Interfaces**, and **Properties** for the device

Data Collection Status

Real time data collection status is shown for each device

- Performance Monitor Status
  - Normal - if every performance monitor job status for this device is Normal
  - Warning - if at least one performance monitor job status for this device is Warning and none is critical
  - Critical - if at least one performance monitor job status for this device is Critical
- Event Receive Status
  - Normal - if the event receive status of every protocol for this device is Normal
  - Warning - if the event receive status of at least one protocol for this device is Warning and none is critical
  - Critical - if the event receive status of at least one protocol for this device is Critical

Performance Monitor Job Status is computed as follows. Two **global** constants are defined in **Admin > Device Support > Custom Properties.**

1. Performance Monitoring Time Gap Warning Threshold - multiples of polling interval (default 3)
2. Performance Monitoring Time Gap Critical Threshold - multiples of polling interval (default 5)

Event Receive Job Status is computed as follows. Two **global** constants are defined in **Admin > Device Support > Custom Properties.**

1. Event Receive Time Gap Warning Threshold in minutes (default 10)
2. Event Receive Time Gap Critical Threshold in minutes (default 20)

These constants can also be specified at a per device level from **CMDB** > **Device** > **Bottom pane Edit** > **Properties.** Write new values for these thresholds in the edit box and click **Save.**

Metric	Status	Condition
Performance Monitor Job Status	Normal	Performance Monitoring Time Gap **LESS** THAN Performance Monitoring Time Gap **Warning** Threshold
Performance Monitor Job Status	Warning	Performance Monitoring Time Gap **GREATER** THAN Performance Monitoring Time Gap **Warning** Threshold BUT **LESS** THAN Performance Monitoring Time Gap **Critical** Threshold
Performance Monitor Job Status	Critical	Performance Monitoring Gap **GREATER** THAN Performance Monitoring Time Gap **Critical** Threshold
Event Receive Job Status	Normal	Event Receive Time Gap **LESS** THAN Event Receive Time Gap Warning Threshold
Event Receive Job Status	Warning	Event Receive Time Gap **GREATER** THAN Event Receive Time Gap **Warning** Threshold BUT **LESS** THAN Event Receive Time Gap Critical Threshold
Event Receive Job Status	Critical	Event Receive Time Gap **GREATER** THAN Event Receive Time Gap **Critical** Threshold

The following table shows how the various job types are classified into Performance Monitor or Event Received types

Job Type	Classification in CMDB > Device > Monitor
Jobs defined in Admin > Setup wizard > Monitor Change/performance	Performance Monitor
Jobs defined in Admin > Setup wizard > Pull Events (e.g.	Event Receive
Protocols via which data is pushed to us - syslog, SNMP Trap, Netflow, SFlow, Windows Agents etc	Event Receive

The following rules trigger when certain data collection exceptions happen.

Rule	When does it trigger?	When does it clear?
Missing specific performance metric from a device	Triggers when **Performance Monitor** is Critical for one job for a monitored device	Clears when **Performance Monitor** is Normal for that job from that device
No performance metrics from a device	Triggers when **Performance Monitor** is Critical for ALL jobs for a monitored device	Clears when **Performance Monitor** is Normal for all jobs from that device
FortiSIEM Performance Monitoring Relay Not Working - All Devices delayed	Triggers when **Performance Monitor** is Critical for all devices monitored by a Worker/Collector (that is acting as a Performance Monitoring Relay)	Clears when **Performance Monitor**is Normal for all devices from that Worker/Collector
No logs from a device	Triggers when **Event Receive Job Status** is Critical for one device	Clears when **Event Receive Job Status** is Normal for that device
FortiSIEM Log Relay Not Working - All Devices delayed	Triggers when **Event Receive Job Status** is Critical for all devices to a specific Worker/collector (that is acting as a Log Relay)	Clears when **Event Receive Job Status** is Normal for all devices from that Worker/Collector

Managing CMDB Objects

CMDB objects include discovered devices and their network relationships, as well as system objects like rules and events. You can find the full list of these objects in the **Device View** of the **CMDB** tab, and you can add objects to the database or edit ones that are already there.

- Anonymity Networks and Groups
- Applications
- Blocked Domains
- Blocked IP Addresses
- Blocked URLs
- Blocked Processes
- Country Groups
- Creating CMDB Groups and Adding Objects to Them
- Default Passwords
- Devices
- Event Types
- Malware Hashes
- Networks
- Protocols
- User Agents
- Users
- Watch Lists

Anonymity Networks and Groups

An anonymity network is used to hide one's network identity, and is typically used by malware to hide its originating IP address. Enterprise network traffic should not be originating from or destined to Anonymity network.

When FortiSIEM discovers traffic destined to or originating from anonymity networks, it triggers these rules:

- Inbound Traffic from Tor Network
- Outbound Traffic to Tor Network
- Inbound Traffic from Open Proxies
- Outbound Traffic to Open Proxies

Adding an Anonymity Network

1. Log into your Supervisor node.
2. Go to **CMDB > Anonymity Networks**.
3. Create a group to add the new network to if you are not adding it to an existing group.
   System-Defined Anonymity Network Groups: FortiSIEM provides two default groups for Anonymity Networks:
   - **Open Proxies**: A set of open proxies in the internet. This is a static group.
   - **Tor Nodes** : This group is dynamically updated from https://check.torproject.org/exit-addresses . You can schedule regular updates for this group by clicking on the group name, then click **Update** and provide update scheduling information.
4. Click **Anonymity Network**.
5. Select the group where you want to add the anonymity network.
6. Click **New**.
7. Enter **IP**, **Port**, and **Country** information about the anonymity network.
8. Click the **Calendar** icon to enter the date you created or updated this entry.
9. Click **Save**.

Adding Anonymity Networks to a Group with a CSV File

Instead of manually adding anonymity networks to a group individually, you can upload a CSV file with multiple entries to the group by selecting the group and then clicking **Upload**. You will need to format the file with these fields:

IP Address,Port,Country,Last Update Time and Date

For example:

```
99.99.99.99,,USA,10:00:00 10/02/2014
```

Adding Anonymity Networks to Watch Lists

You can easily add an anonymity network IP address to your watch lists. Hover you mouse cursor over the anonymity network IP address until the icon for the **Options** menu appears, and then select **Add to Watchlist.**

Setting Up an External Data Source for Anonymity Networks

This topic describes how to import anonymity networks information into FortiSIEM from external threat feed websites. Anonymity networks are used by malware to hide their own identity. Two prominent examples of anonymity networks are Open Proxies and TOR Nodes.

Prerequisites

Before proceeding gather the following information about a threat feed web site.

- The website URL
- Credentials required to access the website (optional)
- If the website is not supported by FortiSIEM, you may need to understand the format of the data returned by the URL.
    - if the data is in the comma separated value format (the separator need not be a comma but could be any separator, then a simple integration is possible.
    - If the data is any other format, e.g. XML, then some code needs to be written for integration using the FortiSIEM provided framework

Procedure

Websites with built in support

The following websites are supported

- Threat Stream Open Proxy (https://api.threatstream.com)
- Threat Stream TOR Node (https://api.threatstream.com)

To import data from these websites, follow these steps

1. In the **CMDB > Anonymity Network**, find the website you need to import data from.
2. Select the folder.
3. Click **Update**.
4. Select **Update via API**. The link should show in the edit box.
5. Enter a schedule by clicking on the "+" icon.
6. Enter the schedule parameters - when to start and how often to import. FortiSIEM recommends no more frequent than hourly.
7. Select the type of template you want to create.

Custom websites - CSV data - one-time manual import

This requires that the data to be imported is already in a file in comma separated value format. The required format is

```
IP, Port, Malware Type, Confidence, Severity, Asn, Org, Country, Description, Data Found(MM/DD/YYYY), Last Seen(MM/DD/YYYY)
```

Although many fields are possible, only IP is required

1. Select **CMDB>Anonymity Network**.
2. Click on the "**+**" button on the left navigation tree to bring up the **Create New Anonymity Network Group** dialog.
3. Enter **Group** and add **Description**. Click **OK** to create the folder under **Anonymity Networks**.
4. Select the folder just created.
5. Select **Import from a file.**
6. Click **Browse**; enter the file name and click **Upload**.
7. The imported data will show on the right pane.

Custom websites - CSV data - programmatic import

This requires that the web site data is

- file in comma separated value format (separator can be any special character such as space, tab, hash, dollar etc.)
- one entry is in one line

Note: Although many fields are possible, only the IP is required.

Follow these steps.

1. Select **CMDB>Anonymity Networks**.
2. Click on the "**+**" button on the left navigation tree to bring up the **Create New Anonymity Network Group** dialog.
3. Enter **Group** and add **Description**. Click **OK** to create the folder under **Anonymity Networks**.
4. Select the folder just created.
5. Select **Update via API**
6. For **Website,** Click **Add**.
7. In the **Data Mapping** dialog:
    1. Enter the **URL** of the website
    2. Enter **User Name** and **Password** (optional)
    3. For **Plugin class**, the default class **com.FortiSIEM.service.threatfeed.impl.ThreatFeedWithMappingPolicyService** is shown. Do not modify this for this case.
    4. Enter the correct **Field separator** (by default it is a comma)
    5. Select CSV as the **Data Format**
    6. Enter the **Data Mapping** by choosing the mapped field and the corresponding position in the website data. For example if the IP is in third position, then choose 3 in the **Position** column.
    7. Click **Save**
8. Select a import schedule by clicking **+** on the **Schedule Summary**. Select when to start the import and how often to import to get new data from the website.
9. The imported data will show on the right pane after some time.

New Websites - non-CSV data - programmatic import

This is the most general case where the website data format does not satisfy the previous conditions. In this case, user has to write a Java plugin class by modifying the default system provided one. After the class has been written and fully tested for correctness, follw these steps.

1. Select **CMDB>Anonymity Networks**.
2. Click on the "**+**" button on the left navigation tree to bring up the **Create New Anonymity Network Group** dialog.
3. Enter **Group** and add **Description**. Click **OK** to create the folder under **Anonymity Networks**.
4. Select the folder just created.
5. Select **Update via API**
6. For **Website,** Click **Add**.
7. In the **Data Mapping** dialog:
    a. Enter the **URL** of the website
    b. Enter **User Name** and **Password** (optional)
    c. For **Plugin class**, the custom Java class for this case.
    d. Enter the correct **Field separator** (by default it is a comma)
    e. Select CSV as the **Data Format**
    f. Enter the **Data Mapping** by choosing the mapped field and the corresponding position in the website data. For example if the IP address is in third position, then choose 3 in the **Position** column.
    g. Click **Save**
8. Select a import schedule by clicking **+** on the **Schedule Summary**. Select when to start the import and how often to import to get new data from the website.
9. The imported data will show on the right pane after some time.

Applications

Applications in the CMDB are grouped at the highest level by **Infrastructure** and **User** apps, with further sub-categorization in each of those two categories.

Adding an Application

1. Log in to your Supervisor node.
2. Go to **CMDB > Applications**.
3. Create a new application group or select an existing one.
4. Click **New**.
5. Enter an **Application Name** and **Process**.
6. Enter any other information for the application.
7. Click **Save**.

Malware Domains

The CMDB **Malware Domains** page lists domains that are known to generate spam, host botnets, create DDoS attacks, and generally contain malware. The three default groups included in your FortiSIEM deployment, **MalwareDomainList, Zeus Domains,** and **SANS Domains**, contain malware domains that are derived from the websites malwaredomainlist.com, zeustracker.abuse.ch, and isc.sans.edu. Because malware domains are constantly shifting, FortiSIEM recommends maintaining a dynamically generated list of IP addresses provided by services such as these that is updated on a regular schedule, but you can also add or remove blocked IP addresses from these system-defined groups, and create your own groups based on manual entry of IP addresses or file upload.

Updating System Defined Malware Domain Groups

System defined groups are **MalwareDomainList, Zeus Domains,** and **SANS Domains**, which are updated by their corresponding services. You can set these to update automatically on a schedule, or add or remove individual IP addresses from them.

Setting Schedule

1. Log in to your Supervisor node.
2. Click **CMDB**.
3. Select a system-defined group.
4. Click **Update**.
5. Select **Update Automatically** to open the update scheduler and verify the URI of the update service.
6. Set the schedule for how often you want the list to update from the service.
7. Click **Save**.

Adding/Removing entries

1. If you want to remove a domain or set of domains from the group, clear the **Enable** selection next to the domain name, and then click **Continue** to confirm.
   The domain will still be listed in the group, but it will no longer be blocked. Select **Enable** to resume blocking it.
2. If you want to add a blocked domain to the group, make sure the group is selected, click **New**, and enter information about the blocked IP address.

Changing to STIX/TAXII

If the system defined threat feeds are available via STIX/TAXII, then check the STIX/TAXII box.

Manually Creating Blocked Domains and Groups

1. Create a group under Blocked Domains as described in Creating CMDB Groups and Adding Objects to Them.
2. Select the group you created and click **New**.
3. Enter information for the Blocked Domain you want to add, and then click **Save**.

Adding Blocked Domains to a Group with a CSV File

Instead of manually adding a blocked domain to a user-defined or system group individually, you can upload a CSV file with multiple entries to the group by selecting the group, clicking **Update**, and then selecting **Import Manually**. You will need to format the file with these fields:

`#domain,malware name,date(MM/DD/YYYY),IP,Reverse IP lookup,ASN` For example:
`t3tr.co.cc,Blackhole exploit kit,1/23/2011,173.201.33.90,ip-173-201-33-90.ip.secureserver.net.,26496sq2s.co.cc,Blackhole exploit kit,1/23/2011,173.201.33.90,ip-173-201-33-90.ip.secureserver.net.,26496`

Custom Malware Domain Threat Feed

This topic describes how to import malware domain information into FortiSIEM from external threat feed websites.

Pre-requisites

Before proceeding gather the following information about a threat feed web site.

- The website URL
- Credentials required to access the website (optional)
- If the website is not supported by FortiSIEM, you may need to understand the format of the data returned by the URL.
    - if the data is in the comma separated value format (the separator need not be a comma but could be any separator, then a simple integration is possible.
    - If the data is any other format, e.g. XML, then some code needs to be written for integration using the FortiSIEM provided framework

Threat feed Websites with built in support

The following websites are supported

- Malware domain list (http://www.malwaredomainlist.com)
- Zeus domains (https://zeustracker.abuse.ch)
- SANS Domains (https://isc.sans.edu/feeds/)
- Threat Stream Domains (https://api.threatstream.com)
- Hail-A-TAXII Domains (http://hailataxii.com/)

For Threat Stream the following malware domain types are included

- Command and Control Domain
- Compromised Domain
- Malware Domain
- Dynamic DNS Domain
- APT Domain

To import data from these websites, follow these steps

1. In the **CMDB > Malware Domains**, find the website you need to import data from.
2. Select the folder.
3. Click **Update**.
4. Select **Update via API**. The link should show in the edit box.
5. Enter a schedule by clicking on the "+" icon.

6. Enter the schedule parameters - when to start and how often to import. FortiSIEM recommends no more frequent than hourly.
7. Select the type of template you want to create.

Custom threat feed websites - CSV data - one-time manual import

This requires that the data to be imported is already in a file in comma separated value format. The required format is

```
Domain Name, IP, Reverse Lookup, Malware Type, Confidence, Severity, ASN, Org,
Country,Description, Date Found(MM/DD/YYYY),Last Seen(MM/DD/YYYY)
```

Although many fields are possible, only the Domain Name is required

1. Select **CMDB>Malware Domains**.
2. Click on the "**+**" button on the left navigation tree to bring up the **Create New Malware Domain Group** dialog.
3. Enter **Group** and add **Description**. Click **OK** to create the folder under **Malware Domains**.
4. Select the folder just created.
5. Select **Import from a file.**
6. Click **Browse**; enter the file name and click **Upload**.
7. The imported data will show on the right pane.

Custom threat feed websites - CSV data - programmatic import

This requires that the web site data is

- file in comma separated value format (separator can be any special character such as space, tab, hash, dollar etc.)
- one entry is in one line

Although many fields are possible, only the Domain Name is required.

Follow these steps.

1. Select **CMDB>Malware Domains**.
2. Click on the "**+**" button on the left navigation tree to bring up the **Create New Malware Domain Group** dialog.
3. Enter **Group** and add **Description**. Click **OK** to create the folder under **Malware Domains**.
4. Select the folder just created.
5. Select **Update via API**.
6. For **Website,** Click **Add**.
7. In the **Data Mapping** dialog:
    a. Enter the **URL** of the website
    b. Enter **User Name** and **Password** (optional)
    c. For **Plugin class**, the default class **com.FortiSIEM.service.threatfeed.impl.ThreatFeedWithMappingPolicyService** is shown. Do not modify this for this case.
    d. Enter the correct **Field separator** (by default it is a comma)
    e. Select CSV as the **Data Format**
    f. Enter the **Data Mapping** by choosing the mapped field and the corresponding position in the website data. For example if the domain name is in third position, then choose 3 in the **Position** column.
    g. Click **Save**.

8. Select a import schedule by clicking **+** on the **Schedule Summary**. Select when to start the import and how often to import to get new data from the website.
9. The imported data will show on the right pane after some time.

Custom threat feed websites - non-CSV data - programmatic import

This is the most general case where the website data format does not satisfy the previous conditions. In this case, user has to write a Java plugin class by modifying the default system provided one. Follow instructions in the FortiSIEM ServiceAPI available at FortiSIEM support portal under FortiSIEM ServiceAPI section. After the class has been written and fully tested for correctness, follow these steps.

1. Select **CMDB>Malware Domains**.
2. Click on the "**+**" button on the left navigation tree to bring up the **Create New Malware Domain Group** dialog.
3. Enter **Group** and add **Description**. Click **OK** to create the folder under **Malware Domains**.
4. Select the folder just created.
5. Select **Update via API**.
6. For **Website,** Click **Add**.
7. In the **Data Mapping** dialog:
    a. Enter the **URL** of the website.
    b. Enter **User Name** and **Password** (optional).
    c. For **Plugin class**, choose the custom Java class for this case.
    d. Enter the correct **Field separator** (by default it is a comma).
    e. Select CSV as the **Data Format**.
    f. Enter the **Data Mapping** by choosing the mapped field and the corresponding position in the website data. For example if the domain name is in third position, then choose 3 in the **Position** column.
    g. Click **Save**.
8. Select a import schedule by clicking **+** on the **Schedule Summary**. Select when to start the import and how often to import to get new data from the website.
9. The imported data will show on the right pane after some time.

Custom threat feed websites - STIX formatted data and TAXII import

In this case, the threat feed data is available formatted as STIX and follows the TAXII protocol.

1. Select **CMDB>Malware Domains**.
2. Click on the "**+**" button on the left navigation tree to bring up the **Create New Malware Domain Group** dialog.
3. Enter **Group** and add **Description**. Click **OK** to create the folder under **Malware Domains**.
4. Select the folder just created.
5. Select **Update via API**.
6. For **Website,** Click **Add**.
7. In the **Data Mapping** dialog:
    a. Enter the **URL** of the website.
    b. Enter **User Name** and **Password** (optional).
    c. For **Plugin class**, choose STIX-TAXII and Full.
    d. Click **Save**.
8. Select a import schedule by clicking **+** on the **Schedule Summary**. Select when to start the import and how often

to import to get new data from the website.

9. The imported data will show on the right pane after some time.

Blocked IP Addresses

The CMDB **Blocked IP Addresses** page lists IP addresses that are known to generate spam, host botnets, create DDoS attacks, and generally contain malware. The two default groups included in your FortiSIEM deployment, **Emerging Threats** and **Zeus**, contain IP addresses that are derived from the websites rules.emergingthreats.net and zeustracker.abuse.ch. Because malware IP addresses are constantly shifting, FortiSIEM recommends maintaining a dynamically generated list of IP addresses provided by services such as these that is updated on a regular schedule, but you can also add or remove blocked IP addresses from these system-defined groups, and create your own groups based on manual entry of IP addresses or file upload.

Updating System-Defined Blocked IP Groups

System defined groups are **Emerging Threats** and **Zeus**, which are updated by their corresponding services. You can set these to update automatically on a schedule, or add or remove individual IP addresses from them.

1. Log in to your Supervisor node.
2. Click **CMDB**.
3. Select a system-defined group.
4. Click **Update**.
5. Select **Update Automatically** to open the update scheduler and verify the URI of the update service.
6. Set the schedule for how often you want the list to update from the service.
7. Click **Save**.
8. If you want to remove an IP address or set of IP addresses from the group, clear the **Enable** selection next to the IP address, and then click **Continue** to confirm.
   The IP address will still be listed in the group, but it will no longer be blocked. Select **Enable** to resume blocking it.
9. If you want to add a blocked IP address to the group, make sure the group is selected, click **New**, and enter information about the blocked IP address.

Manually Creating Blocked IP Addresses and Groups

1. Create a group under Blocked IPs as described in Creating CMDB Groups and Adding Objects to Them.
2. Select the group you created and click **New**.
3. Enter information for the Blocked IP address you want to add, and then click **Save**.

Adding Blocked IP Addresses to a Group with a CSV File

Instead of manually adding blocked IP address to a user-defined or system group individually, you can upload a CSV file with multiple entries to the group by selecting the group, clicking **Update**, and then selecting **Import Manually**. You will need to format the file with these fields:

```
#domain,malware name,date(MM/DD/YYYY),IP,Reverse IP lookup,ASN
For example:
t3tr.co.cc,Blackhole exploit kit,1/23/2011,173.201.33.90,ip-173-201-33-
90.ip.secureserver.net.,26496
sq2s.co.cc,Blackhole exploit kit,1/23/2011,173.201.33.90,ip-173-201-33-
90.ip.secureserver.net.,26496
```

Custom Malware IP Threat Feed

This topic describes how to import Malware IP information into FortiSIEM from external threat feed websites.

Prerequisites

Before proceeding gather the following information about a threat feed web site.

- The website URL
- Credentials required to access the website (optional)
- If the website is not supported by FortiSIEM, you may need to understand the format of the data returned by the URL.
    - if the data is in the comma separated value format (the separator need not be a comma but could be any separator, then a simple integration is possible.
    - If the data is any other format, e.g. XML, then some code needs to be written for integration using the FortiSIEM provided framework

Websites with built in support

The following websites are supported

- Emerging threat (http://rules.emergingthreats.net)
- Zeus (https://zeustracker.abuse.ch)
- Threat Stream Malware IP (https://api.threatstream.com)
- Hail-A-TAXII Malware IP (http://hailataxii.com/)

For Threat Stream Malware IP, the following Malware types are imported

- Bot IP
- Actor IP
- APT Email
- APT IP
- Bruteforce IP
- Compromised IP
- Malware IP
- DDoS IP
- Phishing email IP
- Phish URL IP
- Scan IP
- Spam IP

To import data from these websites, follow these steps

1. In the **CMDB > Malware IPs**, find the website you need to import data from.
2. Select the folder.
3. Click **Update**.
4. Select **Update via API**. The link should show in the edit box.
5. Enter a schedule by clicking on the "+" icon.
6. Enter the schedule parameters - when to start and how often to import. FortiSIEM recommends no more frequent than hourly.
7. Select the type of template you want to create.

Custom threat feed websites - CSV data - one-time manual import

This requires that the data to be imported is already in a file in comma separated value format. The required format is

```
Name, Low IP, High IP, Malware Type, Confidence, Severity, ASN, Org, Country ,De-
scription,Data Found(MM/DD/YYYY),Last Seen(MM/DD/YYYY)
```

Although many fields are possible, only Low IP is required. If High IP is not provided, then it is set to Low IP.

1. Select **CMDB>Malware IP**s.
2. Click on the "**+**" button on the left navigation tree to bring up the **Create New Malware IP Group** dialog.
3. Enter **Group** and add **Description**. Click **OK** to create the folder under **Malware IPs**.
4. Select the folder just created.
5. Select **Import from a file.**
6. Click **Browse**; enter the file name and click **Upload**.
7. The imported data will show on the right pane.

Custom threat feed websites - CSV data - programmatic import

This requires that the web site data is

- file in comma separated value format (separator can be any special character such as space, tab, hash, dollar etc.)
- one entry is in one line

Although many fields are possible, only Low IP is required. If High IP is not provided, then it is set to Low IP.

Follow these steps.

1. Select **CMDB>Malware IPs**.
2. Click on the "**+**" button on the left navigation tree to bring up the **Create New Malware IP Group** dialog.
3. Enter **Group** and add **Description**. Click **OK** to create the folder under **Malware IPs**.
4. Select the folder just created.
5. Select **Update via API**
6. For **Website,** Click **Add**.
7. In the **Data Mapping** dialog:
    a. Enter the **URL** of the website.
    b. Enter **User Name** and **Password** (optional).
    c. For **Plugin class**, the default class **com.FortiSIEM.service.threatfeed.impl.ThreatFeedWithMappingPolicyService** is shown. Do not modify this for this case.

d. Enter the correct **Field separator** (by default it is a comma).
e. Select CSV as the **Data Format**.
f. Enter the **Data Mapping** by choosing the mapped field and the corresponding position in the website data. For example if the IP is in third position, then choose 3 in the **Position** column.
g. Click **Save**.

8. Select a import schedule by clicking **+** on the **Schedule Summary**. Select when to start the import and how often to import to get new data from the website.
9. The imported data will show on the right pane after some time.

Custom threat feed websites - non-CSV data - programmatic import

This is the most general case where the website data format does not satisfy the previous conditions. In this case, user has to write a Java plugin class by modifying the default system provided one. Follow instructions in the FortiSIEM ServiceAPI available at FortiSIEM support portal under FortiSIEM ServiceAPI section.

After the class has been written and fully tested for correctness, follow these steps.

1. Select **CMDB>Malware IPs**.
2. Click on the "**+**" button on the left navigation tree to bring up the **Create New Malware IP Group** dialog.
3. Enter **Group** and add **Description**. Click **OK** to create the folder under **Malware IPs**.
4. Select the folder just created.
5. Select **Update via API**.
6. For **Website,** Click **Add**.
7. In the **Data Mapping** dialog:
   a. Enter the **URL** of the website
   b. Enter **User Name** and **Password** (optional)
   c. For **Plugin class**, the custom Java class for this case.
   d. Enter the correct **Field separator** (by default it is a comma)
   e. Select CSV as the **Data Format**
   f. Enter the **Data Mapping** by choosing the mapped field and the corresponding position in the website data. For example if the Low IP is in first position, then choose 1 in the **Position** column.
   g. Click **Save**
8. Select a import schedule by clicking **+** on the **Schedule Summary**. Select when to start the import and how often to import to get new data from the website.
9. The imported data will show on the right pane after some time.

Custom threat feed websites - STIX formatted data and TAXII import

In this case, the threat feed data is available formatted as STIX and follows the TAXII protocol.

1. Select **CMDB>Malware IPs**.
2. Click on the "**+**" button on the left navigation tree to bring up the **Create New Malware IP Group** dialog.
3. Enter **Group** and add **Description**. Click **OK** to create the folder under **Malware IPs**.
4. Select the folder just created.
5. Select **Update via API**
6. For **Website,** Click **Add**.

7. In the **Data Mapping** dialog:
    a. Enter the **URL** of the website
    b. Enter **User Name** and **Password** (optional)
    c. For **Plugin class**, choose STIX-TAXII and Full
    d. Click **Save**.
8. Select a import schedule by clicking **+** on the **Schedule Summary**. Select when to start the import and how often to import to get new data from the website.
9. The imported data will show on the right pane after some time.

Blocked URLs

The CMDB **Blocked URLs** page lists URLs that are known to host malware.

The **Threat Stream Blocked URL** group is included in your FortiSIEM deployment.

Updating System-Defined Blocked URL Group

System defined groups are **Threat Stream Blocked URL**, which are updated by its own service. You can set these to update automatically on a schedule.

1. Log in to your Supervisor node.
2. Click **CMDB**.
3. Select **Threat Stream Blocked URL**.
4. Click **Update**.
5. Set Schedule
    1. Select **Update Automatically** to open the update scheduler and verify the URI of the update service.
    2. Set the **schedule** for how often you want the list to update from the service.
    3. Click **OK**.
    4. Click **Save**
6. Set user name and password
    1. Select the link (https://api.threatstream.com/api/v1/intelligence/)
    2. Click **Edit**
    3. Enter **User Name** and **Password**
    4. Set **Data Format** to **Custom** and **Incremental**
    5. Click **Save**

Custom Malware URL Threat Feed

This topic describes how to import Malware URL information into FortiSIEM from external threat feed websites.

Prerequisites

Before proceeding gather the following information about a threat feed web site.

- The website URL
- Credentials required to access the website (optional)
- If the website is not supported by FortiSIEM, you may need to understand the format of the data returned by the URL.
    - If the data is in comma separated value (CSV) format, then a simple integration is possible. Note that the separator need not be a comma but could be any separator.
    - If the data is any other format, e.g. XML, then some code needs to be written for integration using the FortiSIEM provided framework

Threat feed websites with built in support

The following websites are supported

- Threat Stream Malware URL (https://api.threatstream.com)
- FortiSandbox Malware URL
- Hail-A-TAXII Malware IP (http://hailataxii.com/)

To import data from these websites, follow these steps

1. In the **CMDB > Malware URLs**, find the website you need to import data from.
2. Select the folder.
3. Click **Update**.
4. Select **Update via API**. The link should show in the edit box.
5. Enter a schedule by clicking on the "+" icon.
6. Enter the schedule parameters - when to start and how often to import. FortiSIEM recommends no more frequent than hourly.

Custom threat feed websites - CSV data - one-time manual import

This requires that the data to be imported is already in a file in comma separated value format. The required format is

```
URL, Malware Type, Confidence, Description,Last Seen(MM/DD/YYYY)
```

1. Select **CMDB>Malware URL**s.
2. Click on the "**+**" button on the left navigation tree to bring up the **Create New Malware URL Group** dialog.
3. Enter **Group** and add **Description**. Click **OK** to create the folder under **Malware URLs**.
4. Select the folder just created.
5. Select **Import from a file.**
6. Click **Browse**; enter the file name and click **Upload**.
7. The imported data will show on the right pane.

Custom threat feed websites - CSV data - GUI import

This requires that the web site data has the following structure.

- The file in comma separated value format (separator can be any special character such as space, tab, hash, dollar etc.)
- One line has only one entry

Follow these steps.

1. Select **CMDB>Malware URLs**.
2. Click on the "**+**" button on the left navigation tree to bring up the **Create New Malware URL Group** dialog.
3. Enter **Group** and add **Description**. Click **OK** to create the folder under **Malware URLs**.
4. Select the folder just created.
5. Select **Update via API**
6. For **Website,** Click **Add**.
7. In the **Data Mapping** dialog:
    a. Enter the **URL** of the website
    b. Enter **User Name** and **Password** (optional)
    c. For **Plugin class**, the default class **com.FortiSIEM.service.threatfeed.impl.ThreatFeedWithMappingPolicyService** is shown. Do not modify this for this case.
    d. Enter the correct **Field separator** (by default it is a comma)
    e. Set **Data Format** to CSV.
    f. Enter the **Data Mapping** by choosing the mapped field and the corresponding position in the website data. For example if the URL is in third position, then choose 3 in the **Position** column.
    g. Click **Save**
8. Select a import schedule by clicking **+** on the **Schedule Summary**. Select when to start the import and how often to import to get new data from the website.
9. The imported data will show on the right pane after some time.

Custom threat feed websites - non-CSV data - programmatic import

This is the most general case where the website data format is not CSV. In this case, user has to write a Java plugin class by modifying the default system provided one. Follow instructions in the FortiSIEM ServiceAPI available at FortiSIEM support portal under FortiSIEM ServiceAPI section.

After the class has been written and fully tested for correctness, follow these steps.

1. Select **CMDB>Malware URLs**.
2. Click on the "**+**" button on the left navigation tree to bring up the **Create New Malware URL Group** dialog.

3. Enter **Group** and add **Description**. Click **OK** to create the folder under **Malware URLs**.
4. Select the folder just created.
5. Select **Update via API**
6. For **Website,** Click **Add**.
7. In the **Data Mapping** dialog:
    a. Enter the **URL** of the website
    b. Enter **User Name** and **Password** (optional)
    c. For **Plugin class**, the custom Java class for this case
    d. Select Custom as the **Data Format**.
    e. Click **Save**
8. Select a import schedule by clicking **+** on the **Schedule Summary**. Select when to start the import and how often to import to get new data from the website.
9. The imported data will show on the right pane after some time.

Custom threat feed websites - STIX formatted data and TAXII import

In this case, the threat feed data is available formatted as STIX and follows the TAXII protocol.

1. Select **CMDB>Malware URLs**.
2. Click on the "**+**" button on the left navigation tree to bring up the **Create New Malware URL Group** dialog.
3. Enter **Group** and add **Description**. Click **OK** to create the folder under **Malware URLs**.
4. Select the folder just created.
5. Select **Update via API**.
6. For **Website,** Click **Add**.
7. In the **Data Mapping** dialog:
    a. Enter the **URL** of the website.
    b. Enter **User Name** and **Password** (optional).
    c. For **Plugin class**, choose STIX-TAXII and Full.
    d. Click **Save**
8. Select a import schedule by clicking **+** on the **Schedule Summary**. Select when to start the import and how often to import to get new data from the website.
9. The imported data will show on the right pane after some time.

Country Groups

The Country Groups page contains a list of all the country names in the FortiSIEM geolocation database. You can also create folders that represent different organizations of countries for use in Analytics.

Adding a New Country or Country Group

1. Log in to your Supervisor node.
2. Go to **CMDB > Country Groups**.
3. Select an existing country group, or create a new one.
4. Click **New**.
5. Enter a name and description for the new country.
6. Click **Save**.

Creating CMDB Groups and Adding Objects to Them

In the CMDB browser pane you will see several categories, or groups, for each type of CMDB object. For example, under **Applications**, you will see the groups **Infrastructure App**, **User App**, and **Ungrouped**, with additional subcategorization within each of those groups. You can create your own groupings and add CMDB objects to them.

1. Log in to your Supervisor node.
2. Click **CMDB**.
3. In the CMDB browser pane, select the type of CMDB object you want to create a group for, and then click **+**.
4. Enter a **Group** name and **Description**.
5. Under **Select Group Members**, select any existing groups from which you would like to add objects to your new group.
   The group containing all the CMDB objects of this type is selected by default.
6. Select the objects you want to add to the group, and then click **>>** to add them to the group.
7. Click **OK**.
   Your new group, and the objects it contains, will be listed under that CMDB object type in the CMDB browser pane. You can add objects directly to the group by selecting it in the CMDB browser pane, and then following the process for adding a new object.

Default Passwords

The CMDB Default Password page contains a list of default vendor credentials. These well-known credentials should never be used in production. During device discovery FortiSIEM checks if the device credentials are still set to default , and the system rule `Default Password Detected by System` triggers an incident if they are.

A sample raw event log for a default password incident:

```
<174>Oct 20 22:50:03 [PH_AUDIT_DEFAULT_PWD_MATCH]:
[phEventCategory]=2,[appTransportProto]=SNMP,
[reptModel]=Firewall-1 SPLAT,[srcIpAddr]=192.168.19.195,
[phCustId]=1,[sessionId]=0f8bdee2b6a265c4bd075fc777ed,
[procName]=AppServer,[reptVendor]=Checkpoint,
[hostIpAddr]=172.16.0.1,[hostName]=SJ-QA-F-Lnx-CHK,
[eventSeverity]=PHL_INFO,[user]=,[phLogDetail]=Default
password matches for the same composite key (Vendor, Model,
Access method, User Name, Password)
```

Adding a New Default Password

1. Log in to your Supervisor node.
2. Go to **CMDB > Default Passwords**.
3. Select a group where you want to add the default password, or create a new one.
4. Click **New**.
5. Select the **Vendor** and **Model** of the device for which you want to enter a default password.
6. Select the **Access Protocol** that is used to connect to the device.
7. Enter the default **User Name** and **Password** for the device.
8. Click **Save**.

Adding Default Passwords to a Group with a CSV File

You can upload a CSV file with multiple entries to the a default password group by selecting the group, clicking **Import**, and then browsing to a CSV file. You will need to format the file with these fields:

```
Vendor,Model,Access Protocol,User Name,Password
```

For example:

```
Microsoft,Windows,WMI,Administrator,Administrator
```

Devices

You would typically add devices to the CMDB through the Discovering Infrastructure process. However, there may be situations in which you want to add devices to the CMDB manually. For example, you may not have access credentials for a device but still want to be able to include network information about it so that logs received by FortiSIEM can be parsed properly. These topics describe those situations and provide instructions for how to successfully add a device to the CMDB:

- Adding Devices to the CMDB Outside of Discovery
- Adding a Synthetic Monitoring Test to a Business Service

Event Types

The CMDB **Event Types** page lists the types of events that are collected for supported devices.

Adding a New Event Type

1. Log in to your Supervisor node.
2. Go to **CMDB > Event Types**.
3. Select a group to add the new event to, or create a new one.
4. Click **New**.
5. Enter a **Name, Display Name,** and **Description** for the event type.
6. Select the **Device** to associate with this event type.
7. Select the level of **Severity** associated with this event type.
8. For **CVE IDs**, enter links to any vulnerabilities associated with this event type as cataloged by the National Vulnerability Database.
9. Click **Save**.

Malware Hashes

The CMDB **Malware Hash** page can be used to define a list of malware files and their hash functions. When FortiSIEM monitors a directory, it generates these directory events:

Directory Event	Generated by
PH_DEV_MON_CUST_FILE_CREATE	New file creation
PH_DEV_MON_CUST_FILE_SCAN	Directory is scanned
PH_DEV_MON_CUST_FILE_CHANGE_CONTENT	Changes in file content

When FortiSIEM scans a file and collects its hash, it uses the system rule `Malware Hash Check` to check the list of malware hashes, and triggers an alert if a match is found.

Adding a New Malware Hash

1. Log in to your Supervisor node.
2. Go to **CMDB > Malware Hash**.
3. Select a group where you want to add the malware hash, or create a new one.
4. Click **New**.
5. Enter information for the malware hash.
6. Click **Save**.

Adding Malware Hashes to a Group with a CSV File

You can upload a CSV file with multiple entries to the a default password group by selecting the group, clicking **Import**, and then browsing to a CSV file. You will need to format the file with these fields:

```
BotNet name, Algorithm, Hash Code, Controller IP, Country, Confidence, Last Seen Time
```

For example:

```
MyBotnet,SHA,aaecdrgt0987995dae567812,101.1.1.2,87,China,100,10/20/2014
```

Updating System Defined Malware Hash Group

Current system defined groups are updated by its own service

- Threat Stream Malware Hash
- FortiSandbox Malware Hash

You only need to set these to update automatically on a schedule.

- Log in to your Supervisor node.
- Click **CMDB**.
- Select a system-defined group.

- Click **Update**.
- Select **Update Automatically** to open the update scheduler and verify the URI of the update service.
- Set the schedule for how often you want the list to update from the service.
- Click **Save**.
- If you want to remove an IP address or set of IP addresses from the group, clear the **Enable** selection next to the IP address, and then click Continue to confirm.
- The IP address will still be listed in the group, but it will no longer be blocked. Select **Enable** to resume blocking it.
- If you want to add a malware IP address to the group, make sure the group is selected, click **New**, and enter information about the blocked IP address.

Manually Creating Manual Hash

1. Create a group under Malware Hash as described in Creating CMDB Groups and Adding Objects to Them.
2. Select the group you created and click New.
3. Enter information for the Malware Hash you want to add, and then click Save.

Adding Blocked URLs to a Group with a CSV File

Instead of manually adding a blocked domain to a user-defined or system group individually, you can upload a CSV file with multiple entries to the group by selecting the group, clicking Update, and then selecting Import from a file. You will need to format the file with these fields:

```
Botnet Name, Algorithm, Has Code, Controller IP, Malware Type, Confidence,
Severity, Asn, Org, Country, Description, Data Found(MM/DD/YYYY), Last Seen
(MM/DD/YYYY)
```

Custom Malware Hash Threat Feed

This topic describes how to import Malware Hash information into FortiSIEM from external threat feed websites.

Prerequisites

Before proceeding gather the following information about a threat feed web site.

- The website URL
- Credentials required to access the website (optional)
- If the website is not supported by FortiSIEM, you may need to understand the format of the data returned by the URL.
    - if the data is in the comma separated value format (the separator need not be a comma but could be any separator, then a simple integration is possible.
    - If the data is any other format, e.g. XML, then some code needs to be written for integration using the FortiSIEM provided framework

Threat feed websites with built in support

The following websites are supported

- ThreatStream Malware Hash (https://api.threatstream.com)
- FortiSandbox Malware Hash
- Hail-A-TAXII Malware IP (http://hailataxii.com/)

To import data from these websites, follow these steps

1. In the **CMDB > Malware Hash**, find the website you need to import data from.
2. Select the folder.
3. Click **Update**.
4. Select **Update via API**. The link should show in the edit box.
5. Enter a schedule by clicking on the "+" icon.
6. Enter the schedule parameters - when to start and how often to import. FortiSIEM recommends no more frequent than hourly.
7. Select the type of template you want to create.

Custom threat feed websites - CSV data - one-time manual import

This requires that the data to be imported is already in a file in comma separated value format. The required format is:

```
Botnet Name, Algorithm, Hash Code, Controller IP, Malware Type, Confidence, Severity,
Asn, Org, Country, Description, Data Found(MM/DD/YYYY), Last Seen(MM/DD/YYYY), High IP,
Malware Type, Confidence, Severity, ASN, Org, Country ,Description,Data Found
```

```
(MM/DD/YYYY),Last Seen(MM/DD/YYYY)
```

Note: Although many fields are possible, only Botnet Name and Hash Code are required.

1. Select **CMDB > Malware Hash**.
2. Click on the "+" button on the left navigation tree to bring up the "Create New Malware Hash Group" dialog.
3. Enter Group and add Description. Click OK to create the folder under Malware Hash.
4. Select the folder just created.
5. Select Import from a file.
6. Click Browse; enter the file name and click Upload.
7. The imported data will show on the right pane.

Custom threat feed websites - CSV data - programmatic import

This requires that the web site data is:

- file in comma separated value format (separator can be any special character such as space, tab, Hash, dollar etc.)
- one entry is in one line

Note: Although many fields are possible, only Botnet Name and Hash Code are required.

Follow these steps.

1. Select **CMDB > Malware Hash**.
2. Click on the "+" button on the left navigation tree to bring up the "Create New Malware Hash Group" dialog.
3. Enter **Group** and add **Description**. Click OK to create the folder under Malware Hash.
4. Select the folder just created.
5. Select **Update via API**
6. For Website, Click **Add**.
7. In the Data Mapping dialog:
    1. Enter the URL of the website
    2. Enter **User Name** and **Password** (optional)
    3. For Plugin class, the default class com.FortiSIEM.service.threatfeed.impl.ThreatFeedWithMappingPolicyService is shown. Do not modify this for this case.
    4. Enter the correct Field separator (by default it is a comma)
    5. Select **CSV** as the Data Format
    6. Enter the **Data Mapping** by choosing the mapped field and the corresponding position in the website data. For example if the Hash is in third position, then choose 3 in the Position column.
    7. Click **Save**.
8. Select a import schedule by clicking + on the Schedule Summary. Select when to start the import and how often to import to get new data from the website.
9. The imported data will show on the right pane after some time.

Custom threat feed websites - non-CSV data - programmatic import

This is the most general case where the website data format does not satisfy the previous conditions. In this case, user has to write a Java plugin class by modifying the default system provided one. Follow instructions in

the FortiSIEM ServiceAPI available at FortiSIEM support portal under FortiSIEM ServiceAPI section. After the class has been written and fully tested for correctness, follow these steps.

1. Select **CMDB>Malware Hash**.
2. Click on the "+" button on the left navigation tree to bring up the "Create New Malware Hash Group" dialog.
3. Enter **Group** and add **Description**. Click **OK** to create the folder under Malware Hash.
4. Select the folder just created.
5. Select **Update** via API
6. For Website, Click **Add**.
7. In the Data Mapping dialog:
    a. Enter the URL of the website
    b. Enter User Name and Password (optional)
    c. For Plugin class, the custom Java class for this case.
    d. Enter the correct Field separator (by default it is a comma)
    e. Select CSV as the Data Format
    f. Enter the Data Mapping by choosing the mapped field and the corresponding position in the website data. For example if the Low Hash is in first position, then choose 1 in the Position column.
    g. Click Save
8. Select a import schedule by clicking + on the Schedule Summary. Select when to start the import and how often to import to get new data from the website.
9. The imported data will show on the right pane after some time.

Custom threat feed websites - STIX formatted data and TAXII import

In this case, the threat feed data is available formatted as STIX and follows the TAXII protocol.

1. Select **CMDB>Malware Hash**.
2. Click on the "**+**" button on the left navigation tree to bring up the Create New Malware Hash Group dialog.
3. Enter **Group** and add **Description**. Click **OK** to create the folder under Malware Hash.
4. Select the folder just created.
5. Select Update via API
6. For Website, Click **Add**.
7. In the Data Mapping dialog:
    a. Enter the URL of the website
    b. Enter **User Name** and **Password** (optional)
    c. For Plugin class, choose STIX-TAXII and Full
    d. Click **Save**.
8. Select a import schedule by clicking + on the Schedule Summary. Select when to start the import and how often to import to get new data from the website.
9. The imported data will show on the right pane after some time.

Networks

The CMDB **Networks** page lists the defined networks in your IT infrastructure

Adding a New Network

1. Log in to your Supervisor node.
2. Go to **CMDB > Networks**.
3. Create a new network group or select an existing one.
4. Click **New**.
5. Enter an **Network Name** and **the Low IP address of the network IP range.**
6. Enter any other information about the network.
7. Click **Save**.

Protocols

The CMDB **Protocols** page lists the protocols used by applications and devices to communicate with the FortiSIEM virtual appliance.

Adding a Protocol

1. Log in to your Supervisor node.
2. Go to **CMDB > Protocols**.
3. **Create a new protocol group or select an existing one.**
4. Click **New**.
5. Enter an **Name** and **Description** for the protocol.
6. Click **+** to select a protocol and associate it with a port
7. Select or create an **Apps Group** to associate with the protocol.
8. Click **Save**.

User Agents

The CMDB User Agent page lists common and uncommon user agents in HTTP communications. The traditional use case for a user agent is to detect browser types so the server can return an optimized page. However, user agents are often misused by malware, and are used to communicate the identity of the client to the BotNet controller over HTTP(S). FortiSIEM monitors HTTP(S) logs and the system rule `Blacklist User Agent Match` uses regular expression matching to detect blacklisted user agents.

Adding User Agents

1. Log in to your Supervisor node.
2. Go to **CMDB > User Agents**.
3. Select the User Agent group where you want to add the new user agent.
4. Click **New**.
5. Enter the **User Agent** using regular expression notation**.**
6. Click **Save**.

Adding User Agents to a Group with a CSV File

Instead of manually adding user agents to a user-defined or system group individually, you can upload a CSV file with multiple entries to the group by selecting the group, clicking **Update**, and then selecting **Import Manually**. You will need to format the user agent password with regular expression notation:

```
^Really\s+Bad\s+User\s+Agent
```

Users

The CMDB Users page contains information about users of your system. For more information about adding users, see Adding a Single User.

Watch Lists

A Watch List is a smart container of similar items such as host names, IP addresses, or user names, that are of significant interest to an administrator and need to be watched. Examples of watch lists that are already set up in FortiSIEM are

- **Frequent Account Lockouts** - users who are frequently locked out
- **Host Scanners** - IP addresses that scan other devices
- **Disk space issues** - hosts with disks that are running out of capacity
- **Denied countries** - countries with an excessive number of access denials at the firewall
- **Blacklisted WLAN endpoints** - Endpoints that have been blacklisted by Wireless IPS systems

Typically items are added to a watch list dynamically when a rule is triggered, but you can also add items to a watch list manually. When you define a rule, you can also choose a watch list that will be populated with a specific incident attribute, as described in Adding a Watch List to a Rule, and you can use watch lists as conditions when creating reports, as described in Using Watch Lists as Conditions in Rules and Reports. You can also define when an entry leaves a watch list. Typically this is time based. For example, if the rule does not trigger for that attribute for defined time-period, then the entry is removed from the watch list. Watch lists are also multi-tenant aware, with organization IDs tracked in relation to watch list items.

- Creating a Watch List
- System-Defined Watch Lists

Related Links

- Using Watch Lists as Conditions in Rules and Reports
- Adding a Watch List to a Rule
- Overview of the CMDB User Interface

Creating a Watch List

1. Log in to your Supervisor node.
2. Go to **CMDB > Watch Lists**.
3. Click +.
4. Choose an **Organization** to associate with the watch list.
5. Enter a **Group** name and **Description** for the watch list.
6. Select an object **Type** for the incident attribute that will be saved to the watch list.
7. Select **Case Sensitive** if the object type is **String** and you want to use case sensitivity to compare strings.
8. For **Values Expire in**, set the time period in which items will expire from the watch if there is no activity for that time.
9. Click **OK**.

You can now add your new watch list to a rule, so that when the rule is triggered, items will be added to the watch list. You can also use your watch list as a condition in historical search. See Adding a Watch List to a Rule and Using Watch Lists as Conditions in Rules and Reports for more information.

Related Links

- Adding a Watch List to a Rule
- Using Watch Lists as Conditions in Rules and Reports

System-Defined Watch Lists

FortiSIEM includes several pre-defined watch lists that are populated by system-defined rules.

Watch list	Description	Attribute Type	Triggering Rules
Accounts Locked	Domain accounts that are locked out frequently	User (STRING)	Account Locked: Domain
Application Issues	Applications exhibiting issues	Host Name (STRING)	IIS Virtual Memory Critical SQL Server Low Buffer Cache Hit Ratio SQL Server Low Log Cache Hit Ratio SQL Server Excessive Deadlock SQL Server Excessive Page Read/Write SQL Server Low Free Pages In Buffer Pool SQL Server Excessive Blocking Database Server Disk Latency Critical SQL Server Excessive Full Scan SQL Server scheduled job failed High Oracle Table Scan Usage High Oracle Non-System Table Space Usage Oracle database not backed up for 1 day Exchange Server SMTP Queue High Exchange Server Mailbox Queue High Exchange Server RPC Request High Exchange Server RPC Latency High Oracle DB Low Buffer Cache Hit Ratio Oracle DB Low Library Cache Hit Ratio Oracle DB Low Row Cache Hit Ratio Oracle DB Low Memory Sorts Ratio Oracle DB Alert Log Error Excessively Slow Oracle DB Query Excessively Slow SQL Server DB Query Excessively Slow MySQL DB Query

Watch list	Description	Attribute Type	Triggering Rules
Availability Issues	Servers, networks or storage devices or Applications that are exhibiting availability issues	Host Name (STRING)	Network Device Degraded - Lossy Ping Response Network Device Down - No Ping Response Server Degraded - Lossy Ping Response Server Down - No Ping Response Server Network Interface Staying Down Network Device Interface Flapping Server Network Interface Flapping Important Process Staying Down Important Process Down Auto Service Stopped Critical network Interface Staying Down EC2 Instance Down Storage Port Down Oracle Database Instance Down Oracle Listener Port Down MySQL Database Instance Down SQL Server Instance Down Service Staying Down - Slow Response To STM Service Down - No Response to STM Service Staying Down - No Response to STM
DNS Violators	Sources that send excessive DNS traffic or send traffic to unauthorized DNS gateways	Source IP	Excessive End User DNS Queries to Unauthorized DNS servers Excessive End User DNS Queries Excessive Denied End User DNS Queries Excessive Malware Domain Name Queries Excessive uncommon DNS Queries Excessive Repeated DNS Queries To The Same Domain

Watch list	Description	Attribute Type	Triggering Rules
Denied Countries	Countries that are seeing a high volume of denials on the firewall	Destination Country (STRING)	Excessive Denied Connections From An External Country
Denied Ports	Ports that are seeing a high volume of denies on the firewall	Destination Port (INT)	Excessive Denied Connection To A Port
Environmental Issues	Environmental Devices that are exhibiting issues	Host name (String)	UPS Battery Metrics Critical UPS Battery Status Critical HVAC Temp High HVAC Temp Low HVAC Humidity High HVAC Humidity Low FPC Voltage THD High FPC Voltage THD Low FPC Current THD High FPC ground current high NetBoz Module Door Open NetBotz Camera Motion Detected Warning APC Trap Critical APC Trap
Hardware Issues	Servers, networks or storage devices that are exhibiting hardware issues	Host Name (String)	Network Device Hardware Warning Network Device Hardware Critical Server Hardware Warning Server Hardware Critical Storage Hardware Warning Storage Hardware Critical Warning NetApp Trap Critical Network Trap

Watch list	Description	Attribute Type	Triggering Rules
Host Scanners	Hosts that scan other hosts	Source IP	Heavy Half-open TCP Host Scan Heavy Half-open TCP Host Scan On Fixed Port Heavy TCP Host Scan Heavy TCP Host Scan On Fixed Port Heavy UDP Host Scan Heavy UDP Host Scan On Fixed Port Heavy ICMP Ping Sweep Multiple IPS Scans From The Same Src
Mail Violators	End nodes that send too much mail or send mail to unauthorized gateways		Excessive End User Mail to Unauthorized Gateways Excessive End User Mail
Malware Found	Hosts where malware found by Host IPS /AV based systems and the malware is not remediated	Host Name (String)	Virus found but not remediated Malware found but not remediated Phishing attack found but not remediated Rootkit found Adware process found

Watch list	Description	Attribute Type	Triggering Rules
Malware Likely	Hosts that are likely to have malware - detected by network devices and the determination is not as certain as host based detection	Source IP or Destination IP	Excessive Denied Connections From Same Src Suspicious BotNet Like End host DNS Behavior Permitted Blacklisted Source Denied Blacklisted Source Permitted Blacklisted Destination Denied Blacklisted Destination Spam/malicious Mail Attachment found but not remediated Spyware found but not remediated DNS Traffic to Malware Domains Traffic to Emerging Threat Shadow server list Traffic to Emerging Threat RBN list Traffic to Emerging Threat Spamhaus list Traffic to Emerging Threat Dshield list Traffic to Zeus Blocked IP list Permitted traffic from Emerging Threat Shadow server list Permitted traffic from Emerging Threat RBN list Permitted traffic from Emerging Threat Spamhaus list Permitted traffic from Emerging Threat Dshield list Permitted traffic from Zeus Blocked IP list
Port Scanners	Hosts that scan ports on a machine	Source IP	Heavy Half-open TCP Port Scan: Single Destination Heavy Half-open TCP Port Scan: Multiple Destinations Heavy TCP Port Scan: Single Destination Heavy TCP Port Scan: Multiple Destinations Heavy UDP Port Scan: Single Destination Heavy UDP Port Scan: Multiple Destinations

Watch list	Description	Attribute Type	Triggering Rules
Policy Violators	End nodes exhibiting behavior that is not acceptable in typical Corporate networks	Source IP	P2P Traffic detected IRC Traffic detected P2P Traffic consuming high network bandwidth Tunneled Traffic detected Inappropriate website access Inappropriate website access - multiple categories Inappropriate website access - high volume Inbound clear text password usage Outbound clear text password usage Remote desktop from Internet VNC From Internet Long lasting VPN session High throughput VPN session Outbound Traffic to Public DNS Servers

Watch list	Description	Attribute Type	Triggering Rules
Resource Issues	Servers, networks or storage devices that are exhibiting resource issues: CPU, memory, disk space, disk I/O, network I/O, virtualization resources - either at the system level or application level	Host Name (STRING)	High Process CPU: Server High Process CPU: Network High Process Memory: Server High Process Memory: Network Server CPU Warning Server CPU Critical Network CPU Warning Network CPU Critical Server Memory Warning Server Memory Critical Network Memory Warning Network Memory Critical Server Swap Memory Critical Server Disk space Warning Server Disk space Critical Server Disk Latency Warning Server Disk Latency Critical Server Intf Util Warning Server Intf Util Critical Network Intf Util Warning Network Intf Util Critical Network IPS Intf Util Warning Network IPS Intf Util Critical Network Intf Error Warning Network Intf Error Critical Server Intf Error Warning Server Intf Error Critical

Watch list	Description	Attribute Type	Triggering Rules
			Virtual Machine CPU Warning Virtual Machine CPU Critical Virtual Machine Memory Swapping Warning Virtual Machine Memory Swapping Critical ESX CPU Warning ESX CPU Critical ESX Memory Warning ESX Memory Critical ESX Disk I/O Warning ESX Disk I/O Critical ESX Network I/O Warning ESX Network I/O Critical Storage CPU Warning Storage CPU Critical NFS Disk space Warning NFS Disk space Critical

Watch list	Description	Attribute Type	Triggering Rules
			NetApp NFS Read/Write Latency Warning NetApp NFS Read/Write Latency Critical NetApp CIFS Read/Write Latency Warning NetApp CIFS Read/Write Latency Critical NetApp ISCSI Read/Write Latency Warning NetApp ISCSI Read/Write Latency Critical NetApp FCP Read/Write Latency Warning NetApp FCP Read/Write Latency Critical NetApp Volume Read/Write Latency Warning
			NetApp Volume Read/Write Latency Critical EqualLogic Connection Read/Write Latency Warning EqualLogic Connection Read/Write Latency Critical Isilon Protocol Latency Warning
Routing Issues	Network devices exhibiting routing related issues	Host Name (STRING)	OSPF Neighbor Down EIGRP Neighbor down OSPF Neighbor Down
Scanned Hosts	Hosts that are scanned	Destination IP	Half-open TCP DDOS Attack TCP DDOS Attack Excessive Denied Connections to Same Destination
Vulnerable Systems	Systems that have high severity vulnerabilities from scanners	Host Name (STRING)	Scanner found severe vulnerability
Wireless LAN Issues	Wireless nodes triggering violations	MAC Address (String)	Rogue or Unsecure AP detected Wireless Host Blacklisted Excessive WLAN Exploits Excessive WLAN Exploits: Same Source

Reporting on CMDB Objects

All of the information in the CMDB can be reported on. FortiSIEM includes a number of pre-defined reports that you can run and export to PDF, and you can also create your own reports.

- CMDB Report Types
- Running, Saving, and Exporting a CMDB Report
- Creating and Modifying CMDB Reports
- Importing and Exporting CMDB Report Definitions

CMDB Report Types

You can find all system-defined reports in **CMDB > CMDB Reports**. The reports are organized into folders as shown in this table. Click on a report to view **Summary** information about it, including the report conditions and the columns included in the report.

Report and Organization Associations for Multi-Tenant Deployments

If you have FortiSIEM Service Provider deployment, the Organization column in the CMDB report table will show whether the report is defined for a specific organization. If it is, then that report is available for both the organization and Super/Global users.

CMDB Report Folder	Object to Report On	Report Name
Overall	Device Approval Status	• Approved Devices • Not Approved Devices
	Users	• Discovered Users • Externally Authenticated FortiSIEM Users • Locally Authenticated FortiSIEM Users • Manually Defined Users
	Rules	• Active Rules • Rules with Exceptions
	Reports	• Scheduled Reports
	Performance Monitors	• Active Performance Monitors
	Task	• All Existing Tasks
	Business Service	• Business Service Membership
Network	Inventory	• Network Device Components with Serial Number • Network Interface Report • Router/Switch Inventory • Router/Switch Image Distribution
	Ports	• Network Open Ports
	Relationship	• WLAN-AP Relationship

CMDB Report Folder	Object to Report On	Report Name
Server	Inventory	• Server Inventory • Server OS Distribution • Server Hardware: Processor • Server Hardware: Memory and Storage
	Ports	• Server Open Ports
	Running Services	• Windows Auto Running Services • Windows Auto Stopped Services • Windows Exchange Running Services • Windows IIS Running Services • Windows Manual Running Services • Windows Manual Stopped Services • Windows SNMP Running Services • Windows VNC Running Services • Windows WMI Running Services
	Installed Software / Patches	Windows Installed Software Windows Installed Patches Windows Installed Software Distribution
Virtualization	Relationship	• VM-ESX Relationship

Running, Saving, and Exporting a CMDB Report

1. Log in to your Supervisor node.
2. Go to **CMDB > CMDB Reports**, and select the report you want to run.
3. Click **Run**.
4. If you have a multi-tenant deployment, you will be prompted to select the organizations for which you want to run the report.
5. Click **Save**if you want to save the report.

   Reports are only saved for the duration of your login session, and you can view saved reports by clicking **Report Results**. Each saved report will be listed as a separate tab, and you can delete them by clicking the X that appears when you hover your mouse over the report name in the tab. You can save up to 5 reports per login session

Creating and Modifying CMDB Reports

There are two ways you can create new CMDB reports: you can create a new report from scratch, or you can clone and modify an existing system or user-defined report.

Creating a New Report

1. Log in to your Supervisor node.
2. Go to **CMDB > CMDB Reports**.
3. Create a group to add the new report to if you are not adding it to an existing group.
4. Click **New**.
5. Enter a **Name** and **Description** for the report.
6. Select the **Conditions** for the report.
   You can use parentheses to give higher precedence to evaluation conditions.
7. Select the **Display Columns**.
   The Display Column attributes contain an implicit "group by" command. You can change the order of the columns with the **Move Row: Up** and **Down** buttons.
8. Click **Save**.

Cloning and Modifying a Report

You can modify user-defined reports by selecting the report and clicking **Edit**. However, you cannot directly edit a system-defined report. Instead, you have to clone it, then save it as a new report and modify it.

1. Log in to your Supervisor node.
2. Go to **CMDB > CMDB Reports**.
3. Select the system-defined you want to modify, and then click **Clone**.
4. Enter a name for the new report, and then click **Save**.
   The cloned report will be added to the folder of the original report.
5. Select the new report, and then click Edit.
6. Edit the report, and then click **Save**.

Importing and Exporting CMDB Report Definitions

Instead of using the user interface to define a report, you can import report definitions, or you can export a definition, modify it, and import it back into your FortiSIEM virtual appliance. Report definitions follow an XML schema.

Importing a Report Definition

1. Log in to your Supervisor node.
2. Go to **CMDB > CMDB Reports**.
3. Select the folder where you want to import the report definition, or create a new one.
4. Click **Import**.
5. Copy your report definition into the text field, and then click **Import**.

Exporting a Report Definition

1. Log in to your Supervisor node.
2. Go to **CMDB > CMDB Reports**.
3. Select the report you want to export, and then click **Export**.
4. Click **Copy to Clipboard**.
5. Paste the report definition into a text editor, modify it, and then follow the instructions for importing it back into your virtual appliance.

XML Schema for Report Definitions

The XML schema for the report definition is:

```
<cmd-
bRe-
ports><cmdbReport><name></name><naturalid></naturalid><description></description><selectCl
```

This is an example for the **Active Rules** report:

```
<cmdbReports><cmdbReport><name>Active Rules</name><naturalId>PH_CMDB_Report_Over-
all_8</naturalId><target>com.ph.phoenix.model.query.Rule</target><description>This
report captures active rules on a per organization basis</-
description><selectClause>ph_drq_rule.ph_incident_category,ph_drq_rule.name,ph_
sys_domain.name</selectClause><orderByClause>ph_drq_rule.ph_incident_category
ASC</orderByClause><whereClause>ph_drq_rule.active =
true</whereClause></cmdbReport></cmdbReports>
```

Importing a CMDB Report Definition

1. Go to Report listing page and select the CMDB Report folder where the report is to be imported.
2. Click Import and paste the report into the window
3. Click Import and see the report showing up in the correct folder.

Exporting a CMDB Report Definition

1. Go to Report listing page
2. Select a CMDB Report and click Export
3. Click "Copy to clipboard" and paste it into a file. Click Close after done.

Creating Event Database Archives

Online v. Offline Storage

The FortiSIEM event database, eventDB, is for near-to-intermediate term storage and querying of events. As an online database, eventDB has fast query performance, but this performance comes with a limited storage capacity, and is expensive in terms of resource consumption. For these reasons, data needs to be periodically purged from eventDB and moved into offline storage, but still be available for querying for forensic analysis. FortiSIEM checks the capacity of the online EventDB storage every 30 minutes, and when approaches capacity, begins to move event information, in daily increments, into the offline storage location.

The FortiSIEM virtual appliance includes a data archiving function that enables you to define an offline storage location, and a policy for the number of days that events will be kept in online or offline storage. This archiving function also includes the ability for compliance auditors to validate logs to ensure that they haven't been tampered with in the offline storage. The data is cryptographically signed (SHA256) at the point of entry, and the checksums are stored in the database. The check sums can be re-verified on demand at any point of time, and if the data has been tampered with, then the check sums will not match. The data integrity reports can be exported in PDF format. If the events in offline storage need to be queried at some point in the future, they can be restored to the FortiSIEM virtual appliance.

Checking Online and Offline Storage Storage Consumption

You can check the amount of storage required for both your online data and your offline archive under the **Event DB Management > Data Manager** tab.

Setting Purge and Archive Policies

Online data is only moved to the archive location when online storage reaches capacity. When you set the archive policy as described in Setting Up an Event Data Archive and Archive Policy, you are setting the amount of time that archived data will be retained before it is purged. For example, if you set the **Data Management Policy** for your deployment or an organization to **90 days**, then maintenance will run every day to purge data that is over 90 days old. If there is not enough offline storage for 90 days, then archived events will be purged from offline storage to create more capacity. If there is enough storage for the 90 days, then events will only be purged after 90 days. For this reason it is very important that you set an archive location that has sufficient capacity to store the amount of data for the number of days that you specify.

For Service Provider deployments, you can set archive policies for each organization. If one organization requires 30 days of storage, and another customer requires 90 days of storage, then FortiSIEM will attempt to enforce these policies in relation to the amount of storage available. For the first organization, events will be deleted from the archive storage location on the 31st day to free up capacity for the organization that has longer storage requirements.

As with the online EventDB data, every 30 minutes FortiSIEM will check the capacity of the offline archive storage, and when the remaining storage capacity reaches a 20GB threshold, it will begin to purge data from the

archive location, beginning with the oldest data, and purging it in daily increments, until the remaining storage capacity is above 20GB.

Archive and Purge Alerts

There are several system alerts that are related to eventDB capacity and the archiving function:

Alert	Description
Online event database close to full (below 20GB)	When the database reaches a point where the remaining storage capacity is below 20GB, its contents will be purged or archived, depending on whether an archive storage location has been defined
Event Archive started	The archive process has been initiated
Event Archive failed	The archive process has failed, likely due to a lack of capacity in the offline storage location
Event Archive purged because of archive purging policy	The contents of the event archive have been purged from offline storage according to the archive purging policy
Event Archive purged because it is full	The contents of the event archive have been purged from offline storage due to capacity issues

Related links

- Setting Up an Event Data Archive and Archive Policy
- Restoring Archived Data
- Validating Log Integrity

Managing Event Data Archive

Prerequisites

- Make sure you read the section on **Setting Archive and Purge Policies** in the topic Creating Event Database Archives before you set up your policy. It is very important that you understand how FortiSIEM moves data into the archive, and purges archived data when the archive destination storage reaches capacity, before you create your policy.
- Make sure that your **Archive Destination** has sufficient storage for your event data + 20GB. When the archive storage reaches 20GB of capacity, FortiSIEM will begin to purge archived data, in daily increments, starting with the oldest data, to maintain a 20GB overhead.

Creating Archive Destination

1. Log in to your Supervisor node.
2. Go to **Admin > Event DB Management**.
3. Click **Retention Policy**.
4. For **Archive Destination**, enter the full path of the file system directory where you want your event data to be archived, and then click **Apply**.

   Offline Storage Capacity for Multi-Tenant Deployments

   Note that all organizations will share the same **Archive Destination**. For this reason, you should make sure that the archive destination has enough capacity to hold the event data for both the number of organizations and the archive retention period that you set for each. If the archive destination does not have enough storage capacity, the archive operation may fail.

Creating Offline (Archive) Retention Policy

This enables you to control which customers data stays in event data archive and for how long.

1. Log in to your Supervisor node.
2. Go to **Admin > Event DB Management**.
3. Click **Retention Policy**.
4. Under **Offline Retention Policies**, click **New**.
5. For multi-tenant installations, select the **Organization** for which this policy will apply.
6. For **Time Period**, enter the number of days that event data should be held in the offline storage before it is purged.
7. Click **Save**.

Managing Online Event Data

Creating Online Event Retention Policy

This enables you to control the content of online event data.

1. Log in to your Supervisor node.
2. Go to **Admin > Event DB Management** .
3. Click **Retention Policy** .
4. Under **Online Retention Policies** , click **Add** .
5. Enter the following information
   a. **Enabled** - Check this box if the policy has to be enforced right away.
   b. **Organizations** - Choose the organizations for which the policy has to be applied (for Service Provide installs)
   c. **Reporting Devices** - Choose the reporting devices relevant to this policy
   d. **Event Type** - Choose the event types or event type groups
   e. **Time period** - enter the number of days that event data specified by the conditions (Organizations, Reporting Devices and Event Type) should be held in the online storage before it is moved to archive or purged.
   f. **Description** - enter a description for the policy
6. Click **Save** .

Viewing Online Event Data Usage

This enables you to see a summarized view of online event data. These views enables you to manage storage more effectively by writing appropriate event dropping policies or online event retention policies.

- **Calendar View** - This view shows you how much storage is used by each organization on a day by day basis

Online Data (Calendar): 4.71 GB

(1 of 1)

Organization	Size
▼ Super	4.71 GB
▼ Year 2017	4.71 GB
▼ February	4.71 GB
February 13	415.70 MB
February 14	458.85 MB
February 15	595.31 MB
February 16	620.57 MB
February 17	616.48 MB
February 18	646.99 MB
February 19	617.59 MB
February 20	300.16 MB
February 21	550.50 MB
February 22	4.02 MB

- **Top (Event Type, Reporting Device) View** - This view shows you the top (Event Type, Reporting Device) tuples consuming most storage for each organization on a month-by-month basis

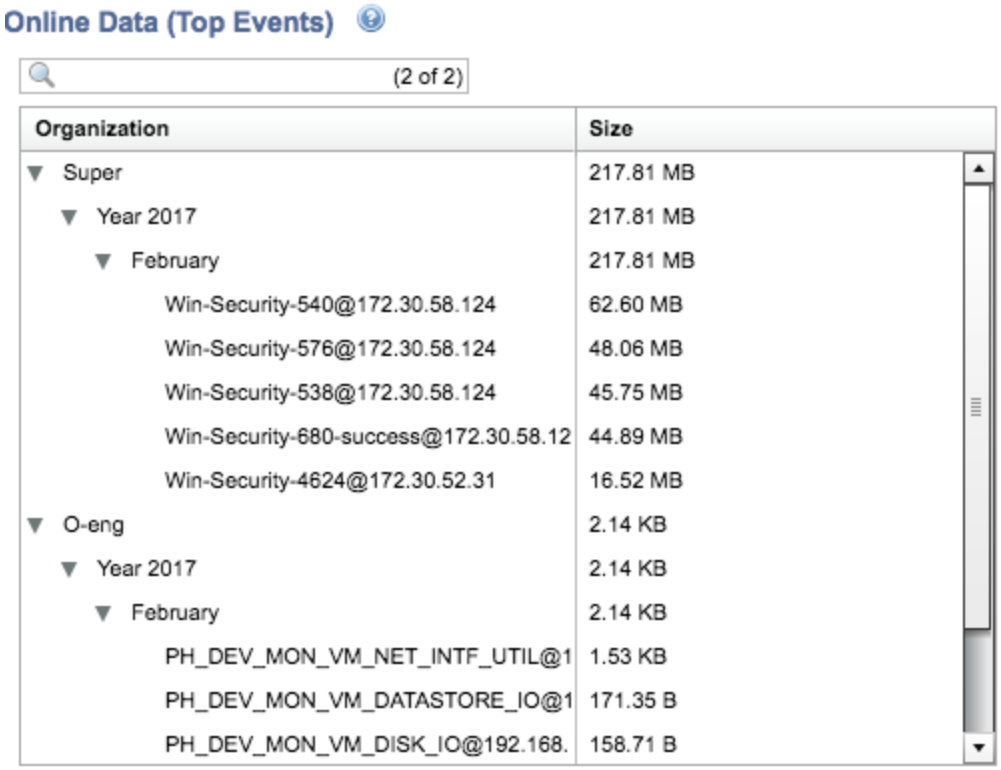

Restoring Archived Data

Once your event data has been moved to an offline archive, you can no longer query that data from within FortiSIEM. However, you can restore it to your virtual appliance, and then proceed with any queries or analysis.

1. Log in to your Supervisor node.
2. Go to **Admin > Event DB Management > Data Manager**.
3. Under **Reserved Restore Space (GB)**, enter the amount of storage space that will be reserved for the restored data.
   This should be equal to or larger than the size of the archive to be restored.
4. Under **Archived Data**, select the archive that you want to restore.
5. Click **Restore**.
   The archive data will be moved to the restore space and can be queried in the usual ways.

Validating Log Integrity

1. Security auditors can validate that archived event data has not been tampered with by using the **Event Integrity** function of **Event DB Management**.
2. Log in to your Supervisor node.
3. Go to **Admin > Event DB Management > Event Integrity**.
4. Select the **Begin Time** and **End Times** for the time period during which log integrity needs to be validated.
5. Click **Show**.
   You will see a table of all the logs that are available for the specified time period
6. Use **Validation Status** to filter the types of logs you want to validate.
7. Select the log you want to validate, and click **Validate**.
   A table showing the validation status of logs will be displayed.

Column	Description
Start Time	The earliest time of the messages in this file. The file does not contain messages that were received by FortiSIEM before this time.
End Time	The latest time of the messages in this file. The file does not contain messages that were received by FortiSIEM after this time.
Category	• **Internal**: these messages were generated by FortiSIEM for its own use. This includes FortiSIEM system logs and monitoring events such as the ones that begin with `PH_DEV_MON`. • **External**: these messages were received by FortiSIEM from an external system • **Incident**: these corresponds to incidents generated by FortiSIEM
File Name	The name of the log file
Event Count	The number of events in the file
Checksum Algorithm	The checksum algorithm used for computing message integrity
Message Checksum	The value of the checksum
Validation Status	• **Not Validated:** the event integrity has not been validated yet • **Successfu**l: the event integrity has been validated and the return was success. This means that the logs in this file were not altered. • **Failed**: the event integrity has been validated and the return was failed. This means that the logs in this file were altered. • **Archived**: the events in this file were archived to offline storage
File Location	• **Local: l**ocal to Supervisor node • **External**: means external to Supervisor node, for example on NFS storage

8. Click **Export** to create a PDF version of the validation results.

Integrating with External CMDB and Helpdesk Systems

Topics in this section include:

FortiSIEM CMDB/Helpdesk System Integration Overview

FortiSIEM integration helps to create a two-way linkage between external ticketing/work flow systems like ServiceNow, ConnectWise and Salesforce. The integration can be for Incidents and CMDB.

This involves two steps:

1. Create an integration.
2. Attach the integration to an Incident Notification Policy or run the integration on a schedule.

Four types of integrations are supported:

- **Incident Outbound Integration**: This creates a ticket in an external ticketing system from FortiSIEM incidents.
- **Incident Inbound Integration**: This updates FortiSIEM incident ticket state from external system ticket states. Specifically, when a ticket is closed in the external ticketing system, the incident is cleared in FortiSIEM and the ticket status is marked closed to synchronize with the external ticketing system.
- **CMDB Outbound Integration**: This populates an external CMDB from FortiSIEM CMDB.
- **CMDB Inbound Integration**: This populates FortiSIEM CMDB from an external CMDB.

FortiSIEM provides a Java based API that can be used to integrate with ticketing systems. Out of the box integration is available for ServiceNow, ConnectWise and Salesforce. Integration with other systems can be built using the API. Contact Fortinet support for assistance.

Configuring external helpdesk systems for FortiSIEM integration

This section specifies how to configure the out of the box external ticketing systems for FortiSIEM integration.

Configuring ServiceNow

1. Login to ServiceNow.
2. For Service Provider Configurations, create Companies by creating Company Name.

Configuring ConnectWise

1. Login to ConnectWise MANAGE.
2. Go to Setup Tables > Integrator Login List.
3. Create a new **Integrator Login** for FortiSIEM:
    a. Enter **Username**.
    b. Enter **Password**.
    c. Set **Access Level** to **Records created by integrator**.
    d. Enable **Service Ticket API** for Incident Integration.
    e. Enable **Configure API** for CMDB Integration.
4. For Service Provider Configurations, create Companies by creating:
    a. **Company Name**
    b. **Company ID**

Configuring Salesforce for FortiSIEM Integration

1. Login to Salesforce.
2. Create a **custom domain**.
3. For Service Provider Configurations, create **Service App** > **Accounts**.
   FortiSIEM will use the **Account Name**.

Incident Outbound Integration

This creates a ticket in an external ticketing system when an incident triggers in FortiSIEM incidents. Built-in integrations are available for ServiceNow, ConnectWise and Salesforce.

The steps are:

1. Create an Incident Outbound integration.
2. Link the integration to one or more Incident Notification Policies.

When an incident triggers, the notification policy will be invoked and a ticket will be created in the external system.

Create an Incident Outbound integration

1. Log into your Supervisor node with administrator credentials.
2. Go to **Admin > General Settings > Integration**.
3. Click **Add**.
4. For **Type**, select **Device**.
5. For **Direction**, select **Outbound**.
6. For **Vendor**, select the vendor of the system you want to connect to. ServiceNow is supported out of the box. When you select the Vendor:
    a. An **Instance** is created - this is the unique name for this policy. For example if you had 2 ServiceNow installations, each would have different Instance names.
    b. A default **Plugin Name** is populated - this is the Java code that implements the integration including connecting to the external help desk systems and synching the CMDB elements. The plugin is automatically populated for ServiceNow and ConnectWise. For other vendors, you have to create your own plugin and type in the plugin name here.
7. For **Host/URL**, enter the host name or URL of the external system (see section Configuring external helpdesk systems)
    a. For ServiceNow, select the login URL
    b. For ConnectWise, select the login URL.
    c. For Salesforce:
        i. Login to Salesforce.
        ii. Go to **Setup** > **Settings**.
        iii. Use the **Custom URL** under **My Domain**, typically it is xyz.my.salesforce.com
8. For **User Name** and **Password**, enter a user name and password that the system can use to authenticate with the external system.
    a. For ServiceNow, select the login credentials.
    b. For ConnectWise, select the credentials created in Step 3.
    c. For Salesforce, select the login credentials.
9. For **Incidents Comments Template**, specify the formatting of the incident fields.
10. For **Org Mapping**, click **Edit** to create mappings between the organizations in your FortiSIEM deployment and the names of the organization in the external system.
    a. For ServiceNow, select the Company names as in Step 2.
    b. For ConnectWise, select the Company name in Step 4.

c. For Salesforce:
    a. Go to **Service App** > **Accounts**.
    b. Use **Account Name**.

11. For **Run For**, choose the organizations for whom tickets will be created.
12. Click **Save**.

Link the integration to one or more Incident Notification Policies

1. Log into your FortiSIEM Supervisor with administrator credentials.
2. Go to **Incidents** > **Incident Notification Policy**.
3. Create a new policy or edit an existing policy.
4. Select **Actions** > **Invoke an Integration Policy** and choose a specific integration.
5. Click **Save**.

Incident Inbound Integration

This updates the FortiSIEM incident state and clears the incident when the incident is cleared in the external help desk system. Built-in integrations are available for ServiceNow, ConnectWise and Salesforce.

The steps are:

1. Create an Incident Inbound integration schedule.
2. Create a schedule for automatically running the Incident Inbound integration.

This will update the FortiSIEM incident inbound integration schedule and clears the incident when the incident is cleared in the external help desk system.

Create an Incident Inbound integration

1. Log into your Supervisor node with administrator credentials.
2. Go to **Admin > General Settings > Integration**.
3. Click **Add**.
4. For **Type**, select **Incident**.
5. For **Direction**, select **Inbound**.
6. For **Vendor**, select the vendor of the system you want to connect to. ServiceNow is supported out of the box. When you select the Vendor:
    a. An **Instance** is created - this is the unique name for this policy. For example if you had 2 ServiceNow installations, each would have different Instance names.
    b. A default **Plugin Name** is populated - this is the Java code that implements the integration including connecting to the external help desk systems and synching the CMDB elements. The plugin is automatically populated for ServiceNow and ConnectWise. For other vendors, you have to create your own plugin and type in the plugin name here
7. For **Host/URL**, enter the host name or URL of the external system (see section Configuring external helpdesk systems).
    a. For ServiceNow, select the login URL.
    b. For ConnectWise, select the login URL.
    c. For Salesforce:
        i. Login to Salesforce.
        ii. Go to Setup > Settings
        iii. Use the **custom URL** under **My Domain** – typically it is xyz.my.salesforce.com
8. For **User Name** and **Password**, enter a user name and password that the system can use to authenticate with the external system.
    a. For ServiceNow, select the login credentials
    b. For ConnectWise, select the credentials created in Step 3
    c. For Salesforce, select the login credentials.
9. For **Time Window**, select the number of hours for which incident states will be synched. For example, if time windows is set to 10 hours, the states of incidents that occurred in the last 10 hours will be synched.
10. Click **Save**.

Create an Incident Inbound integration schedule

This will update FortiSIEM following incident fields when ticket state is updated in the external ticketing system.

- External Ticket State
- Ticket State
- External Cleared Time
- External Resolve Time

1. Log into your FortiSIEM Supervisor with administrator credentials.
2. Go to Admin > General Settings > Integration.
3. Click Schedule and then click **+**.
    a. Select the integration policy.
    b. Select a schedule.

CMDB Outbound Integration

CMDB Outbound Integration populates an external CMDB from FortiSIEM's own CMDB. Built in integrations are available for ServiceNow, ConnectWise and Salesforce.

The steps are:

1. Create a CMDB Outbound integration.
2. Create a CMDB Outbound integration schedule.

Create a CMDB Outbound integration

1. Log into your Supervisor node with administrator credentials.
2. Go to **Admin > General Settings > Integration**.
3. Click **Add**.
4. For **Type**, select **Device**.
5. For **Direction**, select **Outbound**.
6. For **Vendor**, select the vendor of the system you want to connect to. ServiceNow is supported out of the box. When you select the Vendor:
    a. An **Instance** is created - this is the unique name for this policy. For example if you had 2 ServiceNow installations, each would have different Instance names.
    b. A default **Plugin Name** is populated - this is the Java code that implements the integration including connecting to the external help desk systems and synching the CMDB elements. The plugin is automatically populated for ServiceNow and ConnectWise. For other vendors, you have to create your own plugin and type in the plugin name here.
7. For **Host/URL**, enter the host name or URL of the external system (see section Configuring external helpdesk systems)
    a. For ServiceNow, select the login URL
    b. For ConnectWise, select the login URL.
    c. For Salesforce:
        i. Login to Salesforce.
        ii. Go to **Setup** > **Settings**.
        iii. Use the **Custom URL** under **My Domain**, typically it is xyz.my.salesforce.com
8. For **User Name** and **Password**, enter a user name and password that the system can use to authenticate with the external system.
    a. For ServiceNow, select the login credentials.
    b. For ConnectWise, select the credentials created in Step 3.
    c. For Salesforce, select the login credentials.
9. Enter the **Maximum** number of devices to send to the external system.
10. For **Org Mapping**, click **Edit** to create mappings between the organizations in your FortiSIEM deployment and the names of the organization in the external system.
    a. For ServiceNow, select the Company names as in Step 2.
    b. For ConnectWise, select the Company name in Step 4.
    c. For Salesforce:
        a. Go to **Service App** > **Accounts**.
        b. Use **Account Name**.
11. For **Run For**, choose the organizations for whom tickets will be created.

12. For **Groups**, select the FortiSIEM CMDB Groups whose member devices would be synched to external CMDB.
13. For ConnectWise, it is possible to define a Content Mapping.
    a. Enter **Column Mapping** values:
        i. To add a new mapping, click on the + button.
        ii. Choose FortiSIEM CMDB attribute as the Source Column.
        iii. Enter external (ConnectWise) attribute as the Destination Column.
        iv. Specify Default Mapped Value as the value assigned to the Destination Column if the Source Column is not found in Data Mapping definitions.
        v. Select Put to a Question is the Destination Column is a custom column in ConnectWise.
    b. Enter **Data Mapping** values:
        i. Choose the (Destination) Column Name.
        ii. Enter From as the value in FortiSIEM.
        iii. Enter To as the value in ConnectWise.
14. Select **Run after Discovery** if you want this export to take place after you have run discovery in your system. This is the only way to push automatic changes from FortiSIEM to the external system.
15. Click **Save**.

Create a CMDB Outbound integration schedule

Updating external CMDB automatically after FortiSIEM discovery

1. Create an integration policy.
2. Make sure Run after Discovery is checked.
3. Click **Save.**

Updating external CMDB on a schedule

1. Log into your FortiSIEM Supervisor with administrator credentials.
2. Go to Admin > General Settings > Integration.
3. Click **Schedule** and then click **+**.
    a. Select the integration policies.
    b. Select a schedule.

Updating external CMDB on-demand (one-time)

1. Log into your FortiSIEM Supervisor with administrator credentials.
2. Go to Admin > General Settings > Integration.
3. Select a specific integration policy and click **Run**

CMDB Inbound Integration

CMDB Inbound Integration populates FortiSIEM CMDB from an external CMDB.

The steps are:

1. Create a CMDB Inbound integration.
2. Create a CMDB Inbound integration schedule.

Create a CMDB Inbound integration

You need to have created a CSV file for mapping the contents of the external database to a location on your FortiSIEM Supervisor, which will be periodically updated based on the schedule you set.

1. Log into your Supervisor node with administrator credentials.
2. Go to **Admin > General Settings > Integration**.
3. Click **Add**.
4. For **Type**, select **Device**.
5. For **Direction**, select **Inbound**.
6. Select the **Vendor** of the external system you want to connect to.
7. Enter the **File Path** to the CSV file.
8. For **Column Mapping**, click **+** and enter the mapping between columns in the **Source** CSV file and the **Destination** CMDB.
   For example, if the source CSV has a column **IP**, and you want to map that to the column **Device IP** in the CMDB, you would enter **IP** for **Source Column**, and select **Device IP** for **Destination Column**.
   a. Enter Source CSV column Name for **Source Column**
   b. Check **Create Property if it Does not Exist** to create the new attribute in FortiSIEM if it does not exist
      i. Enter a name for the **Destination Column** of the property in the CMDB.
      ii. Select a **Property type**.
      iii. Enter the **Display Name** for the property.
      iv. Select **Overwrite Existing Value** if the property exists in the CMDB, but you want to overwrite it's current value.
   c. If the property exists in the CMDB, select FortiSIEM CMDB attribute for **Destination Column**.
   d. Select **Overwrite Existing Value** if the property exists in the CMDB, but you want to overwrite its current value.
   e. Click **OK**.
9. For **Data Mapping**, click **+** and enter the mapping between data values in the external system and the destination CMDB.
   For example, if you wanted to change all instances of **California** in the entries for the **State** attribute in the external system to **CA** in the destination CMDB, you would select the **State** attribute, enter **California** for **From**. and **CA** for **To**.
10. Click **OK**.
11. Click **Save**.

Create a CMDB Inbound integration schedule

Updating FortiSIEM CMDB on a schedule

1. Log into your FortiSIEM Supervisor with administrator credentials.
2. Go to Admin > General Settings > Integration.
3. Click **Schedule** and then click **+**.
   a. Select the integration policies.
   b. Select a schedule..

Updating FortiSIEM CMDB on-demand (one-time)

1. Log into your FortiSIEM Supervisor with administrator credentials.
2. Go to Admin > General Settings > Integration.
3. Select a specific integration policy and click **Run**.

Exporting Events to External Systems via Kafka

This section describes procedures for exporting FortiSIEM events to an external system via the Kafka message bus.

Prerequisites

- Make sure you have set up a Kafka Cloud (here) with a specific Topic for FortiSIEM events.
- Make sure you have identified a set of Kafka brokers that FortiSIEM is going to send events to.
- Make sure you have configured Kafka receivers which can parse FortiSIEM events and store in a database. An example would be Logstash receiver (see here) that can store in a Elastic Search database.
- Supported Kafka version: 0.8

Procedure

1. Go to **Admin** > **General Settings** > **Kafka Configuration**.
2. Select **Enable Kafka**.
3. Select a **Topic**.
4. Go to `/etc/hosts` file in the command line and add the IP address and host name of brokers.
5. Add **Brokers** by clicking on **+** icon.
    - Enter IP address or Host name of the broker.
    - Enter Broker port (default 9092).
6. Click **Save**.

Note: Enter multiple broker addresses for redundancy. If one broker is not available, FortiSIEM is going to try the next broker in the list. The full list of brokers does not need to be specified.

Backing Up and Restoring FortiSIEM Directories and Databases

Backing Up and Restoring SVN

Backup and restore SVN

FortiSIEM uses an inbuilt SVN to store network device configuration and installed software versions.

Backup

The SVN files are stored in `/data/svn`. Copy the entire directory to another location.

```
# cd /data
# cp -r svn /<another>/<mount>/<point>
```

Restore

Copy the entire `/data/svn` from the backup location and rename the directory to `/data/svn`.

```
# cd /<another>/<mount>/<point># cp -r svn /data
```

Backing Up and Restoring the CMDB

The FortiSIEM Configuration Management Database (CMDB) contains discovered information about devices, servers, networks and applications. You should create regular backups of the CMDB that you can use to restore it in the event of database corruption.

Backup

The database files are stored in `/data/cmdb/data`. FortiSIEM automatically backs up this data twice daily and the backup files are stored in /data/archive/cmdb. To perform a backup, move these files to another location. For example:

```
[root@SaaS-Sup cmdb] #cd /data/archive/cmdb
[root@SaaS-Sup cmdb] #cp phoenixdb* /<another>/<mount>/<point>
```

If your `/data` disk is on an external NFS mount then your CMDB backup is already separate from the VM infrastructure.

```
[root@SaaS-Sup cmdb]# pwd
/data/archive/cmdb
[root@SaaS-Sup cmdb]# ls -lt
total 1213952
-rw-rw-rw- 1 root root 95559457 Apr 20 03:02 phoenixdb_2011-04-20T03-00-01
-rw-rw-rw- 1 root root 93010144 Apr 19 13:04 phoenixdb_2011-04-19T13-00-02
-rw-rw-rw- 1 root root 91142941 Apr 19 03:02 phoenixdb_2011-04-19T03-00-01
-rw-rw-rw- 1 root root 89686080 Apr 18 13:03 phoenixdb_2011-04-18T13-00-02
```

Restore

If your database becomes corrupted, restore it from backup by performing these steps on you Supervisor node.

1. Stop all processes with this phTools command:

   ```
   #phtools --stop all
   ```

2. Check that all processes have stopped.

   ```
   #phstatus
   ```

   These processes will continue to run, which is expected behavior:

   ```
   phMonitor      1-01:55:17      0              992m          540m
   Apache         1-01:56:45      0              236m          9720
   AppSvr         1-01:56:35      0              3908m         758m
   DBSvr          1-01:57:06      0              383m          6656
   ```

3. Copy the latest `phoenixdb_<timestamp>` file to a directory like `/tmp` on the Supervisor host.
4. Go to `/opt/phoenix/deployment`.
5. Run `db_restore /tmp/phoenixdb_<timestamp>`.
6. When this process completes, reboot the system.

   ```
   #reboot
   ```

Backing Up and Restoring the Event Database

Backup

The event data is stored in `/data/eventdb`. Since this data can become very large over time, you should use a program such as rsync to incrementally move the data to another location. From version 4.2.1 the rsync program is installed on FortiSIEM by default.

Use this command to back up the eventdb.

```
#rsync -a --status /data/eventdb /<another>/<mount>/<point>
```

Restore

To restore eventdb there are two options:

- Mount the directory where the event database was backed up.
- Copy the backup to the **/data/eventdb** directory.

These instructions are for copying the backup to the **/data/eventdb** directory.

1. Stop all running processes.

   ```
   #phtools --stop all
   ```

2. Check that all processes have stopped.

   ```
   #phstatus
   ```

   You will see that these processes are still running, which is expected behavior.

   ```
   phMonitor      1-01:55:17      0              992m          540m
   Apache         1-01:56:45      0              236m          9720
   AppSvr         1-01:56:35      0              3908m         758m
   DBSvr          1-01:57:06      0              383m          6656
   ```

3. Copy the the event DB to the event DB location `/data/eventdb` If you use the `cp` command it may appear that the command has hung if there is a lot of data to copy

   ```
   #cp -a /backup/eventdb /data/eventdb
   ```

   Alternatively you can use `rsync` and display the process status.

   ```
   #rsync -a --status /backup/eventdb /data/eventdb
   ```

4. Once complete, restart all processes.

   ```
   #phtools --start all
   ```

   Check that all processes have started.

   ```
   #phstatus
   ```

Monitoring Operations with FortiSIEM

This chapter describes the following:

Dashboards - Flash version

FortiSIEM includes several different types of dashboards and views to monitor your IT infrastructure. Topics in this section provide an overview of the **General** and **VM View** dashboards available in the **Dashboard** tab, along with their user interface controls and customization options.

Dashboard Overview

FortiSIEM includes two types of component dashboards: **General**, which are used to monitor IT infrastructure components, and **VM View**, which focus specifically on information about virtual machines in your infrastructure. These two types of component dashboards also include two types of dashboads for collecting different types of information:

- **Summarydashboards** that provide single-line entries for IT infrastructure components based on their system status (**Critical**, **Criitcal + Warning**, **All**) in operational time
- **Widget-based dashboards** that provide metrics and analytics for functional areas using historical data

In addition to the summary and widget-based dashboards, FortiSIEM also includes a specialized **Incident dashboard**, with features that are detailed in the Incidents - Flash version section.

Topics in this section provide an overview of the Summary and Widget dashboards, as well as how to use the Analysis menu to gain more information about your IT infrastructure components.

- Summary Dashboard User Interface Overview
- VM Dashboard User Interface Overview
- Widget Dashboard User Interface Overview
- Network Topology View of Devices
- How Values in Dashboard Columns are Derived
- Using the Analysis Menu

Summary Dashboard User Interface Overview

Dashboard Overview

Summary dashboards are best used for gathering information about individual infrastructure components in operational time. Summary dashboards include the **Exec Summary** dashboard, and all the dashboards in the **Summary Dashboards** and **Availability/Performance** folders of the **Dashboards > General** pane. In the **Dashboards > VM View pane**, summary dashboards include the **ESX Host Type** dashboards (**All ESX Hosts** and **Standalone ESX Hosts**, for example). Metrics for these dashboards are displayed either on a real-time basis, or as an average of ten minute intervals.

This screenshot shows an example of a **Biz Service Summary** dashboard for a multi-tenant deployment. It contains all the standard user interface controls found in summary dashboard, though some additional UI controls are found in other summary dashboards as described in the table **Columnar Dashboard UI Controls**. Selecting a business service in the top pane loads all the components associated with that service into the panes below.

Summary Dashboard UI Controls

UI Control	Description
Status Filter	Filters the view of the components based on component status: **Critical**, **Critical + Warning**, **All**
Organizations Filter	For multi-tenant deployments, filter components based on the organization they belong to
Service Info	For the **Business Services** summary dashboard, shows the **Quick Info** for the business service. For other components, an **Info** link is provided in the same location in the UI.
Analysis Menu	The Analysis menu contains a number of options for component analytics, depending on the component selected. See Using the Analysis Menu for more information. You can also access the Analysis menu for a component by hovering your mouse over the component's **Device IP** menu until the blue **Quick Info** icon appears, and then clicking the icon.
Customize Columns	The **Custom Columns** control lets you change the columns that are displayed in the dashboard. See Adding Custom Columns to Dashboards for more information.
Performance Summaries	Most columns contain a summary or trend view of their display information. Hover your mouse over the metric until a trend line icon appears, and then click to view the summary or trend information. Note that many of these summary pop-ups have their own navigational controls, for example to set the time interval for the summary.
Incident Summary	The incident summary shows the number and type of incidents associated with the component. Hover over the number to view a quick summary of the incidents, click on the incident number to view incident details.

UI Control	Description
Quick Info	The **Quick Info** view of a device, which you can also access through the **Analysis** menu or hovering your mouse cursor over the **Device IP** column, displays **General and Health information** for the device, and when appropriate, **Identity and Location** information. It also contains links to additional information about the device: • **Incidents** An exportable summary of incidents associated with the device • **Health** **Availability**, **Performance**, and **Security** health information for the device**.** You can also access this information by clicking the **Device Health** user interface control, or by selecting **Device Health** in the **Analysis** menu. • **BizService** Any business services impacted by the device. You can also access this information by selecting **Impacted Business Services** in the **Analysis** menu. • **Applications** Displays a report on the top 10 applications associated with the device by **Average CPU Utilization** over the past hour • **Vulnerability and IP Status** (Not used in the **Dashboard** view) Displays the vulnerability status reports that are also available by selecting **Vulnerability and IPS Status** in the **Analysis** menu • **Hardware Health** (Used only for the **CMDB/Storage** view) Displays health information for the hardware being used for storage • **Interfaces** Displays a report on the top 10 interfaces associated with the device by average throughput • **Topology** Shows the device's location in the network topology**.** You can also access this information by selecting **Topology** in the **Analysis** menu. The Quick Info view also contains two links, **Goto Config Item**, which links to the device entry in the CMDB, and **Goto Identity**, which links to **Analytics > Identity and Location Report**, where you can edit this information for the device.
Component Health	**Availability** , **Performance** , and **Security** health reports for the device. You can also access this information by selecting a device in the Summary dashboard, and then click **Health,** or by going to **Quick Info > Health** after selecting the device . If any **Incidents** are displayed, click the number to view the **Incident Summary** . Depending on the reported metric, you can zoom in for a closer look at graphs and reports by clicking the **Magnifying Glass** icon that appears when you hover your mouse cursor over them.
Location Selection	Filters components by their geographic locations. See Setting Device Location Information for more information.
Time View and Refresh Interval	The Time View has two options for whether you want to view **Real Time** or **Average-10 mins** metrics for your component, and for the interval and which you want them to refresh.{to

VM Dashboard User Interface Overview

The **Dashboard > VM View** provides a complete overview of your virtual infrastructure, including Data Centers, Standalone ESX Hosts, Resource Pools, Clusters, ESXs, and VMs. Over 400 VMs can be discovered, and their metrics pulled via VCenter in under three minutes during initial discovery. As you navigate the **Virtual Infrastructure** hierarchy, you will see Summary dashboards similar to those in the **General > Dashboard** view for **VM Clusters**, **All ESX Hosts**, and **Standalone ESX Hosts**, while widget dashboards that provide performance metrics for CPU Utilization, Memory, Network Interface, Disk I/O and Data Store Utilization are available at the level of **VM**, **ESX**, **Resource Pool** and **Cluster**.

- VM Summary Dashboards Overview
- The ESX Hosts View
- The ESX and VM View

VM Summary Dashboards Overview

This screenshot shows the **All ESX Hosts** summary dashboard, which includes a summary pane for All ESXs at the top, and a summary pane for individual VM instances for selected ESXs at the bottom. The user interface controls for the Virtual Infrastructure summary dashboards are very similar to those in the **General** summary dashboards.

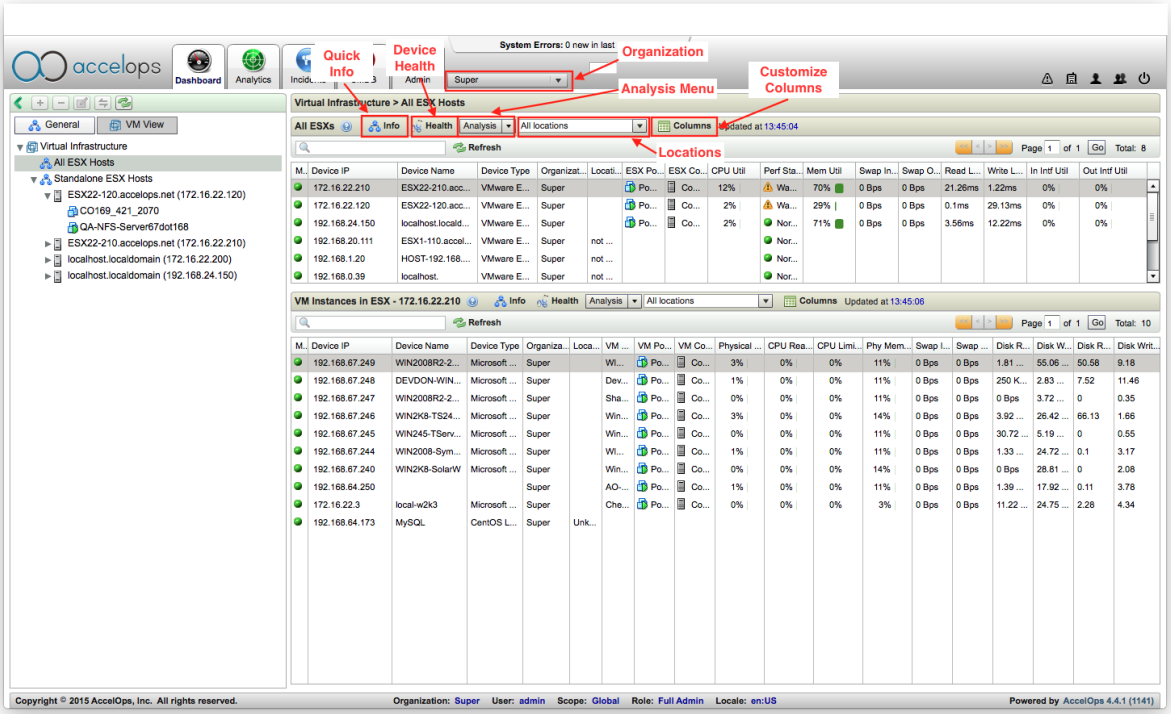

UI Controls for Virtual Infrastructure Summary Dashboards

Ui Control	Description
Organizations Filter	For multi-tenant deployments, filter components based on the organization they belong to
Quick Info	The **Quick Info** view of a device, which you can also access through the **Analysis** menu or hovering your mouse cursor over the **Device IP** column, displays **General and Health information** for the device, and when appropriate, **Identity and Location** information. It also contains links to additional information about the device: • **Incidents** An exportable summary of incidents associated with the device • **Health** **Availability**, **Performance**, and **Security** health information for the device**.** You can also access this information by clicking the **Device Health** user interface control, or by selecting **Device Health** in the **Analysis** menu. • **BizService** Any business services impacted by the device. You can also access this information by selecting **Impacted Business Services** in the **Analysis** menu. • **Applications** Displays a report on the top 10 applications associated with the device by **Average CPU Utilization** over the past hour • **Vulnerability and IP Status** (Not used in the **Dashboard** view) Displays the vulnerability status reports that are also available by selecting **Vulnerability and IPS Status** in the **Analysis** menu • **Hardware Health** (Used only for the **CMDB/Storage** view) Displays health information for the hardware being used for storage • **Interfaces** Displays a report on the top 10 interfaces associated with the device by average throughput. • **Topology** Shows the device's location in the network topology**.** You can also access this information by selecting **Topology** in the **Analysis** menu. The Quick Info view also contains two links, **Goto Config Item**, which links to the device entry in the CMDB, and **Goto Identity**, which links to **Analytics > Identity and Location Report**, where you can edit this information for the device.

Ui Control	Description
Device Health	**Availability** , **Performance** , and **Security** health reports for the device. You can also access this information by selecting a device in the Summary dashboard, and then click **Health,** or by going to **Quick Info > Health** after selecting the device . If any **Incidents** are displayed, click the number to view the **Incident Summary** . Depending on the reported metric, you can zoom in for a closer look at graphs and reports by clicking the **Magnifying Glass** icon that appears when you hover your mouse cursor over them.
Analysis Menu	The Analysis menu contains a number of options for component analytics, depending on the component selected. See Using the Analysis Menu for more information. You can also access the Analysis menu for a component by hovering your mouse over the component's **Device IP** menu until the blue **Quick Info** icon appears, and then clicking the icon.
Locations	Filters components by their geographic locations. See Setting Device Location Information for more information.
Customize Columns	The **Custom Columns** control lets you change the columns that are displayed in the dashboard. See Adding Custom Columns to Dashboards for more information.

ESX Hosts View

When you select an individual ESX Host in the Virtual Infrastructure hierarchy, the ESX Health tab will be selected and you will see a widget dashboard with reports for **ESX Statistics**, **Active Incidents**, **Performance Metrics**, **Memory Utilization**, and **Disk Rate**. Additional tabs are **VM Summary** and **Top VMs**.

Tab Name	Description
ESX Health	A widget dashboard with reports for **ESX Statistics**, **Active Incidents**, **Performance Metrics**, **Memory Utilization**, and **Disk Rate**
VM Summary	A summary dashboard for VMs on the ESX host.
Top VMs	A widget dashboard with reports for **Top VMs by CPU Utilization**, **Top VMs by Memory Utilization**, **Top VMs by Disk Write Request Rates**, **Top VMs by CPU Ready Percentage**, and **Top VMs by Disk Read Request Rate**, all updated hourly

ESX and VM View

When you select an ESX or VM in the Virtual Infrastructure hierarchy, you will see a widget dashboard that contains reports for **VM Statistics**, **Active Incidents**, and **Performance Metrics**.

Widget Dashboard User Interface Overview

Widget dashboards are best for viewing aggregated metrics based on historical search, which are generally presented in the form of a graph or chart. From the widget view of information, you can drill down to view and modify the underlying historical search. Examples of widget dashboards include **Availability/Performance > Avail/Perf Widgets**, the **Security Dashboard**, **BizService Dashboard > Avail/Perf Widgets** and **Security Widgets,** and all the dashboards listed under **Dashboards by Function**.

This screenshot shows an edited view of the **Availability/Performance >Avail/Perf** widgets dashboard. It contains all the standard user interface controls found in widget dashboards.

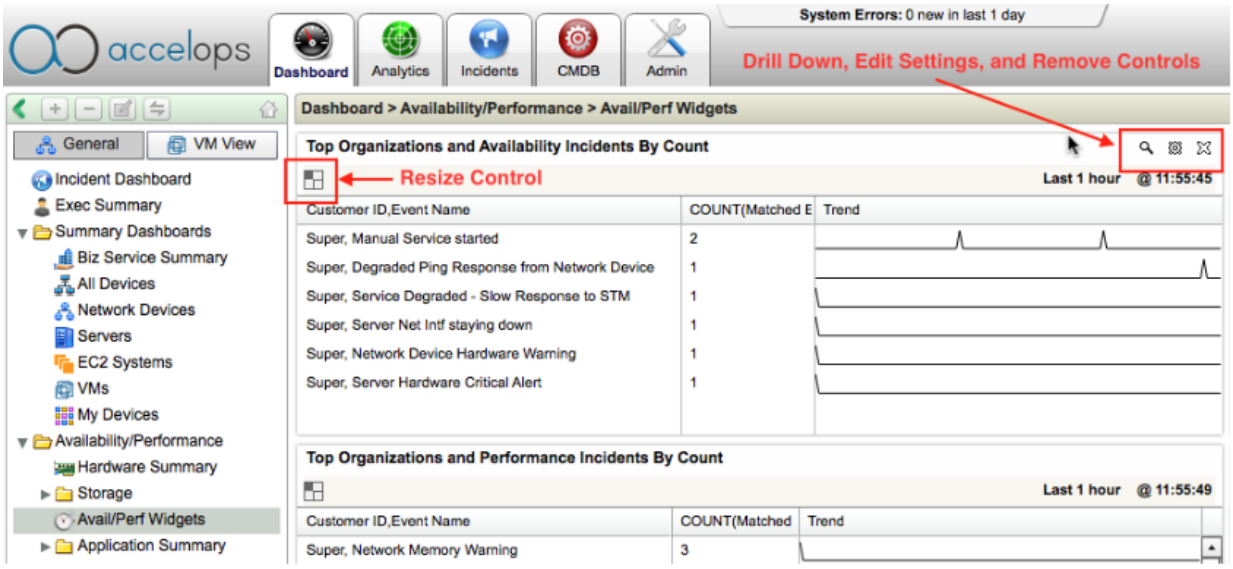

This screenshot shows the **Event Info** menu that you open by hovering your mouse cursor over an event within a widget until the menu icon appears.

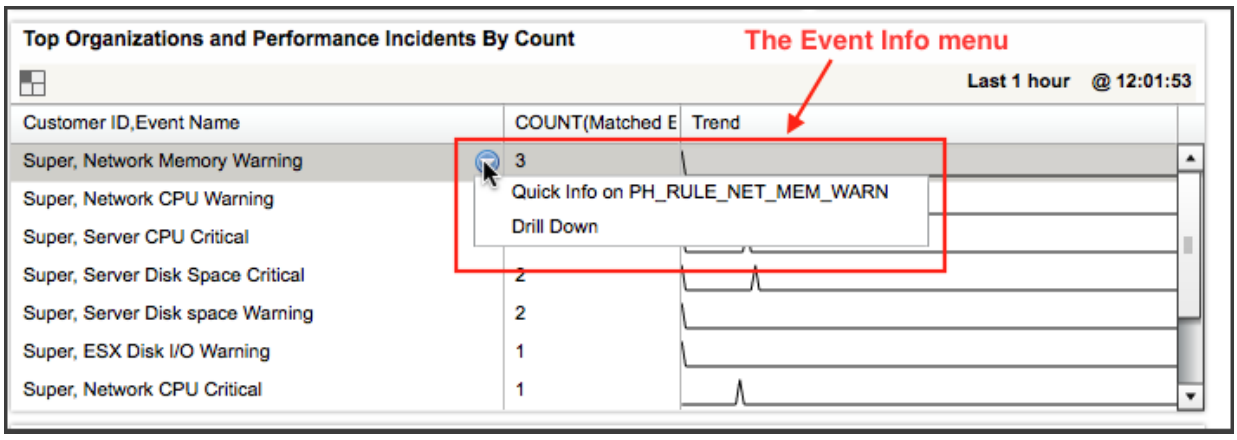

Widget Dashboard UI Controls

UI Control	Description
Resize	You can resize the widget by clicking on this control, and then indicating how many tile spaces you want that widget to use in the dashboard
Drill Down	Hover your mouse cursor over the right upper corner of the widget to access this control. Select a line displayed in the widget to drill down to the historical search associated with that metric. You can then run or modify the search. See Refining the Results from Historical Searchfor more information. This is also the same functionality as the **Drill Down** option in the **Event Info** menu.
Edit Settings	Hover your mouse cursor over the right upper corner of the widget to access this control. Edit the settings associated with the widget. These include: • **Title** - the title of the report • **Description** - a summary description of the report • **Condition** - filters within the report. Look up the report in **CMDB > CMDB Reports** to view the filter conditions it uses. • **Display** - select the type of chart you would like the widget to display • **Time** - the time interval to use in gathering data • **Refresh Interval** - how often the data should be refreshed • **Result Limit** - how many results should be included in the report • **Run report for -** for multi-tenant deployments, select the organization that the widget should report on
Remove	Hover your mouse cursor over the right upper corner of the widget to access this control. Click this control to remove the widget from the dashboard
Event Info	Hover your mouse cursor over a line in a report to view the **Quick Info** for the associated Event Type, or select **Drill Down** to view, edit, and run the associated historical search. See Refining the Results from Historical Search for more information.
Add Report	At the bottom of each widget dashboard is a button to add more widgets to the dashboard.

Related Links

- Refining the Results from Historical Search

Network Topology View of Devices

FortiSIEM provides two ways to view the topology of your IT infrastructure, one at the CMDB level that shows all devices, and another at the level of device groups and individual devices.

- How is Network Topology Discovered and Visualized?
- CMDB All Devices View
- Device Group and Device View
- Viewing Device Information in the Topological Map

How is Network Topology Discovered and Visualized?

FortiSIEM discovers network topology at two levels, layer 3 and layer 2. Layer 3 connectivity involves IP addresses, while Layer 2 connectivity

The layer 3 topology is discovered by obtaining network interface IP address and masks for all devices via SNMP (RFC 1213). The local networks e.g. loopback (127.0.0.0/8), link local addresses (169.254.0.0/16) are filtered out and the distinct networks segments are identified.

A layer 3 topology is visualized on the FortiSIEM Topology map by drawing:

- Network segment and devices as node and
- Srawing line segments from the network segment nodes to every device node that have an interface with IP address in that network segment.

The devices are represented by vendor specific icons and the network nodes are represented by a line and labeled as “Net-<net>/<maskbits>”. For visual clarity:

- Only the network devices are drawn by default. A network device is one that belongs to row Network Device tab in the CMDB.
- Only those networks are drawn that have devices discovered by FortiSIEM (and are in CMDB). There is a “ ” button next to those networks. Clicking on the “” button displays those hosts in the topology graph. Clicking on the “-“ button hides those hosts.

When an enterprise network has Layer 2 switches and hubs, a layer 3 topology misses the connectivity between servers to layer 2 switches and the trunk port connectivity between layer 2/3 switches. Layer 2 discovery is difficult and, more importantly, vendor dependent as vendors have different implementations of the Spanning Tree Protocol (STP).

For Cisco switches, the layer 2 topology is obtained via SNMP (IEEE spanning tree MIB as found in RFC1493 and CISCO-VTP-MIB) as follows:

For every switch,

1. Identify all active VLANs on that switch
2. For every active VLAN:
   a) Get MAC forwarding table
   b) Get STP table to identify trunk ports and directly connected trunk port on adjacent switches

The MAC forwarding table obtained in Step 2a provides the server to switch port connectivity (after eliminating the trunk port entries obtained in step 2b). The trunk port connectivity between switch ports is directly obtained from Step 2b.

The Layer 2 topology is visualized on the FortiSIEM topology diagram by choosing the layer 2 mode. Then by clicking the "+" next to a device, the VLANs on that switch are displayed. Also, the trunk port connectivity is shown in an orange color and a tool tip provides the VLANs over this trunk link.

Then by clicking on the "+" of a VLAN, the hosts belonging to that VLAN and also the switch ports they connect to are displayed.

The host to switch port connectivity can also be seen in a tabular form by first clicking the switch and then clicking the "Port Mapping Table".

CMDB All Devices View

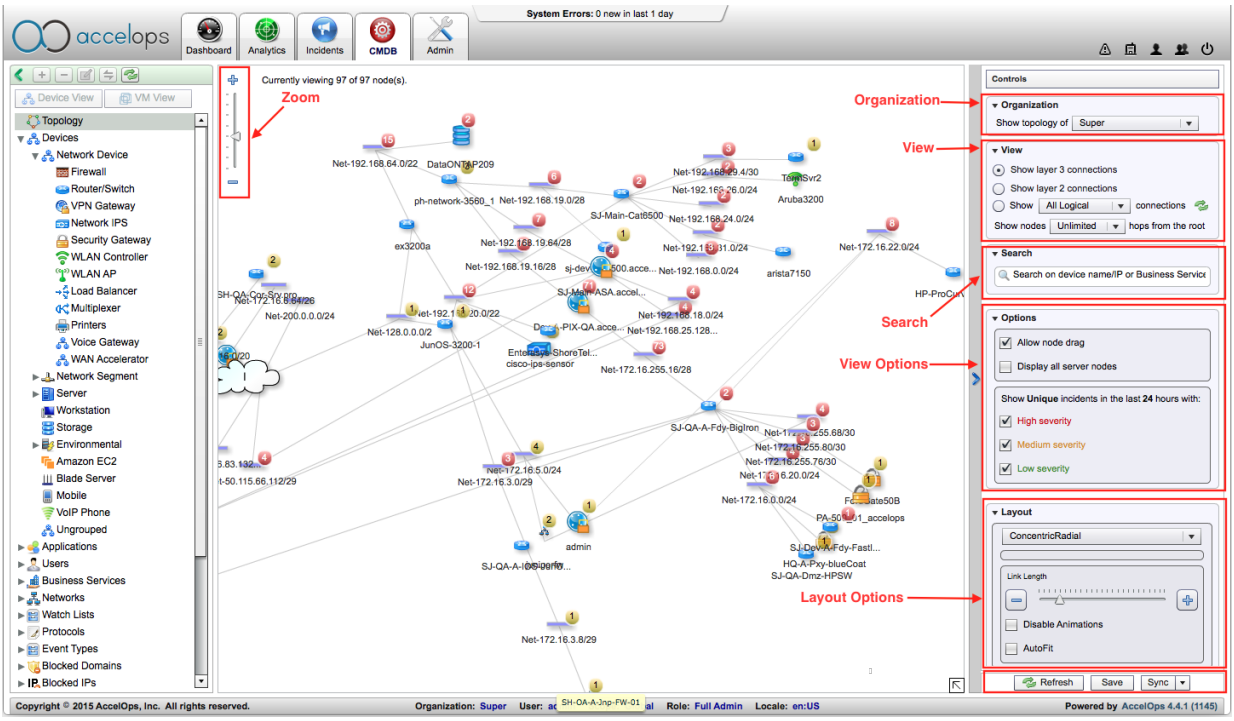

This screenshot shows the **CMDB** tab selected, and in the **Device View**, **Topology** is selected. This topology map shows all the devices for the selected organization, and provides controls for editing the topology views that will be available to users from that organization.

CMDB All Devices User Interface Controls

UI Control	Description
Zoom	Use the slider to increase or decrease the zoom level of the map
Organizations Filter	For multi-tenant deployments, filter devices based on the organization they belong to
View	Select the layers, connection types, and number of hops from the host to display in the map

UI Control	Description
Search	Search for specific devices based on name, IP, or Business Service
View Options	Set the display options, including severity levels, for the map
Layout Options	Set the type of topological map to display, as well as the length of links between devices
Save and Update	• **Refresh** When you make a change to the map settings, click Refresh to see them reflected in the map • **Save** Save your **Layout** and **View Options** to use them in other topographical maps associated with this organization • **Sync** If you make changes to your infrastructure or add devices to the CMDB, click Sync to see them reflected in the map

Device Group and Device View

You can access the device group view of the topological map by selecting a group of devices in the Device View, and then clicking the **Topo** button in the Summary pane. Select an individual device, and then click the Topo button in the Details pane to view that device within the topological map.

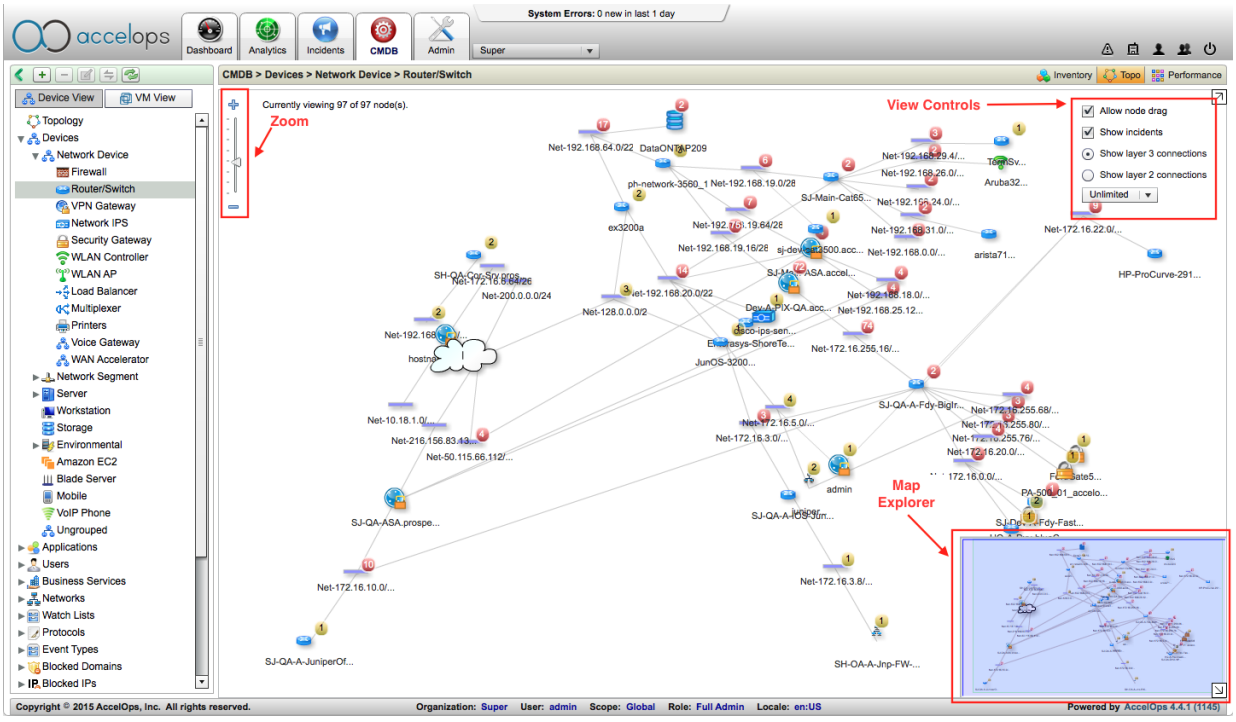

Device Group and Device View User Interface Controls

UI Control	Description
Zoom	Use the slider to increase or decrease the zoom level of the map
View Controls	Click on the arrow icon in the upper-right corner of the map to open these controls. Options to enable/disable node dragging, incident display, connection layer display, and the number of hops from the host to display.
Map Explorer	Click o the arrow icon in the lower-right corner of the map to open the Map Explorer. As you zoom into the map, the map explorer will show you the area that you are currently viewing. You can move to another area by clicking and dragging the highlighted section of the map explorer to that area.

Viewing Device Information in the Topological Map

Devices within the topological map have additional icons to represent information about the device.

Icon	Name	Description
	**Show Connected Hosts**	If a device has a green **+** icon in the topographic map, you can click on that icon to see hosts that are connected to that device
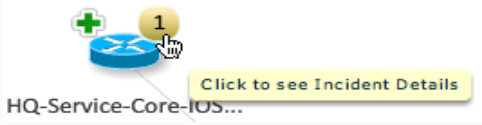	**Show Incident Details**	Incidents for a device are displayed as a number in a circle to the right of the device icon, with the color of the circle (red, yellow, green) indicating the severity of the incidents. Click the number to view the **Incident Summary** for the device, and then click on individual incident to view the **Incident Details** in the Overview of the Incident Dashboard. In the Incident Summary you can also view and apply a subset of options from the Analysis Menu by having your mouse cursor over the **Incident Source** or **Incident Target** entries for the incident.
	**Show Device Details**	Click on the name of the device to view details about it. The kind of information displayed will depend the type of device you select.

How Values in Dashboard Columns are Derived

The values in Summary dashboard columns are either derived from system information (for example, the IP address for a device), or are metrics associated with events and their attributes. This topic uses the example of the **CPU Util** column in many summary dashboards to explain the relationship between event attributes and display columns, and how values in those columns are calculated.

1. Log into you your Supervisor node.
2. Go to **Dashboard > Device View > All Devices**.
3. Click **Select Columns**.
   You will see a list of all the columns used in this dashboard under **Selected Columns**. Under Selected Columns you'll see **CPU Util**, and next to it, in parentheses, you will see three event types listed, whose attributes are used to create this calculation: `PH_DEV_MON_SYS_CPU_UTIL`, `PH_DEV_MON_EC2_METRIC`, and `PH_DEV_MON_CLARION_SP_UTIL`. The metrics associated with these attributes are displayed in the CPU Util column, but how are metrics collected over time represented as a single value? To answer this question, you need to examine the column settings and Aggregation Method in the **Device Support > Dashboard Columns** page.
4. Go to **Admin > Device Support > Dashboard Columns**.
5. Find **System CPU Utilization** in the list of dashboard columns.
   CPU Util is part of the System CPU Utilization set of metric.
6. Each dashboard column has the same set of attributes:

Column Attribute	Description	Value for System CPU Utilization
Name	The metric collected	System CPU Utilization

Event Type	The type of event that provides the attributes for the metric	`PH_DEV_MON_SYS_CPU_UTIL`
		`PH_DEV_MON_EC2_METRIC`
		`PH_DEV_MON_CLARION_SP_UTIL`

Column Name	The display name in the Summary dashboard for the metric	**CPU Name** **Storage Processor** **CPU Utilization** **Host IP Address** Most events include a **Host** IP address, however there is no Column Name for this metric as FortiSIEM generates the column name **Device IP** in relation to the metric.
Column Attribute	The specific attribute used for each Column Name	**Device IP (system generated name)** - `hostIpAddr` **CPU Name** - `cpuName` **Storage Processor** - `spName` **CPU Util** - `cpuUtil`
Column Type	The type of information that will be displayed in the column for each attribute	**Device IP** (system generated name) - `hostIpAddr` - Host **CPU Name** - `cpuName` - Object **Storage Processor** - `spName` -Object **CPU Util** - `cpuUtil` - Reading
Aggregator	For readings, the mathematical aggregator that will be used to calculate the metric. Options are: AVG, SUM, MAX, MIN, LAST. Using a pipe **|** between two operators indicates that the first operation should be aggregated over time, and the second over the object.	**CPU Util** - `cpuUtil` - Reading - AVG|AVG

With this information, you can see that **CPU Util** metric is derived from the `cpuUtil` attribute of the `PH_DEV_MON_SYS_CPU_UTIL` event, and that the display column is a reading that uses the calculation Average over time and then Average over the object being reported on. Now apply this to the event reports for a host with two CPUs, and you can see how the calculation works.

```
[PH_DEV_MON_SYS_CPU_UTIL]:[eventSeverity]=PHL_INFO,[fileName]=phPerfJob.cpp,
[lineNumber]=3137,
[cpuName]=CPU x 1,[hostName]=win2k8.FortiSIEM.net,[hostIpAddr]=192.168.0.40,
[cpuUtil]=2.000000,[pollIntv]=176,[phLogDetail]=
```

```
[PH_DEV_MON_SYS_CPU_UTIL]:[eventSeverity]=PHL_INFO,[fileName]=phPerfJob.cpp,
[lineNumber]=3137,
[cpuName]=CPU x 1,[hostName]=win2k8.FortiSIEM.net,[hostIpAddr]=192.168.0.40,
[cpuUtil]=4.000000,[pollIntv]=176,[phLogDetail]=

PH_DEV_MON_SYS_CPU_UTIL]:[eventSeverity]=PHL_INFO,[fileName]=phPerfJob.cpp,
[lineNumber]=3137,
[cpuName]=CPU x 2,[hostName]=win2k8.FortiSIEM.net,[hostIpAddr]=192.168.0.40,
[cpuUtil]=20.000000,[pollIntv]=176,[phLogDetail]=

[PH_DEV_MON_SYS_CPU_UTIL]:[eventSeverity]=PHL_INFO,[fileName]=phPerfJob.cpp,
[lineNumber]=3137,
[cpuName]=CPU x 2,[hostName]=win2k8.FortiSIEM.net,[hostIpAddr]=192.168.0.40,
[cpuUtil]=40.000000,[pollIntv]=176,[phLogDetail]=
```

This output shows two samples of `cpuUtil` taken over three minutes for each CPU running on the host `192.168.0.40`. According to the **Aggregator** for this column, FortiSIEM should first average the samples over time for each CPU, and then average those together to derive the metric for the host. The average for the CPU 1 is `3.000000`, and the average for CPU 2 is `30.000000`. These values are combined and averaged again to get the overall metric for the host, which is `16.500000`.

Using the Analysis Menu

The **Analysis** menu located in the Summary dashboards presents a number of options for gathering more information about items selected in the dashboard. You can also access the Analysis menu items by selecting a line in a summary dashboard, and hovering your mouse over the IP address of the device until the blue **Analysis** menu option appears.

Analysis Menu Options

Editing Settings for Displayed Reports

You can edit the settings for any of the reports displayed from Analysis menu options. See **Widget Dashboard UI Controls > Edit Settings** in the topic Widget Dashboard User Interface Overview for more information about widget settings options.

Menu Option	Description
Quick Info	The **Quick Info** view of a device, which you can also access through the **Analysis** menu or hovering your mouse cursor over the **Device IP** column, displays **General and Health information** for the device, and when appropriate, **Identity and Location** information. It also contains links to additional information about the device: • **Incidents** An exportable summary of incidents associated with the device • **Health** **Availability**, **Performance**, and **Security** health information for the device**.** You can also access this information by clicking the **Device Health** user interface control, or by selecting **Device Health** in the **Analysis** menu. • **BizService** Any business services impacted by the device. You can also access this information by selecting **Impacted Business Services** in the **Analysis** menu. • **Applications** Displays a report on the top 10 applications associated with the device by **Average CPU Utilization** over the past hour • **Vulnerability and IP Status** (Not used in the **Dashboard** view) Displays the vulnerability status reports that are also available by selecting **Vulnerability and IPS Status** in the **Analysis** menu • **Hardware Health** (Used only for the **CMDB/Storage** view) Displays health information for the hardware being used for storage • **Interfaces** Displays a report on the top 10 interfaces associated with the device by average throughput • **Topology** Shows the device's location in the network topology**.** You can also access this information by selecting **Topology** in the **Analysis** menu. The Quick Info view also contains two links, **Goto Config Item**, which links to the device entry in the CMDB, and **Goto Identity**, which links to **Analytics > Identity and Location Report**, where you can edit this information for the device.
Topology	Shows the device location within the network topology

Menu Option	Description
Device Health	**Availability** , **Performance** , and **Security** health reports for the device. You can also access this information by selecting a device in the Summary dashboard, and then click **Health,** or by going to **Quick Info > Health** after selecting the device . If any **Incidents** are displayed, click the number to view the **Incident Summary** . Depending on the reported metric, you can zoom in for a closer look at graphs and reports by clicking the **Magnifying Glass** icon that appears when you hover your mouse cursor over them.
Incidents Summary	A summary of incidents associated with the device. Select an incident and then hover your mouse cursor over the **Incident Name** to open the **View Incident Details** option, which will load the selected incident into the **Incident Dashboard**. See the topics under Incidents - Flash version for more information about working with the Incident Dashboard. If you hover your mouse cursor over the **Incident Target** for an incident in the Incident Summary screen, you will see some additional options, including: • **Add to Watch List** - add the incident target to a watch list. See Watch Lists for more information. • **Show Related Real Time Search** - opens a real time search using the Host IP and Name for the incident target • **Show Related Historical Search** - opens an historical search using the Host IP and Name for the incident target
Device Availability	Displays reports for **Availability Trend Status**, **Ping Response Time**, and **Ping Packet Loss** for the device over the past hour, and **Device Uptime** for the device over the past thirty minutes
Device Performance	Displays reports for **Performance Health Trend**, **Avg Memory Utilization**, **Avg CPU Utilization**, and **Avg Disk Utilization** over the past hour for the device
Interface Status	Displays reports for **Interface Utilization Percentage**, **Interface Error Percentage**, **Interface Traffic**, and **Interface Error Count** over the past hour for the device
Application Performance	Displays reports for **Average Application CPU Utilization**, **Application CPU Utilization**, **Average Application Memory Utilization**, and **Application Memory Utilization** over the past hour for the device
Event Status	Displays reports for **Events per Second**, **Top Network Connections**, **Top Events by Severity**, and **Top TCP/UDP Ports** over the past hour for the device
All Events by Group for the Last 10 Minutes	Opens an Historial Search for the selected device using these criteria

Menu Option	Description
Traffic Status	Displays reports for **All Permitted Traffic Sourced From or Destined to** the selected device, and **All Denied Traffic Sourced from or Destined to** the selected device over the previous hour
Vulnerability and IPS Status	Displays reports for **All Vulnerabilities for Last 1 Day** and **All Warning + Critical IPS Events** for the device over the past 24 hours
Impacted Biz Services	Business services that contain the selected device
Real-time Events	Opens a Real-Time Search for the selected device
Historical Events for Last 5 Mins	Opens an historical search for all events associated with the device over the past five minutes

Customizing Dashboards

FortiSIEM includes several dashboards for device types and IT functional areas, but you can also customize and create new dashboards and widgets.

- Adding Custom Columns to Dashboards
- Adding Widgets to Dashboards
- Creating a Customized Dashboard
- Setting a Dashboard to Home

Adding Custom Columns to Dashboards

You may want to add custom columns based on event attributes to a Summary dashboard. This topic explains how to create a custom set of columns using the example of a hardware temperature readout, and then add them to a dashboard.

Prerequisites

- Read the topic How Values in Dashboard Columns are Derived

Procedure

1. Find the event that contains the attribute you want to use.
   In this case, you want to create a hardware temperature reading. The event `PH_DEV_MON_HW_TEMP` contains the attribute `envTempDegC`.

   ```
   [[PH_DEV_MON_HW_TEMP]:[eventSeverity]=PHL_INFO,[fileName]=deviceJunOS.cpp,
   [lineNumber]=619,[hostName]=JunOS-3200-1,
   [hostIpAddr]=172.16.5.64,[hwComponentName]=FPC- EX3200-24T- 8 POE @ 0/*/*,
   [envTempDegC]=33,[phLogDetail]=
   ```

2. Go to **Admin > Device Support > Dashboard Columns**.
3. Click **New**.
4. For **Name**, enter the display name for the new metric you want to collect.
   For this example, enter the name **Temperature Reading**.
5. For **Event Type**, click the **Edit** icon and select the event you want to use.
   For this example, select `PH_DEV_MON_HW_TEMP`.
6. Click the **+** icon to add a column. As you complete each column, click **OK**, then click + to add more columns.
   For each event type, you will typically create three columns: a **Host** column that contains IP information for associated hosts, an **Object** column that includes information about the object being reported on, and a **Reading** column that contains the metric you want to report on. Note that you could create additional Reading columns for other attributes contained in your event.

Column Type	Example Settings
Host	**Attributes :** `hostIpAddr` **Aggregator : N/A** **Display Name : N/A** **Format : N/A** **Trend Chart : N/A** **Type : Host**
Object	**Attributes :** `hwComponentName` **Aggregator : N/A** **Display Name : N/A** **Format : N/A** **Trend Chart : N/A** **Type : Object**
Reading	**Attributes :** `envTempDegC` **Aggregator : AVG|MAX** **Display Name : Temp** **Format : DegreeC** **Trend Chart : Health** **Type : Reading**

7. When you're finished adding columns, click **OK**.
   The new column you created will appear in the **Admin > Device Support > Dashboard Columns.**
8. Select your new column in the list, and then click **Apply**.
9. To add your column to a dashboard, navigate to the dashboard**.**
10. In the dashboard, click **Select Columns**.
11. Under **Event Types**, select the event type you used to create the new column.
    The columns associated with that event type will be listed under **Columns**, and the **Attribute Name** will list the attribute you used to create the column.
12. Under **Columns**, select your column and use the **>>** button to move it into the **Selected Columns**.
13. Use the up and down position buttons to place the column in the order where you want it to appear in the dashboard.
14. Click **OK**.
    Your new column will appear in the dashboard.

Adding Widgets to Dashboards

1. Navigate to the widget dashboard where you want to add the widget.
2. At the bottom of the dashboard click **Add Reports to Dashboard**.
3. For Service Provider deployments, select the **Organization** that you want to have access to the report.
4. Select a **Category** for the type of report you want to add.
5. Under **Available Reports**, select the report you want to add, and then click the **>>** button to add it to the **Selected Reports**.
6. Click **OK**.

To add CMDB Reports, select from the CMDB Reports folder in Step 5.

Creating a Customized Dashboard

You can create both Summary and Widget custom dashboards.

1. In the **Dashboard** tab, select **My Dashboard** in the **General** view.
2. At the top of the **General** view, click the **+** icon.
3. Enter a **Group** to categorize the dashboard, and a **Description**.
4. Select a **Dashboard Type**.
5. Click **OK**.
   The dashboard will be added under **My Dashboard**.
6. Select the dashboard.
7. For a **Device Summary Dashboard**, click **Devices** at the top of the dashboard and select the devices you want to add to the dashboard.
8. For a **Widget Dashboard**, click **Add Reports to Dashboard**, and then select the reports you want to add.

Setting a Dashboard to Home

You can set any system or user-defined dashboard to be your home page when you log into FortiSIEM.

1. In the **Dashboard** view, select the dashboard you want to set for your home page.
2. At the top of the **General** view of the dashboards, click the **Home** icon.
   The Home icon will be filled in rather than greyed out, and the next time you log into FortiSIEM, the page you selected will be your home page.

Creating Dashboard Slideshow

- Go to **Dashboard**
- Click on Slideshow icon on top left. 
    - A check box appears next to each dashboard in the dashboard tree under the **General** tab
    - For each folder, expand the folder to see the check box for each dashboard
    - Select the dashboards for slideshow
- Select the **Interval** for switching between dashboards
- Click **Start** to enter Slideshow mode. Click **Cancel** to not save the current slideshow configuration
- Once in Slideshow mode, click **Escape** button to stop the slideshow

Note: Slideshow configuration is saved on a per user basis. When the user logs back on, same slideshow can be used.

Exporting and Importing Dashboards

It is possible to export and then import the following types of widget dashboards

- My Dashboard
- Availability/Performance > Avail/Perf Widgets
- Biz Svc Dashboard
- Dashboards By Function

To export a dashboard

- Go to a specific dashboard folder
- Click **Export** on top right portion
- An XML file will be created and saved.

To import a dashboard, first have the XML file ready

- Go to a specific dashboard folder
- Click **Import** on top right portion
- Provide the dashboard file in XML format

Link Usage Dashboard

For perimeter network devices such as firewalls and routers, it is important to know which interfaces are busy and which traffic is consuming the most resources. This special dashboard provides this view and enables users to determine which router interfaces are overly utilized, which applications are using them and what is the QoS statistics.

- Go to **CMDB > Devices > Firewall** or **CMDB > Devices > Router/Switch**
- The default is **Inventory View**. Click **Link Usage** to change to this special view
- A three level panel appears on right
  - **Top pane**: Device level view: System level metrics such as CPU, Memory, Connections, Sent/Receive Traffic, Received EPS. Source is SNMP.
  - **Middle pane**: For the selected device in Top pane, it shows metrics for each interface. Source is SNMP.
  - **Bottom pane**: For the selected device in Top pane and interface in middle pane, it shows
    - Application Usage: Top Applications, Top Sources, Top Connections. Source is Netflow.
    - QoS Statistics: QoS statistics. Source is SNMP - Fortinet only

Note: Slideshow configuration is saved on a per user basis. When the user logs back on, same slideshow can be used.

Dashboards - HTML5 version

FortiSIEM includes two types of dashboards:

- **Summary dashboards** that shows multiple metrics for the device in a single line. This enables users to see multiple metrics of the same device in one view.
- **Widget dashboards** that provide separate views of each metric. This enables to see critical devices for a metric at a time.

Multiple dashboards can be grouped into a folder. User first needs to choose the dashboard folder and then select the dashboard within that folder.

Viewing System Dashboards

FortiSIEM provides several built-in dashboard folders covering many functional areas:

- Infrastructure level
  - Network Dashboard
  - Server Dashboard
  - VMWare Dashboard
  - Web Server Dashboard
  - Application Server dashboard
- Cloud Infrastructure level
  - Amazon Web Services Dashboard
- Security Dashboard
- Storage level
  - NetApp Dashboard
  - VNX Dashboard
- Application level
  - Salesforce Dashboard
  - Office 365 Dashboard
  - Google Apps Dashboard
- FortiSIEM Dashboard

To view these dashboards

1. Logon to FortiSIEM
2. Switch to the right organization (for Service Provider version)
3. Click **Dashboard** tab on the main user interface
4. Select the appropriate dashboard folder from the drop down. The dashboards belonging to the selected folder will show and the contents of the first dashboard will display automatically.
5. Select the appropriate dashboard to see its contents.

Creating New Dashboards

Make sure that you are logged on to the right organization (for Service Provider version).

Creating a new dashboard folder

1. Click on the dashboard folder menu and Select **New.**
2. Enter the name of the new dashboard folder
3. The new dashboard will show

Creating a new dashboard within a folder

1. Click on the  icon on the top bar
2. Enter the following information
   - **Name** - the name of the dashboard
   - **Type** - Widget or Summary dashboard
   - **Description**
3. Click Save

Adding reports to a widget dashboard

1. Click on the  icon on the left under the dashboard name
2. Select the report and it will highlight
3. Drag the report to the dashboard and the results will show
4. To customize the chart settings, see here.

To add a CMDB Report, simply add from the CMDB Report folder in Step 2.

Adding devices to a summary dashboard

1. Click on the  icon on the top menu bar
2. Select the device(s) and move them to the right pane by clicking the  button
3. Click OK
4. To customize the columns, see here

Deleting Dashboards

Note that built-in dashboard folders and dashboards can not be deleted.

Deleting user defined dashboards

1. Click on the button next to the dashboard.
2. Click OK.

Deleting user defined dashboard folders

1. Click on the button next to the dashboard folder.
2. Click OK.

Modifying Dashboards

Saving changes

User settings changes are saved - both for built-in and user created dashboards. If the user logs back in, then the changes will be seen. System upgrades will also preserve these customizations.

Modifying widget display

1. Select a widget and click on the  Settings button
2. Customize the fields as appropriate
   - **Title** - the chart name that displays at the top
   - **Display** - select chart type from the possible options
   - **Width** - the size of the chart in horizontal dimension - note that this is relative
   - **Height** - the size of the chart in vertical dimension - note that this is relative
   - **Refresh interval** - how often the chart's content will refresh
   - **Result Limit** - number of rows in the result
3. Click **OK**.

Adding reports to a widget dashboard

1. Click on the  icon on the left under the dashboard name
2. Select the report and it will highlight.
3. Drag the report to the dashboard and the results will show
4. To customize the chart settings, see here.

If you want to add a new report or modify a system report, then follow these steps

1. Create the report in Analytics
2. Then report will show up in the list of reports in Step 2 above.

Modifying widget dashboard layout

There are two possibilities - Tile layout (default) or column layout.

1. To select Tile layout, select **Tile** option from the menu next to 

on top. Tile layout allows you to place widgets of several sizes on the dashboard.
2. To select a column layout, choose the number of columns from the menu next to 

.

Adding, removing and re-ordering columns on a summary dashboard

1. Select the  button the top.
2. To remove one or more columns from display, select them in the **Selected Columns** and then move them to the left by clicking the  button.

3. To add one or more columns to the display:
    a. Select an **Event Type** in the left most column. The corresponding metrics from that event type will show.
    b. Select one or more columns in the middle column.
    c. Move them to the right by clicking the  button.
4. To change the position of the columns
5. Click OK to save the changes.

Sharing Dashboards

The following sharing rules are enforced

- User created dashboard folders and its contents are only visible to the user who created it. If this folder need to be visible to other users, then we recommend:
  - using a shared account or
  - using export/import mechanism to create the folder for that user.
- System dashboard folders are owned by FortiSIEM. Any changes to those dashboards may be lost during upgrade, if FortiSIEM also decides to change those dashboards.

Importing and Export Widget Dashboards

Importing widget dashboards

Widget Dashboards can be imported from another FortiSIEM installation or from another dashboard folder of the same installation. If the two FortiSIEM versions do not have the same version, then the charts may look different because the data definition may be different.

1. Make sure you are viewing the dashboard
2. Click **Import** 

.
3. Select the file from local desktop. It must an XML file suitable for import. Typically this is exported from another FortiSIEM system.
4. Click **Import**.
5. The dashboard will display

Exporting widget dashboards

1. Make sure you are viewing the dashboard.
2. Click **Export** 

Analytics

FortiSIEM Analytics has three components:

Search

FortiSIEM search functionality includes real time and and historical search of information that has been collected from your IT infrastructure. With real time search, you can see events as they happen, while historical search is based on information stored in the event database. Both types of search include simple keyword searching, and structured searches that let you search based on specific event attributes and values, and then group the results by attributes.

Rules

Because FortiSIEM is continuously monitoring your IT infrastructure, you can also set rules so that when specific conditions are met, it triggers an incident, and, in some cases, sends a notification.

Reports

Reports are pre-defined search queries. FortiSIEM includes a large catalog of reports for common devices and IT analysis tasks that you can use and customize, and you can also save searches that you've run as reports to use again later.

- Search
- Rules
- Reports
- Audit
- Visual Analytics
- Real Time Performance Probe

Search

Historical and Real Time search is the core functionality of FortiSIEM analytics, enabling you to analyze, report on, and further improve your IT infrastructure.

- Historical Search
- Real Time Search
- Structured Search Operators
- Selecting Attributes for Structured Searches, Display Fields, and Rules
- Using Expressions in Structured Searches and Rules
- Keywords and Operators for Simple Searches
- Using Geolocation Attributes in Searches and Search Results
- Creating Filter Criteria and Display Column Sets

Historical Search

With the Historical Search feature, you can go back in time and retrieve events from the event database. By using either a simple keyword-based search or a more detailed structured search, you can get quick and valuable insights into events that have occurred over any selected time period.

Overview of the Historical Search User Interface

You can run two types of historical searches on FortiSIEM data: simple searches, in which you use a keyword search, and structured searches, in which you can specify search conditions and how the results should be grouped.

- Simple Historical Search
- Structured Historical Search

Simple Historical Search

When you use simple historical search, you enter a keyword to search for in the logs collected by FortiSIEM, specify any filter criteria, and then run the search, which will produce a chart and a list of results matching your search criteria. You can then use additional user interface controls to change the chart display, filter or find more information about events in the result list, and export or share results.

This screenshot shows the results of simple search using the keyword **TCP**.

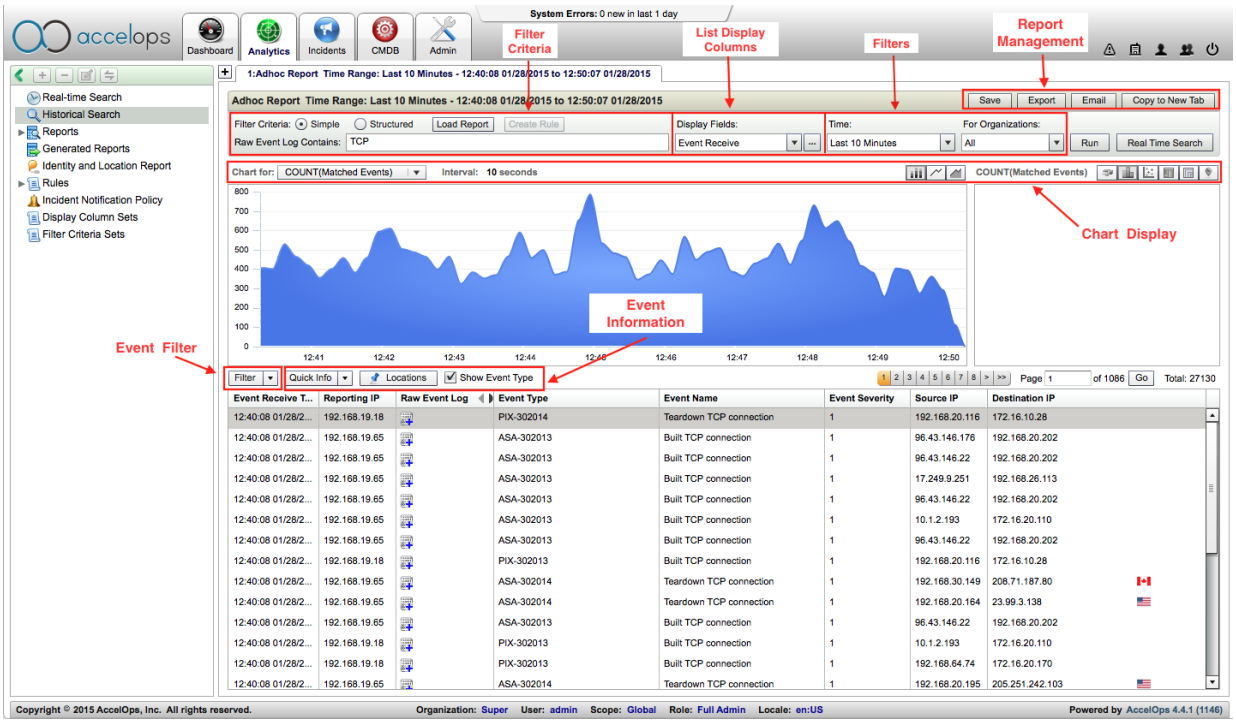

Simple Historical Search User Interface Controls

UI Control	Description
Search Criteria	For simple historical search, use the search box to find keywords in raw event logs. You can also load an existing historical search report to use for your search criteria, or create a rule from your search results.
List Display Columns	Select which columns will be displayed in the search results
Filters	Set the time interval over which you want to search, and, for multi-tenant deployments, which organization's logs you want to search
Report Management	• **Save** Saves the report to Generated Reports where it will be retained for the time period you specify. You can also select whether you want the search criteria to be saved as a report that you can use in the future. • **Export** Export the report, with the option of including the chart, as a PDF or CSV file • **Email** Email the report as a CSV or PDF file, with the option of including the chart • **Copy to a new tab** Load the search into a new tab within FortiSIEM
ChartDisplay	You can set both the data you want to display, and how it should be displayed. See Overview of Historical Search Results and Charts for more about the different chart types.
Event Filter	Select an event from the results, and add its attributes to structured search conditions.
Event Information	Select an event, and view **Quick Info** about it, or view **Location** information about it such as source or destination IPs.

Structured Historical Search

With historical structured search, you can enter conditions for your search based on event attributes, and set which attributes will be used to group the search results in a way that is similar to the use of the of the Group By command in SQL.

This screenshot shows a structured historical search for **All Non-Reporting Modules** selected from the system **Reports > Event Status**. The screenshot below it shows a close-up of the the **Conditions and Group By**

options dialog. See Creating a Structured Historical Search and Structured Search Operators for more information about these options.

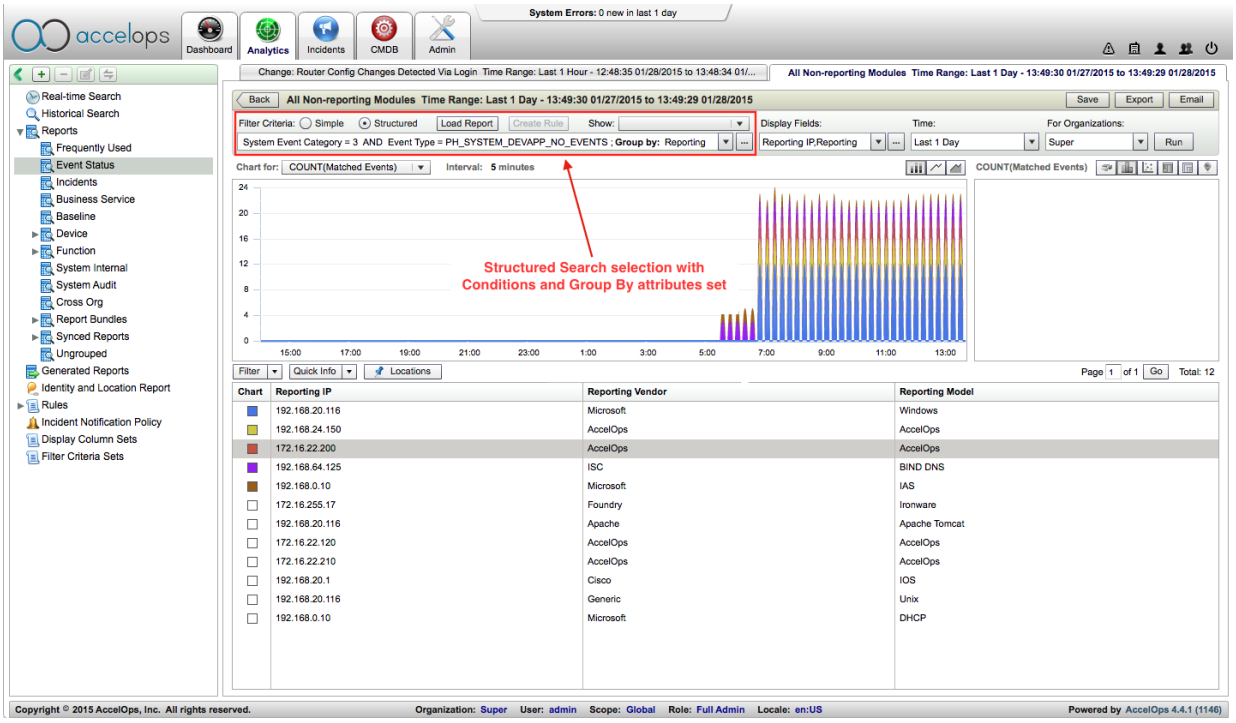

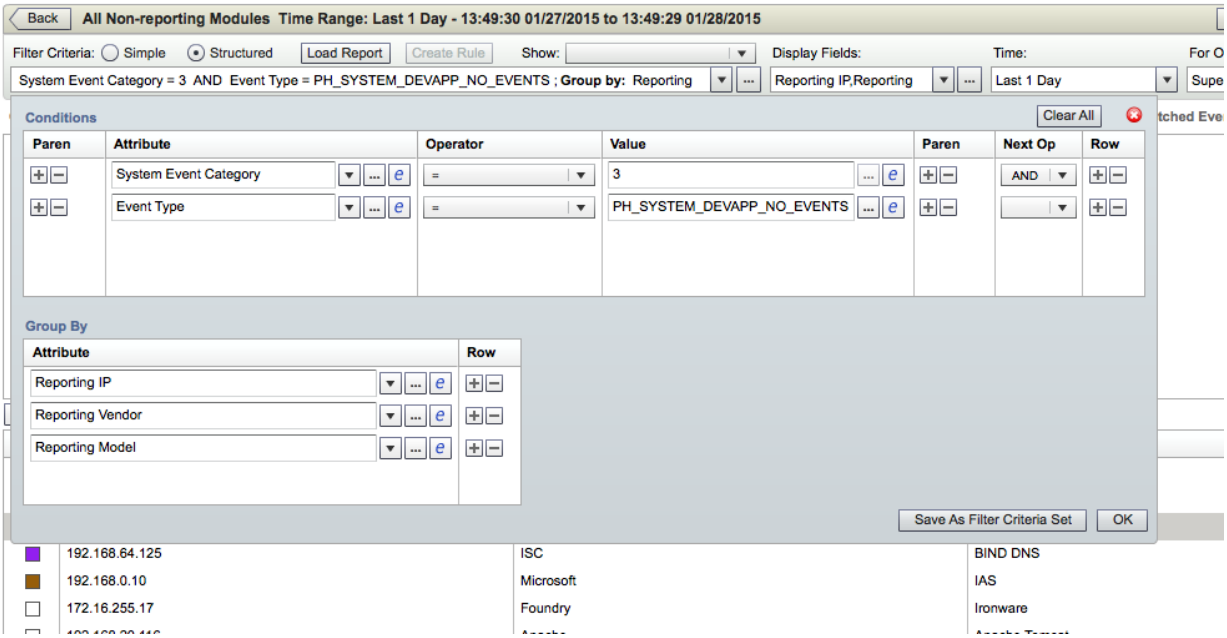

Example of How a Structured Historical Search is Processed

When you run a structured historical search, all events within the specified time window are examined and added to the result set following these steps:

1. The system fetches the next event within the search time window and applies the filtering criteria. If the event does not pass the filtering criteria, the system fetches the next event.
2. If the event passes the filtering criteria, the system then compares the attributes of this event against the other entries in the result set. If the current event contains an attribute that is included in the Group By attribute set, then the results for that attribute are updated. Otherwise, a new entry is created in the result set.
3. After all the events in the search time window are processed, the system sorts the results to produce the final result set.

As an example, consider these events in the event database, and running a search for **Top Firewall Recorded Conversations Ranked By Total Connections (Descending) and Total Bytes (descending)** over them**.**

Event id	Time	Reporting Device	Source IP	Destination IP	Protocol	Source Port	Destination Port	Total Bytes
1	1/1/2010	10.1.1.1	192.168.1.1	192.168.10.4	TCP	2033	80	1024
2	1/2/2010	10.1.1.1	192.168.1.2	192.168.10.4	TCP	3000	443	4096
3	1/3/2010	10.1.1.1	192.168.1.1	192.168.10.4	TCP	2034	80	1024
4	1/4/2010	10.1.1.1	192.168.1.2	192.168.10.5	TCP	3001	443	2048
5	1/4/2010	10.1.1.1	192.168.1.1	192.168.10.4	TCP	2035	80	1024
6	1/5/2010	10.1.1.1	192.168.1.2	192.168.10.6	TCP	3002	443	2048
7	1/5/2010	10.1.1.2	192.168.1.1	192.168.10.4	TCP	9000	80	1024

Search	Search Criteria
Top Firewall Recorded Conversations Ranked By Total Connections (Descending) and Total Bytes (descending)	**Filtering criteria:** Reporting Device IP IN Firewall AND Event Type IN Permit Traffic **Group-By attributes:** Source IP, Destination IP, IP Protocol, Destination Port **Display attributes:** Source IP, Destination IP, IP Protocol, Destination Port, SUM(Matched Events) DESC, SUM(Total Bytes) DESC **Query window:** Between 1/2/10 and 1/5/10

Result

Source IP	Destination IP	Protocol	Destination Port	COUNT (Matched Events)	SUM(Total Bytes)
192.1.1.1	202.1.1.4	TCP	80	3	3072
192.1.1.2	202.1.1.4	TCP	80	1	4096
192.1.1.2	202.1.1.5	TCP	443	1	2048
192.1.1.2	202.1.1.6	TCP	443	1	2048

You could then run another search over these results:

Search	Search Criteria
Top Destination IPs Ranked By Total Connections (Descending) and Total Bytes (descending)	**Filtering criteria:** Reporting Device IP IN Firewall AND Event Type IN Permit Traffic **Group-By attributes:** Destination IP **Display attributes:** Destination IP, SUM(Matched Events) DESC, SUM(Total Bytes) DESC **Query window:** Between 1/2/10 and 1/5/10

Result

Destination IP	COUNT (Matched Events)	SUM(Total Bytes)
202.1.1.4	4	7 KB
202.1.1.5	1	2 KB
202.1.1.6	1	2KB

Sample Historical Searches

- Sample Filter Criteria
- Sample Structured Searches

Sample Filter Criteria

Filter criteria	Type	Meaning
Raw Event Log CONTAINS "login AND failed"	Simple (keyword) search	Only events that contain both the keywords "logon" and "failed" are part of report
Raw Event Log CONTAINS "denied"	Simple (keyword) search	Only events that contain the keyword "denied" are part of report
Reporting Device IP = 10.1.1.1	Structured search	Only events from the device that is reporting with IP address 10.1.1.1 are part of the report
Reporting Device IP IN Firewall	Structured search	Only events from firewall devices in CMDB are part of the report
Reporting Device IP IN Firewall AND Event Type IN Deny Traffic	Structured search	Only firewall deny events from firewall devices in CMDB are part of the report
Reporting Device IP IN Firewall AND Event Type IN Deny Traffic AND (Source IP = 192.1.1.1 OR Dest IP = 192.1.1.1)	Structured search	Denied traffic from 192.1.1.1 or to 192.1.1.1 reported by firewall devices in CMDB are part of the report
Reporting Device IP IN Domain Controller AND Event Type IN User/Group Change AND user NOT IN Domain Admins	Structured search	Domain Controller User/Group Changes not performed by users in the Domain Admin group
Raw Event Log REGEXP "faddr\s+\d+.\d+\d+\d+"	Structured search	Only events that contains strings like "faddr 10.1.1.1", "faddr 192.168.29.1" are included in the report.

Sample Structured Searches

The following examples illustrate how to write a search using the FortiSIEM GUI.

Search	Specification in FortiSIEM GUI
Top Reporting Firewalls ranked by event count in the last hour	**Filter Criteria:** Reporting Device IP IN Firewall **Group By attributes:** Reporting Device IP **Display attributes:** Reporting IP, COUNT(Matched Events) DESC **Query window:** 1 hour
Top Reporting Firewalls and Event Types ranked by event count in the last hour	**Filte Criteria:** Reporting Device IP IN Firewall **Group By attributes:** Reporting Device IP, Event Type **Display attributes:** Reporting IP, Event Type, Severity, COUNT(Matched Events) DESC **Query window:** 1 hour
Top Firewall Denied Source IPs ranked by the total number of attempts in the last hour	**Filter Criteria:** Reporting Device IP IN Firewall AND Event Type IN Deny Traffic **Group By attributes:** Source IP **Display attributes:** Source IP, COUNT(Matched Events) DESC **Query window:** 1 hour
Top Firewall Recorded Conversations Ranked By Sent Bytes (descending), Received Bytes (descending)	**Filter Criteria:** Reporting Device IP IN Firewall AND Event Type IN Permit Traffic **Group By attributes:** Source IP, Destination IP, IP Protocol, Destination Port **Display attributes:** Source IP, Destination IP, IP Protocol, Destination Port, SUM(Sent Bytes) DESC, SUM(Received Bytes) DESC **Query window:** 1 hour
All unauthorized domain user/group changes in the last week	**Filter Criteria:** Reporting Device IP IN Domain Controller AND Event Type IN User/Group Change AND user NOT IN Domain Admins **Group By attributes:** none **Display attributes:** Time, event type, user, computer, domain, target user, target domain **Query window:** 1 week

Creating a Simple Historical Search

Prequisites

If you need to familiarize yourself with how historical search works or the historical search interface, you should read these topics:

- Overview of the Historical Search User Interface
- Example of How a Structured Historical Search is Processed
- Sample Historical Searches
- Structured Search Operators

Procedure

1. Log in to your Supervisor node.
2. Go to **Analytics > Historical Search**.
3. For **Filter Criteria**, select **Simple**.
4. Enter the keywords you want to search for in the raw event logs.
   See Keywords and Operators for Simple Searches for information on keyword searching.
5. Under **Display Fields**, select the attributes you want to use as the columns in your results list.
   See Selecting Attributes for Structured Searches, Display Fields, and Rules and Creating Filter Criteria and Display Column Sets for options for selecting display field attributes and sets.
6. For **Time**, set the interval over which you want the search to run.
7. For Service Provider deployments, select the **Organization** you want to run the search against.
8. Click **Run**.
   The results of your search will be displayed in the chart and search results list.

Creating a Structured Historical Search

Prequisites

If you need to familiarize yourself with how historical search works or the historical search interface, you should read these topics:

- Overview of the Historical Search User Interface
- Example of How a Structured Historical Search is Processed
- Sample Historical Searches
- Structured Search Operators

Procedure

1. Log in to your Supervisor node.
2. Go to **Analytics > Historical Search**.
3. For **Filter Criteria**, select **Structured**.
   The **Conditions and Group By** search window will open.
4. Click the downward arrow in the search window to open the **Conditions and Group By** options.
   Alternatively you can click **...** to use a saved **Filter Criteria Set**.
5. Under **Conditions**, set the **Attribute**, **Operator**, and **Value** for your condition.
   You can also use expressions as search conditions. See Using Expressions in Structured Searches and Rules for more information, and Selecting Attributes for Structured Searches, Display Fields, and Rules for more information about using attributes in conditions.
6. Click **+** under **Row** to add another condition, and set the **Next Operator** to use for that condition.
   You can give precedence to conditions by setting parentheses around them with the **+** button under **Paren**.
7. Under **Group By**, set the event attributes that you want to use to group the results, as described in Example of How a Structured Historical Search is Processed.
8. Click **OK**.
   You can also click **Save as Filter Criteria Set**, and these conditions and group by attributes will be available for future historical searches by clicking **...** next to the search window.
9. Under **Display Fields**, select the attributes you want to use as the columns in your results list.
   See Selecting Attributes for Structured Searches, Display Fields, and Rules for more information about selecting attributes for devices and events to use as display fields.
10. For Service Provider deployments, select the **Organization** you want to run the search against.
11. For **Time**, set the interval over which you want the search to run.
12. Click **Run**.
    The results of your search will appear in the chart and results list.

Using System-Defined Reports for Historical Search

FortiSIEM includes a number of pre-defined reports that you can use as the basis for historical searches.

- Viewing Available Reports
- Using System-Defined Reports in Historical Searches

Viewing Available Reports

1. Log in to your Supervisor node.
2. Go to **Analytics > Reports**.
3. Select a report group in the navigation pane, and then a report.

Each report includes four information tabs:

Tab	Description
Summary	Includes name, description, and all the criteria used in constructing the historical search for the report
Schedule	Any scheduled runs for the report. See Scheduling Reports for more information.
Results	Any saved results from running the report
Defintion	The XML definition of the report

Using System-Defined Reports in Historical Searches

1. Log in to your Supervisor node.
2. Go to **Analytics > Historical Search**.
3. Click **Load Report**.
4. Select the report you want to use, and then click **OK**.
5. Follow the same steps that you would for Creating a Structured Historical Search.

Overview of Historical Search Results and Charts

When your search runs, you will see both a **Results List** in the bottom pane of the screen, and a chart in the middle pane. The types of charts that are displayed depend both on the data being analyzed, and whether or not you have specified any **Group By** conditions in your search. You can also add dimensions to your search results and change the chart display type for further analysis.

- Non-Aggregated Search Results
- Aggregated Search Results

Non-Aggregated Search Results

Non-aggregated searches are searches that don't use any Group By conditions to process the results. These types of searches produce two views of the results:

View	Description	Screen Example
Trend	Shows the trend over time for search results	
Results List	Shows the results of the search based on the Search Display fields you selected	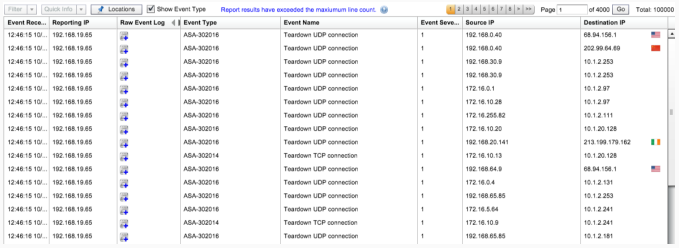

Aggregated Search Results

Aggregated searches are those that use a Group By condition to process the results.

View	Description	Screen Example	Notes
Results List	Shows the results of the search based on the Group By and Display fields you selected	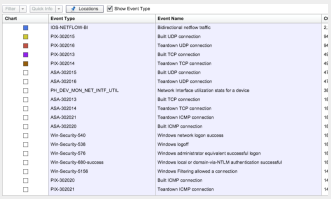	This example shows the search results for **Top Event Types by Count** • **Filter Condition**: Empty • **Group By Condition**: Event Type • **Selected Display Fields**: Event Type and COUNT(Matched Events)
Trend	Shows the time trend of aggregated fields (one at a time)	 	There are two trend views of results for aggregated searches, the line chart, shown here as the first chart, and the stack chart, shown as the second chart. In this example, the line chart illustrates when the events occurred. The stacked display avoids line crossings, but the values have to be read off as the height and not the absolute value. For example, the event count for PIX-302015 at 9:00 hours is 20,000-14000 = 6000.
Pie Chart	Shows the proportion for the COUNT (Matched Events) attribute	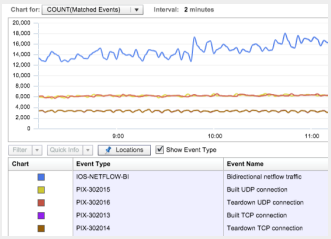	For any set of results where you are charting **Count (Matched Events)**, click the **Pie Chart** icon to view a proportional representation of the results.
Bar Chart	Shows the distribution of aggregated fields	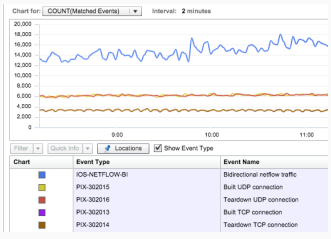	For any set of results where you are charting **Count (Matched Events)** , click the **Bar Chart** icon to view the distribution of events for your results.

View	Description	Screen Example	Notes
Scatter Plot	Shows the correlation between two aggregated fields	 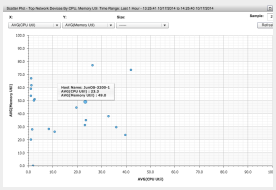	Scatter plots can show the correlation between two aggregated dimensions, effectively converting a one dimensional chart into a two dimensional one. In this case, a report is run with these parameters: • **Filter Condition**: Event Types `PH_DEV_MON_SYS_CPU_UTIL` and `PH_DEV_MON_SYS_MEM_UTIL` • **Group By attribut**: Host Name • **Display Fields**: AVG(CPU Utilization) and AVG(Memory Utilization) The results are first presented as a stacked trend and bar chart. When you click on the **Scatter Plot Chart** icon, you can now see the display fields as two dimensions, which shows that most devices use more memory than CPU. Hovering your mouse cursor over an item in the chart displays the values for the selected host.
Bubble Plot	Shows the correlation between two aggregated fields with a third dimension as size	 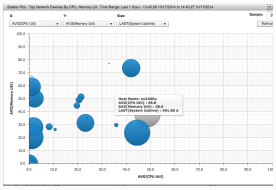	A bubble pot is a scatter plot with a third dimension field added to indicate size. In this example, the same type of search that was used to generate the scatter plot example is run, though the display field **Last (System Uptime)** has been added as a **Size** indicator.

View	Description	Screen Example	Notes

Tree Map	A hierarchical tree-structured visualization that can be used to analyze dominating components of multidimensional data	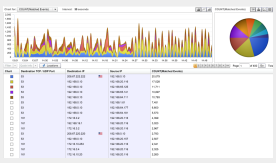 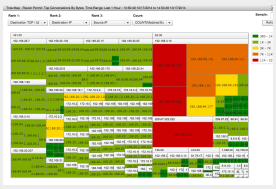 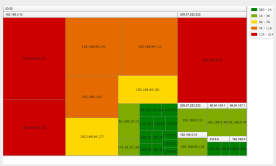	A tree map is a hierarchical tree-structured visualization that you can use to analyze dominant components of multi-dimensional data. A classic example is an attempt to understand Top Talkers in a network.
			
			The results, which run to 400 pages with approximately 10,000 entries, do not provide any information about:
			• The proportion of the Top Destination Port
			• The proportion of Top Source IPs for a given Destination Port
			• The proportion of Top Destination IPs for a given Destination Port and Source IP
			
			By switching to a Tree chart, you can now see:
			• Top ports are 161 (SNMP) and 53 (DNS) - with SNMP taking roughly 1.5 times the connections
			• The top destinations for DNS are:
			  • 192.168.0.10 (Internal DNS)
			  • 208.67.222.222 (External DNS)
			• The top sources going to 192.168.0.10 on the DNS port are 192.168.20.116, 192.168.65.125
			• The top sources going to 208.67.222.222 on DNS port are 192.168.0.10
			
			You can now drill down on port 53 for a closer view by clicking 53.00 in the tree map, which results in the third screenshot in this example.

View	Description	Screen Example	Notes
Heat Map	visualizes calculated measures in two dimensions using a color grade that helps users to understand intensity	 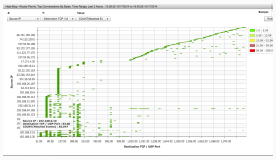 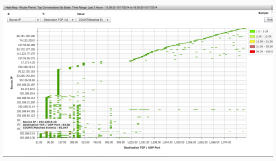	A heat map visualizes two display fields using a color gradient that indicates intensity. A classic example is an attempt to understand which host is talking on which network port. In this example, a search is run with these parameters: • **Filter Conditions**: Group:Permit Traffic • **Group By attributes:** Destination TCP/UDP Port, Source IP • **Display Fields**: Destination TCP/UDP Port, Source IP, COUNT (Matched Events) The first screenshot shows the results as a stacked trend chart. The second shows the results as a heat map with the **Sample** set to 1000. You can now hover your mouse cursor over indicators of higher intensity to view specific information. In this case 192.168.0.10, which appears as a small red bar in the lower left corner, is a heavy contributor to traffic on Port 53. In addition, vertical lines indicate multiple hosts communicating on the same port, for example ports 22, 53, 80, 443, while horizontal lines indicate same host talking across multiple ports.

Refining the Results from Historical Search

Overview of Historical Search Results and Charts describes the charts that you can use to visualize historical search results, but there are also a variety of methods you can use to drill down into search results and refine your queries.

- Charting a Specific Row from Historical Search Results
- Charting Multiple Aggregation Attributes on the Same Historical Search Results Chart
- Drilling Down on Search Results by Time Interval
- Using Search Results to Refine Historical Searches
- Using Tabs to View Multiple Search Results

Charting a Specific Row from Historical Search Results

When your chart loads, the top five items are displayed as color-coded stack charts, as show in the example of this screenshot. However, you may want to remove results from the chart to get a clearer view of what is happening with a specific result. Here, for example, there are spikes for 192.168.19.65 that are clearly visible at various intervals, but the chart results for the other IPs obscure much of what is happening with this source IP.

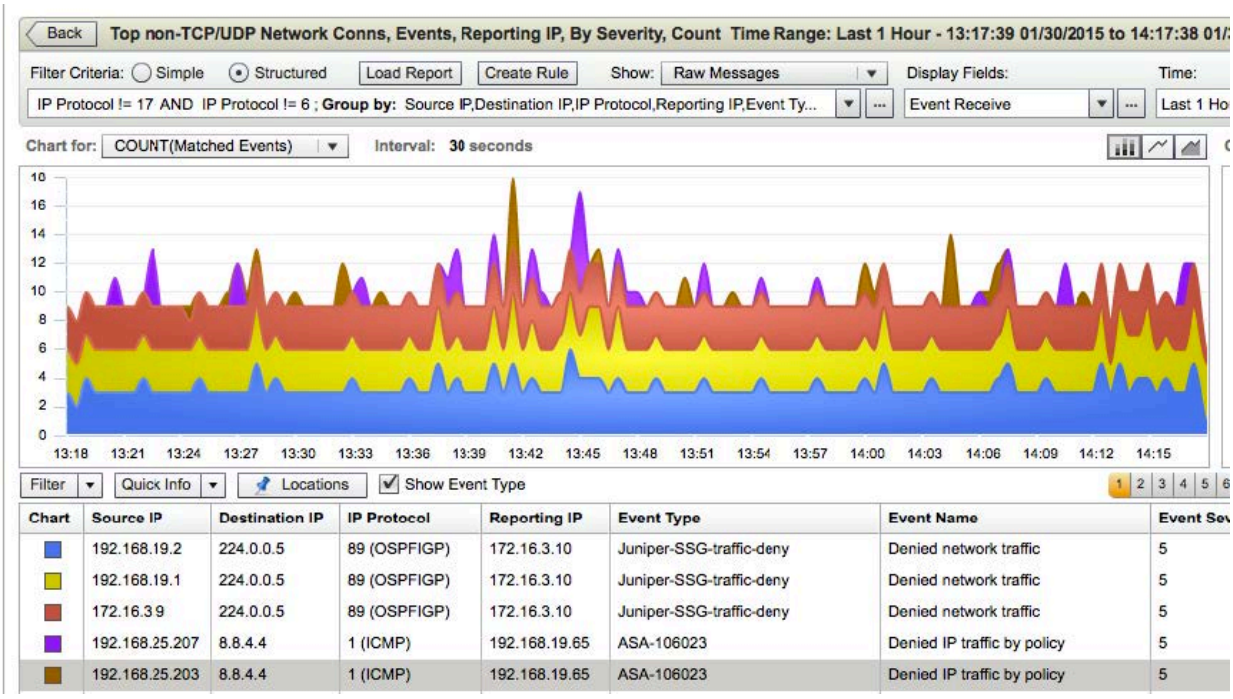

The solution is to remove the other Source IPs from the chart. In the **Chart** column of the **Results List,** click on the items you want to remove from or add to the chart. In this example, all four of the other IPs have been removed from the chart to obtain a clearer visualization of the activity for 192.168.19.65.

Charting Multiple Aggregation Attributes on the Same Historical Search Results Chart

When you run a query, the resulting chart typically displays the first aggregated attribute in the **Results List**. However, if there are other aggregated attribute values in the search results, you can add those to the chart as a second dimension.

This screenshot shows the results for the report **Top Router Network Intf By Util, Error, Discards**, which includes the values for a single aggregated attribute, **AVG(In Intf Util)**, for incoming interface utilization.

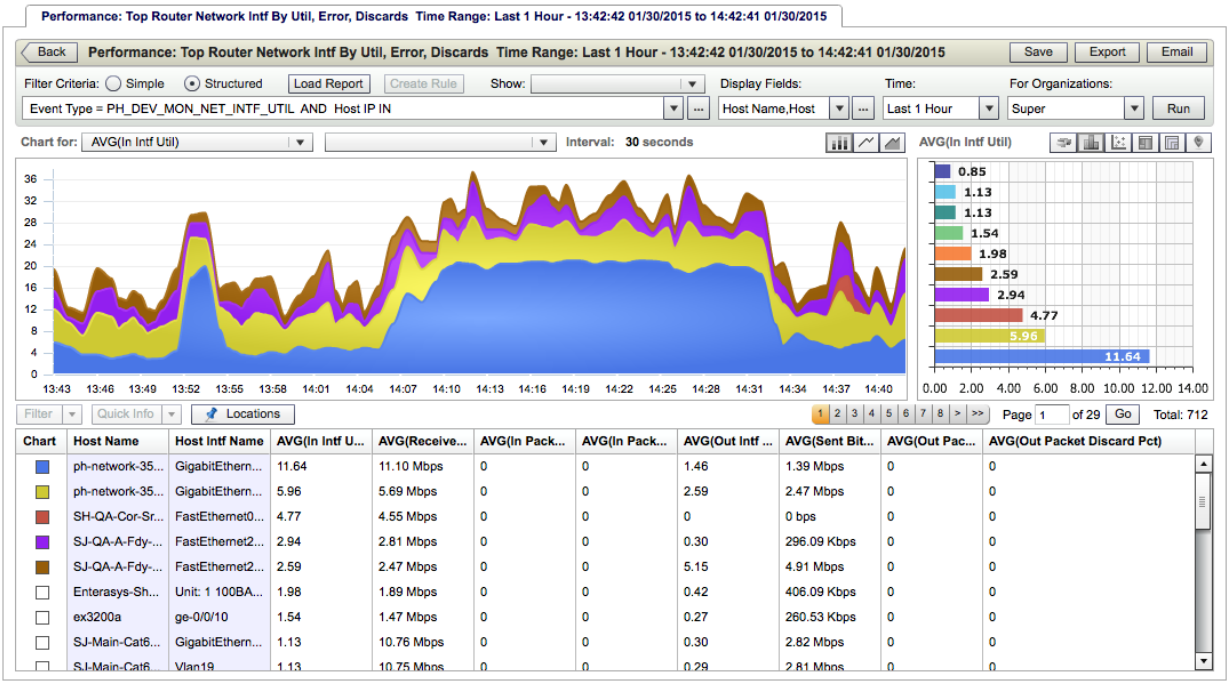

In this case, it could also be informative to understand more about the outbound interface utilization. In the second **Chart For** menu, **AVG(Out Intf Util)** is selected, and this is added as a second dimension to the chart beneath the **0** line, as shown in this screenshot.

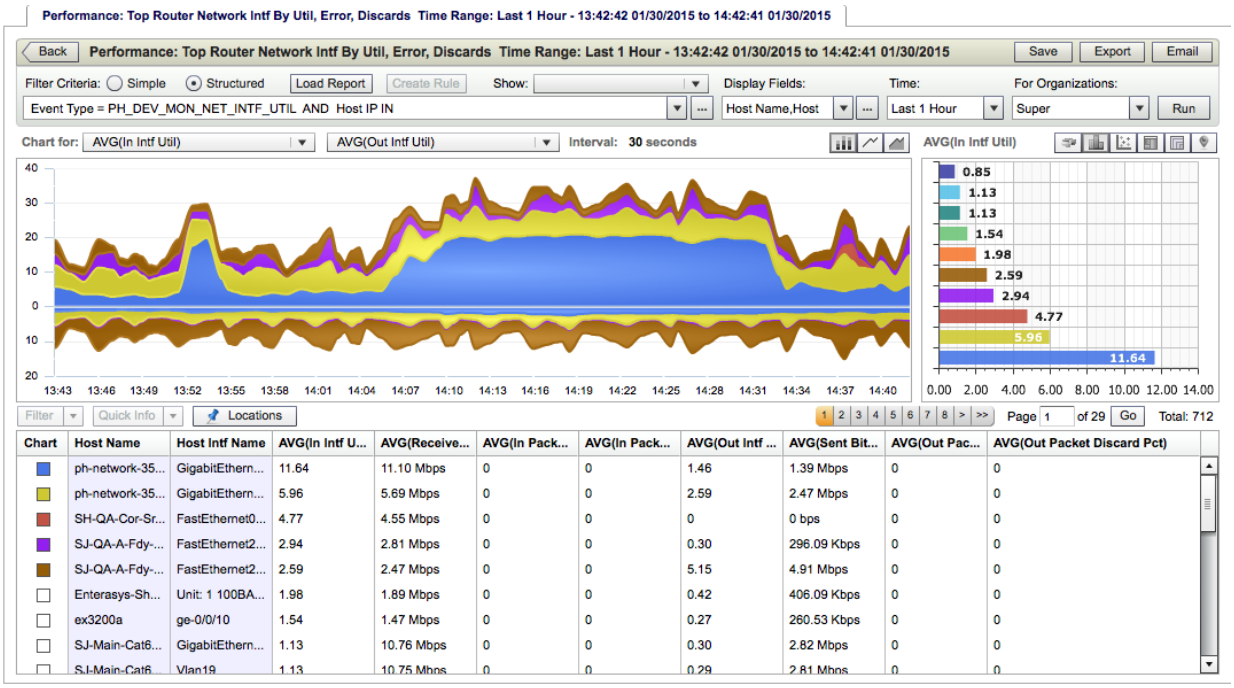

Drilling Down on Search Results by Time Interval

When you run a search, the chart displays results for the time interval you set in your original query. However, you can also drill down to 5 minute, 10 second, and 1 second time intervals for a closer inspection of the results.

1. Hover your mouse cursor over the result and time interval you want to drill down on until the information pop-up appears, as shown in the first example screenshot.
2. Click to drill down and view the results for a 5 minute interval.
3. Follow the same process to drill down to the 10 second and one second intervals.

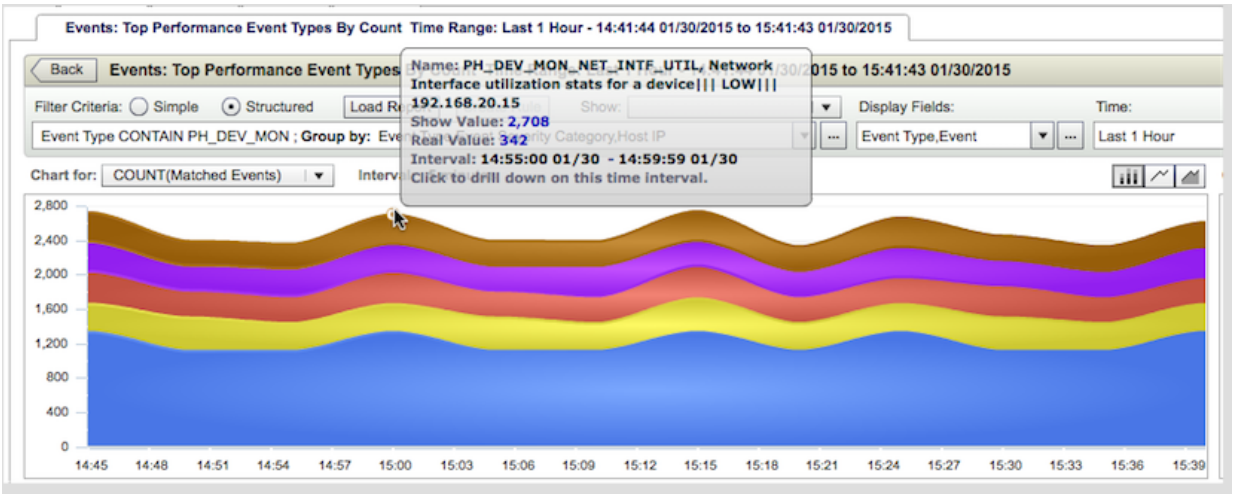

This series of screenshots illustrates starting from the original search results, and then drilling down to the 5 minute interval.

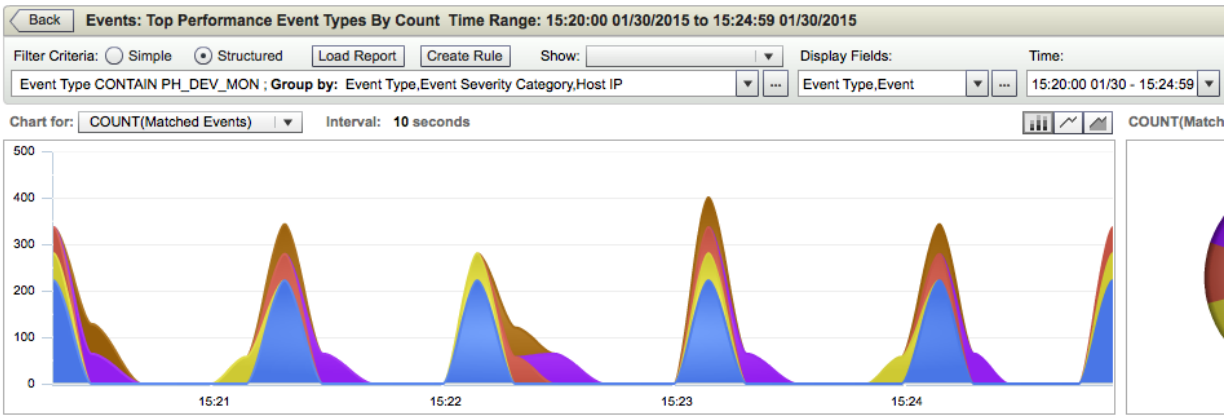
Back
Events: Top Performance Event Types By Count Time Range: 15:20:00 01/30/2015 to 15:24:59 01/30/2015
Filter Criteria: Simple Structured
Load Report
Create Rule
Show:
Display Fields:
Time:
Event Type CONTAIN PH_DEV_MON ; Group by: Event Type,Event Severity Category,Host IP
Event Type,Event
15:20:00 01/30 - 15:24:59
Chart for: COUNT(Matched Events)
Interval: 10 seconds
COUNT(Match
500
400
300
200
100
0
15:21
15:22
15:23
15:24

Using Search Results to Refine Historical Searches

In this screenshot of search results you can see a small but sudden spike in the **SUM(Total Bytes)** for **Destination TCP/UDP Port 20756**, which is represented by the color purple in the chart. In order to understand what is happening in this time interval, you can select this port and the time period of interest, and use these as filter criteria for a deeper investigation.

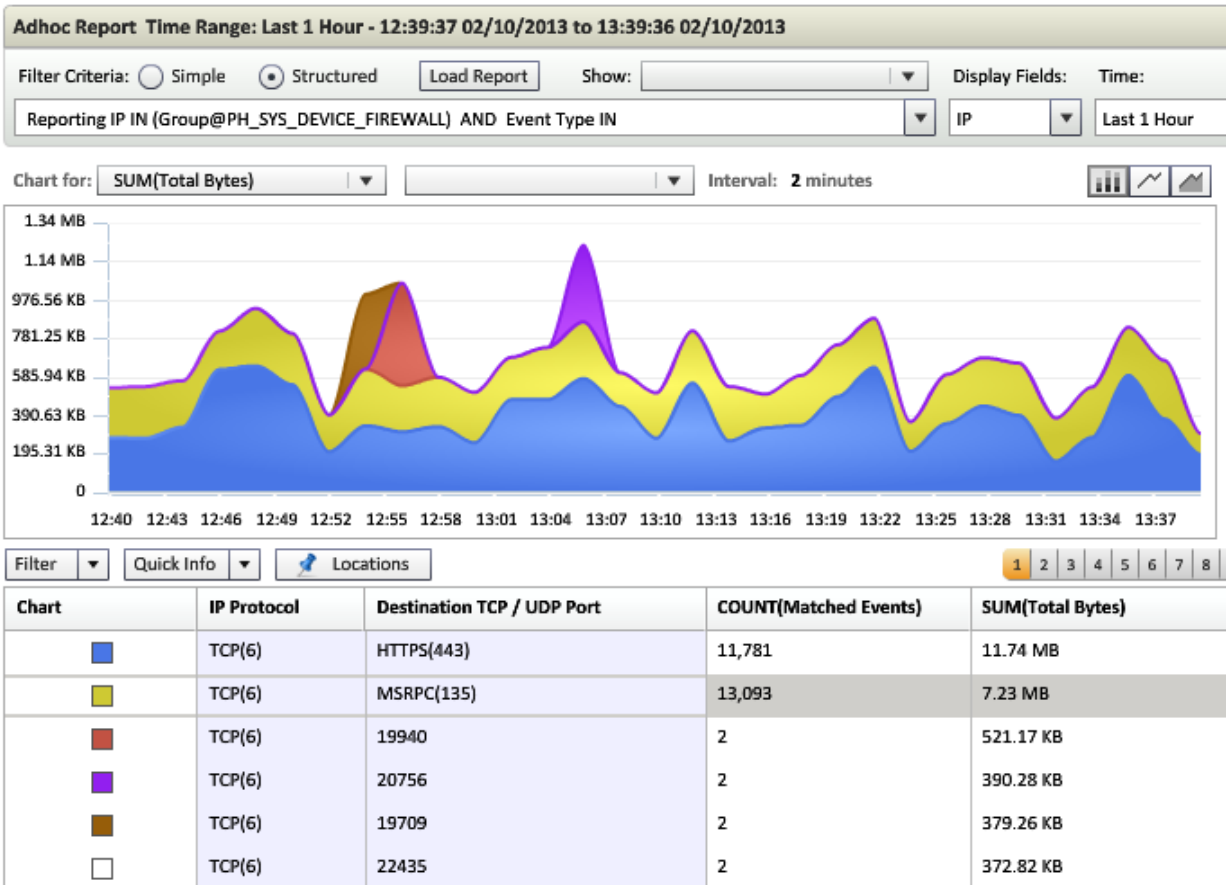

1. In the **Results List**, select the row containing the item of interest.
2. Click the **Filter** menu, and you will see the attributes of the selected item as filter options.
3. Select the attribute you want to use for your filter.
   In this case, you would select **Destination TCP/UDP Port = 20756**.

   Adding a Specific Attribute Value to a Filter

   You can also click in the cell of the **Results List** that contains the attribute value you want to use in your filter, and then select **Add to Filter** from the pop-up menu that appears when you hover your mouse cursor over the attribute value.
4. In the **Show** menu select **Raw Messages.**
   This will include the raw event logs in the **Incident Details**.
5. In the **Display Fields** menu, add or remove any display fields you want for the refined search results.
   In this case two fields are added, **Destination TCP/UDP Port** and **Total Bytes.**
6. In the chart, click on the time period that is of interest to add it to the search criteria.
7. Click **Run**.
   This screenshot shows the results for the selected port and time period, indicating that two events originating from

Seattle WA were responsible for the spike.

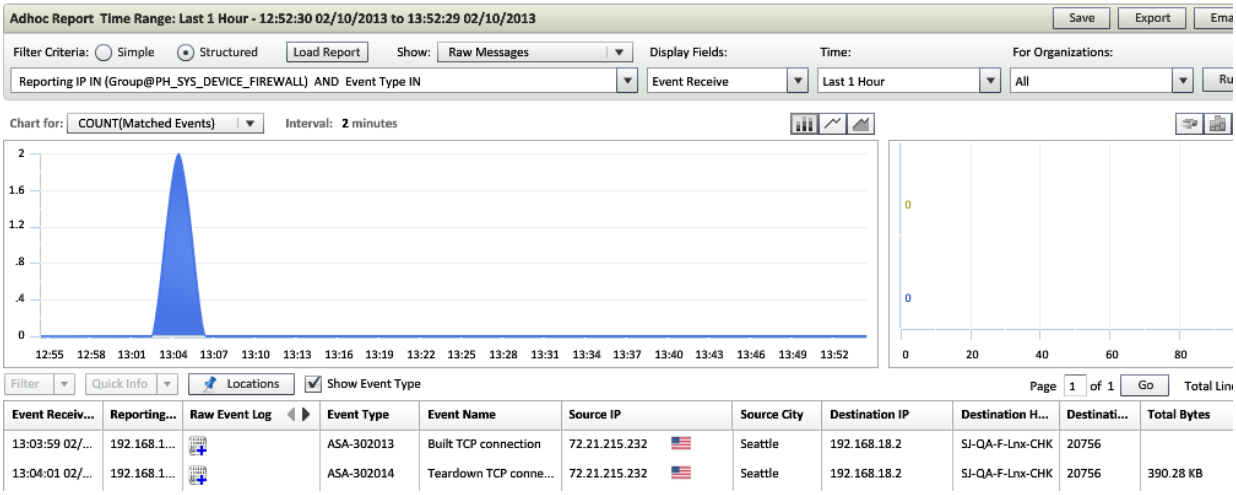

8. Click in the **Raw Event Log** column for an event to view the event details.
   See Selecting Attributes for Structured Searches, Display Fields, and Rules for more information on how to view the attributes for reported events and add them to the display fields for your results.

Using Tabs to View Multiple Search Results

There may be occasions when you want to be able to run and compare the results of multiple searches.

1. Run your first search.
2. In the upper-left corner of the search screen, click **+**.
   A new tab will open up in the Analytics Window.
3. Run your second search in the new tab.

New Tabs for Drill-Down and Refined Searches

If you refine an existing search, zoom in on a time period, or use the time interval drill-down to examine search results, new tabs are automatically generated for each level of drill down, and for each refined search. When you select an attribute to use in a refined search, you can also select **Add to Filter in New Tab** from the **Options** menu.

Converting an Historical Search to a Real Time Search

In the course of running an historical search, you may produce results that you want to examine in real time. For example, suppose that an historical search shows that yesterday there was an excessive amount of outgoing traffic from your home country or countries that you do business with. You may want to know if this same traffic pattern is happening right now, in real time. You can answer this question from within the same historical search that raised your suspicions.

1. In the historical search window, click **Real Time Search**.
   The historical search criteria are loaded into a Real Time Search window and begin to execute.
2. You can now refine your Real Time Search results to reflect your current interest, for example by adding a Destination County attribute to the display results and running the search again.

Converting an Historical Search to a Rule

Example

While using historical search, you may observe a pattern that you want to use as a rule so if the pattern recurs, it will trigger an alert. For example, in an historical search you may notice excessive traffic going outside your country or the countries you do business with. You can generate a rule to watch for this traffic pattern from within the historical search.

These screenshots show the conditions and results for the example of an historical search for excessive outgoing traffic.

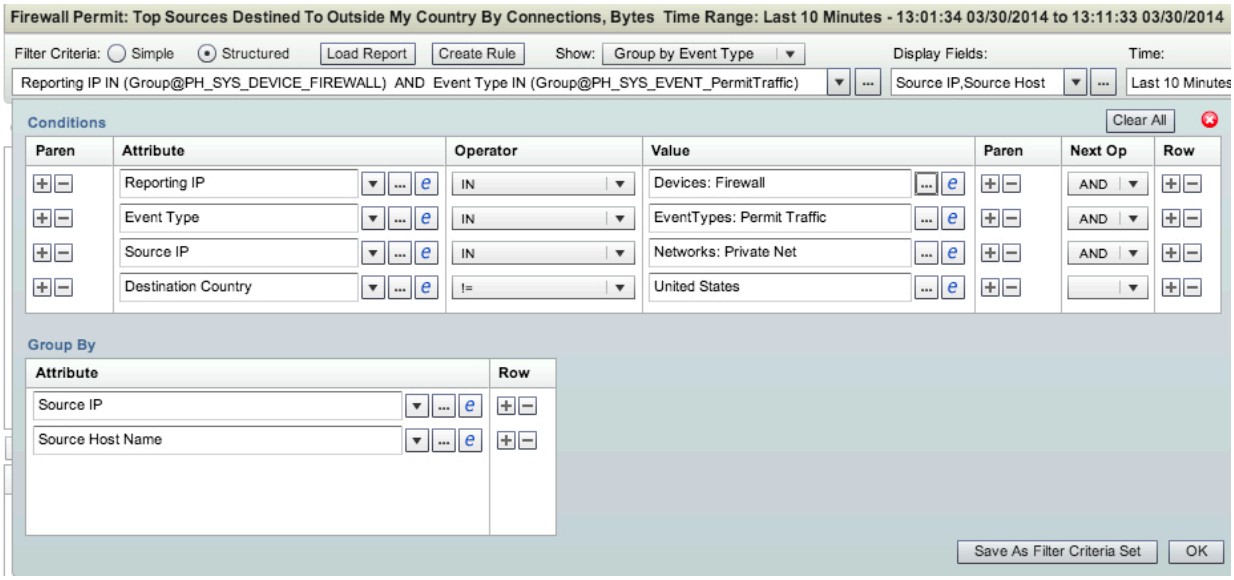

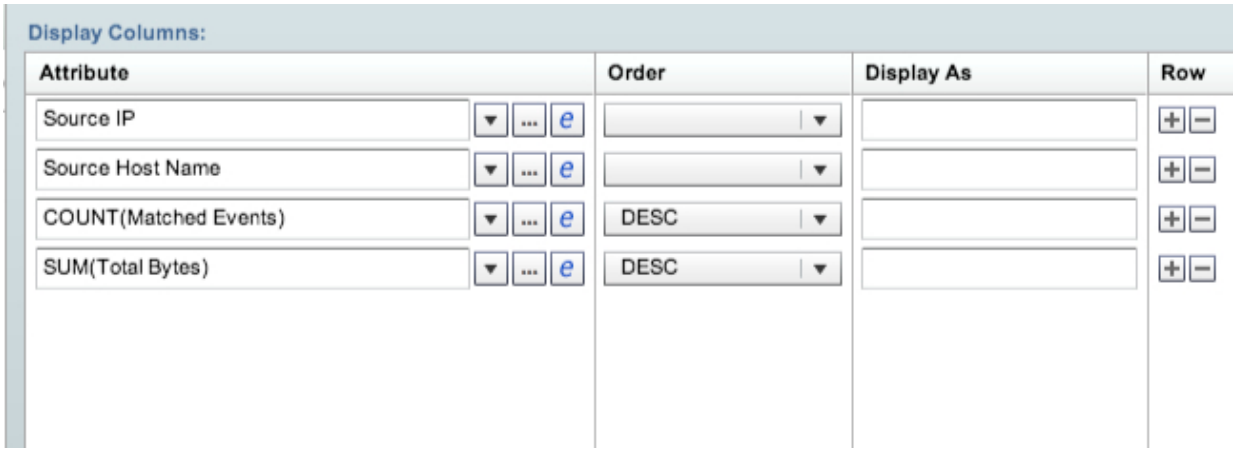

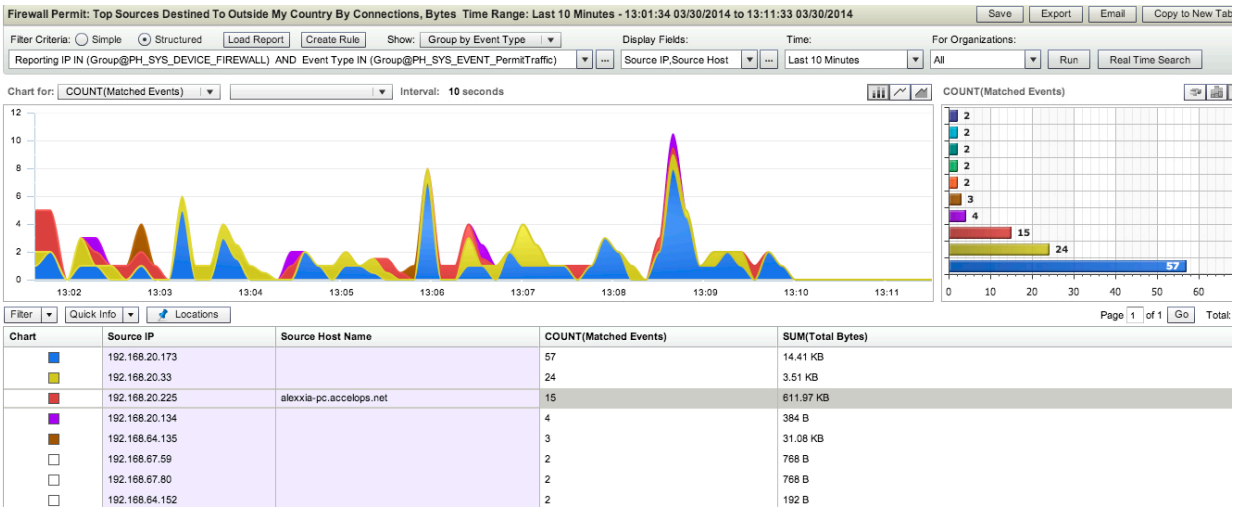

Following this example, you may now want to create a rule that will send you an alert when a particular source sends more than 1000 connections, or more that 5MB of traffic, in five minutes.

Procedure

1. In the historical search that you want to use as the basis for your rule, click **Create Rule**.
   The **Rule Editor** will load, with most information for the rule auto-populated from the search. You can also read the topics under Rules for more information about creating rules.
2. Enter a **Rule Name** and **Description**.
3. Set the **Severity** to associate with incidents generated by this rule.
4. Set the **Incident Category** to associate with incidents generated by this rule.
5. Set the number of seconds for the **Time Window** that this rule should apply to.
   In the example of excessive outgoing traffic over a five minute period, this would be set to **300**.
6. Under the **Conditions**, click the **Edit** icon for **Filter_1**.
   You will see that all your filter conditions for the search have been populated into this sub pattern.
7. You can now edit the Filter and Aggregate conditions for your original search, or change the Group By conditions.
8. Click **Save** when you're done editing the rule.

This screenshot show editing the rule sub pattern **Filter_1** from the original rule conditions, with the **Aggregate Conditions** for **COUNT(Matched Events)** and **SUM(Total Bytes**) to **1000** and **5242880** to match the new alert conditions from the example historical search, and the **AND** operator changed to **OR.**

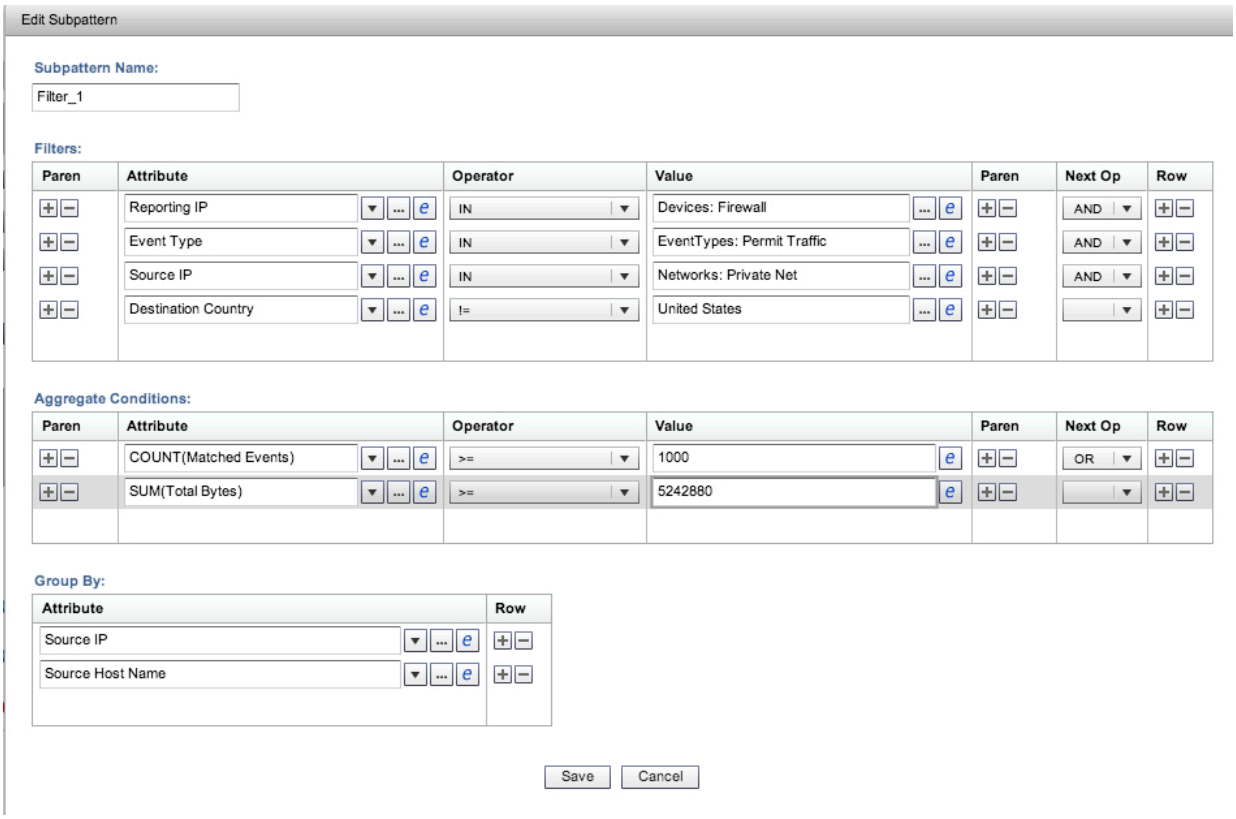
Edit Subpattern
Subpattern Name:
Filter_1
Filters:
Paren
Attribute
Operator
Value
Paren
Next Op
Row
Reporting IP
IN
Devices: Firewall
AND
Event Type
IN
EventTypes: Permit Traffic
AND
Source IP
IN
Networks: Private Net
AND
Destination Country
!=
United States
Aggregate Conditions:
Paren
Attribute
Operator
Value
Paren
Next Op
Row
COUNT(Matched Events)
>=
1000
OR
SUM(Total Bytes)
>=
5242880
Group By:
Attribute
Row
Source IP
Source Host Name
Save
Cancel

Real Time Search

You can use Real Time search to view events as they are occurring in real time within your IT infrastructure. You can use both simple and structured search criteria, as you would with historical search, but instead of the results displayed in a report like you would see with an historical search, real time search results are displayed as a rolling graph and summary of events that you can drill down into.

- Overview of the Real Time Search User Interface
- Creating a Simple Real Time Search
- Creating a Structured Real Time Search
- Viewing and Refining Real Time Search Results

Overview of the Real Time Search User Interface

The real time search interface is very similar to the interface for historical search, with the exception that real time search doesn't have an option to set a search time period. As with historical search, you can also run simple or structured search queries. The main difference between historical and real time search is that real time search displays your results as they are occurring in your IT infrastructure, with a scrolling chart and summary of the results.

- Simple Real Time Search
- Structured Real Time Search

Simple Real Time Search

When you use simple real time search, you enter a keyword to search for in the logs collected by FortiSIEM, set any columns you want to display in the Raw **Event Log Results Summary**, and, for multi-tenant deployments, select any organizations you want to filter the results for. You can then select results in the real time chart to use for historical searches, or you can select results in the Raw Event Log Results Summary to learn more information about them or use them as filters in refining your search.

This screenshot shows the results for searching the raw event logs for occurrences of **TCP**.

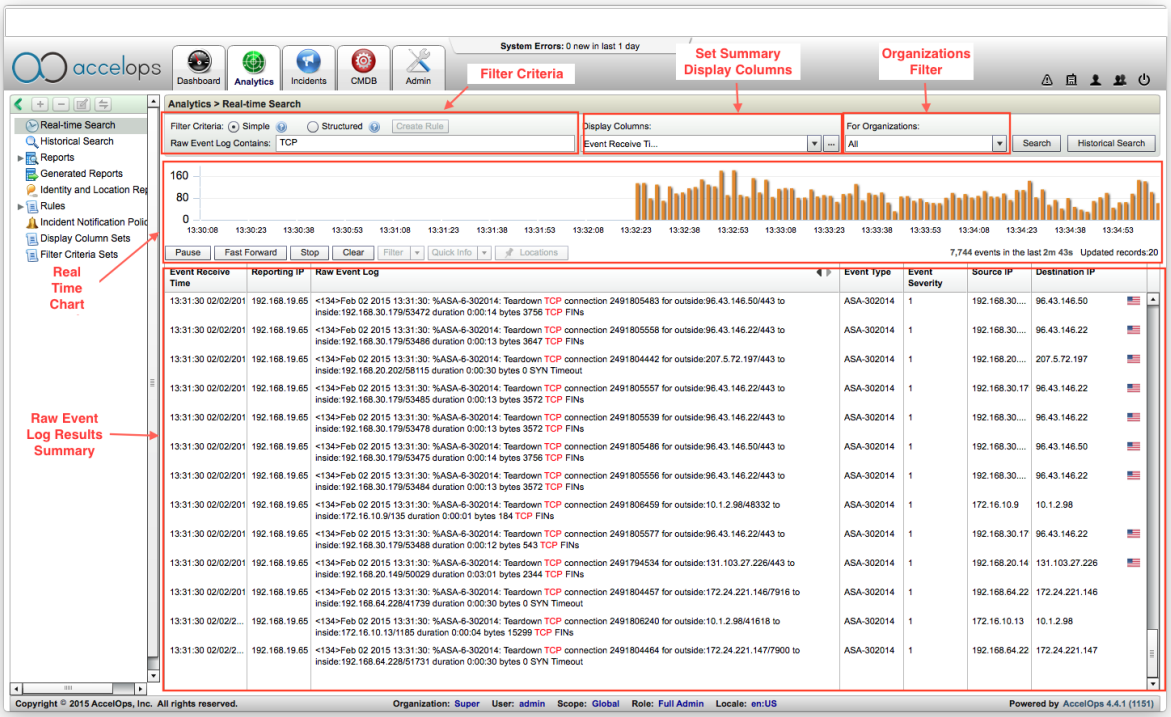

Simple Real Time Search Interface Controls

Ui Control	Description
Filter Criteria	For simple real time search, use the search box to find keywords in raw event logs. You can also create a rule from your search results.
Set Summary Display Columns	Select which columns will be displayed in the **Raw Event Log Results Summary**
Organizations Filter	For multi-tenant deployments, select which organizations you would like to filter the results for
Real Time Chart	Displays results as they occur in real time. Use the Pause, Fast Forward, Stop, and Clear buttons to control the display.
Raw Event Log Results Summary	Displays a summary of the raw event logs for your search results in real time. Click **Pause** in the real time chart and then select an item in the summary results to view attributes such as Reporting and Destination IP, add an IP address to a watch list, add an attribute as a search filter, or get topological information about network devices. Selecting a result from the summary list also enables the **Filter**, **Quick Info**, and **Locations** buttons.

Structured Real Time Search

For structured real time search, you only enter the filter conditions that you want to use, instead of having to also specify aggregation and group by conditions as you would in a structured historical search.

This screenshot shows the **Conditions** dialog for structured real time search. You can select attributes and create expressions to use in structured real time search the same way you would in structured historical search.

This screenshot shows the **Conditions** dialog after having selected **Structured** in the search controls, with two search conditions set.

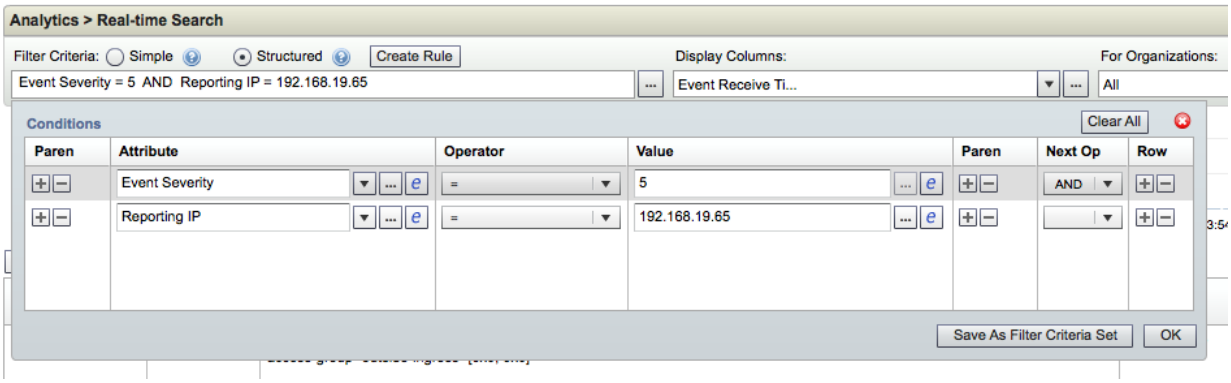

Creating a Simple Real Time Search

1. Log into your Supervisor node.
2. Go to **Analytics > Real Time Search**.
3. In **Filter Criteria**, select **Simple**.
4. Enter the keywords you want to search for in the raw event logs collected by FortiSIEM.
   See Keywords and Operators for Simple Searches for more information about keyword searching.
5. Select the **Display Fields** for the results summary.
   See Selecting Attributes for Structured Searches and Display Fields for more information about selecting attributes that can be displayed for reported events.
6. For Service Provider deployments, select any **Organizations** that you want to filter the results for.
7. Click **Search**.

Related Links

- Keywords and Operators for Simple Searches
- Selecting Attributes for Structured Searches, Display Fields, and Rules

Creating a Structured Real Time Search

1. Log in to your Supervisor node.
2. Go to **Analytics > Real Time Search**.
3. For **Filter Criteria**, select **Structured**.
   The **Conditions** search window will open.
4. Click the downward arrow in the search window to open the **Conditions** options.
   Alternatively you can click **...** to use a saved **Filter Criteria Set**.
5. Under **Conditions**, set the **Attribute**, **Operator**, and **Value** for your condition.
   You can also use expressions as search conditions. See Using Expressions in Structured Searches and Rules for more information, and Selecting Attributes for Structured Searches, Display Fields, and Rules for more information about using attributes in conditions.
6. Click **+** under **Row** to add another condition, and set the **Next Operator** to use for that condition.
   You can give precedence to conditions by setting parentheses around them with the **+** button under **Paren**.
7. Click **OK**.
   You can also click **Save as Filter Criteria Set**, and these conditions will be available for future searches by clicking **...** next to the search window.
8. Under **Display Fields**, select the attributes you want to use as the columns in your results list.
   See Selecting Attributes for Structured Searches, Display Fields, and Rules for more information about selecting attributes for devices and events to use as display fields.
9. For Service Provider deployments, select the **Organization** you want to run the search against.
10. Click **Search**.
    The results of your search will appear in the real time chart and results list.

Viewing and Refining Real Time Search Results

When your real time search runs, you will see the results represented as a scrolling chart across the top of the search results window, and as a scrolling list in the bottom of the window that include the raw event log information for events matching your search criteria. You can select items in the scrolling chart to use in historical search, view more information about individual items in the results list, and add attributes from your search results to your search filters or display fields.

- Selecting Results for Historical Search
- Viewing Information about Real Time Search Results
- Adding Search Results to Search Filters. Watch Lists, or Display Field

Selecting Results for Historical Search

1. When you see a time interval of events that you want to use for historical search appear in the scrolling chart, click **Pause** or **Stop**.
2. Hover your mouse cursor over the bar that represents the time interval until you see the time interval information appears, and then double-click on the bar.
3. The time interval and Event Type will be added to the criteria for an historical search.
   Complete the other criteria you want to use for the search as described in Historical Search.

Viewing Information about Real Time Search Results

1. When you see an event appear in the search results list that you want more information about, click **Pause or Stop**.
2. Select the event row and click **Quick Info** to view the **Reporting IP**, **Event Type**, **Source IP,** and **Destination IP** for that event.
3. To view information about specific attributes of an event, click in the attribute display field and click **Quick Info**. For attributes associated with devices, this will open the **Quick Info** view of the device as described Overview of the Summary Dashboard User Interface. For events types, it will show info such as the severity and device associated with the even type.
4. To view information about a device's location in the network topology, select it in the display field and then select **Topology**.

Adding Search Results to Search Filters. Watch Lists, or Display Fields

- With a search result selected in the results list, click **Filter** to select event attributes to add to the search filter.
- In the expanded **Raw Events Log**, click on items in the text string to include or exclude them as search filter criteria.
- To add a specific result to the search criteria, in the results list, click on an item in a display field to open the options menu, and then select **Add to Filter**.
- To add an IP address to a watch list, click on it to open the options menu, and then select **Add to Watch List**. See Watch Lists for more information.
- See the section on **Selecting Attributes from the Raw Event Log Column in the Results Lists** in the topic Selecting Attributes for Structured Searches and Display Fields for information on how you can view and select the attributes associated with events to use as search filters or display fields from the real time search results list.

Structured Search Operators

Operator	Meaning	Allowed on Event Attribute Types or CMDB Group	Example as seen in GUI
=, !=	Compares whether an attribute is exactly identical or not identical to a specified value.	All except DATE types	Event Type = "PH_DEV_MON_SYS_CPU_UTIL" Source IP != 10.1.1.1
>, >=, <, <=	Compares whether an attribute is less or greater than a specified value	Numeric types: UINT16, UINT32, UINT64, DOUBLE	CPU Util > 10
IN, NOT IN	Determines whether an attribute belongs or does not belong to a set of values. For string valued attributes, the match is case insensitive.	All except DATE type Allows CMDB Groups	System Event Category IN (3,6) Event Type IN ("PH_DEV_MON_SYS_CPU_UTIL","PH_DEV_MON_SYS_MEM_UTIL") Event Type IN ("PH_DEV_MON_SYS_CPU_UTIL",Event Types:Login Failure) Source IP IN Devices:Windows, Devices:Unix Destination IP IN Networks:VPN Pool
BETWEEN, NOT BETWEEN	Determines whether an attribute is between a range of values	All except STRING types	Source IP BETWEEN (10.1.1.1, 10.1.1.255) CPU Util BETWEEN (20.0, 30.0) Event Receive Time BETWEEN (18:35 03/17/2014, 18:35 03/26/2014)
IS (NULL), IS NOT (NULL)	Determines whether an attribute is present or not	All types	Host Name IS NOT NULL

Operator	Meaning	Allowed on Event Attribute Types or CMDB Group	Example as seen in GUI
CONTAINS, NOT CONTAINS	Determines whether a string valued attribute contains a specified sub-string. • For Raw Event Log - the sub-string has to contain the beginning of every word • For all other string type attributes: the sub-string can be in any position	STRING	Event Type CONTAINS "DEV_MON" matches "PH_DEV_MON_CPU" Event Type NOT CONTAINS "DEV_MON" does not matche "PH_DEV_MON_CPU" Reporting Model CONTAINS "dows" matches "Microsoft Windows" Reporting Model CONTAINS "soft win" matches "Microsoft Windows" Raw Event Log CONTAINS "dows" does not match "Microsoft Windows" Raw Event Log CONTAINS "microsoft win" matches "Microsoft Windows 2003" (For more general patterns use regular expressions)
REGEXP, NOT REGEXP	Determines whether a string valued attribute matches a specified pattern. Raw message needs to be UTF-8 encoded.	STRING	Raw Event Log REGEXP "\d+.\d+\d+.\d+" Event Type NOT REGEXP "PH_DEV_MON_.*" - match events with event types not beginning with PH_DEV_MON

Selecting Attributes for Structured Searches, Display Fields, and Rules

For both Real Time and Historical structured searches you have the option to to select event attributes to use in both your search and **Group By** filters, and as display fields in your result lists. Since FortiSIEM recognizes over 130,000 event attributes, the documentation and user interface provides several ways to find the attributes you want to use. These instructions show how to access the **Common Attributes** menu and the CMDB attribute browser through the **Attributes** in search conditions, but you can access the same functionality in the **Display Fields** menu for searches, and when you create a new rule. They also contain information on how you can access the attributes associated with reported events through the **Raw Event Logs** column of results lists.

- The Event Dictionary and Master Attribute List
- Selecting Attributes in the Common Attributes Menu
- Selecting Event Attributes from the CMDB
- Selecting Attributes from the Raw Events Log Column of the Results Lists

The Event Dictionary and Master Attribute List

This documentation includes an Event Dictionary that describes events and their attributes, and an attribute master list, which lists the primary event attributes and their data type, along with a brief description of what values FortiSIEM expects to see when that attribute information is returned.

Selecting Attributes in the Common Attributes Menu

This screenshot shows the **Common Attributes** menu open in the Conditions Builder for an Historical search. Open the menu by clicking the downward arrow next to an **Attribute** text field. You can scroll through the list of event attributes to select the one you want, or begin typing an attribute name and the menu will sort based on your entry.

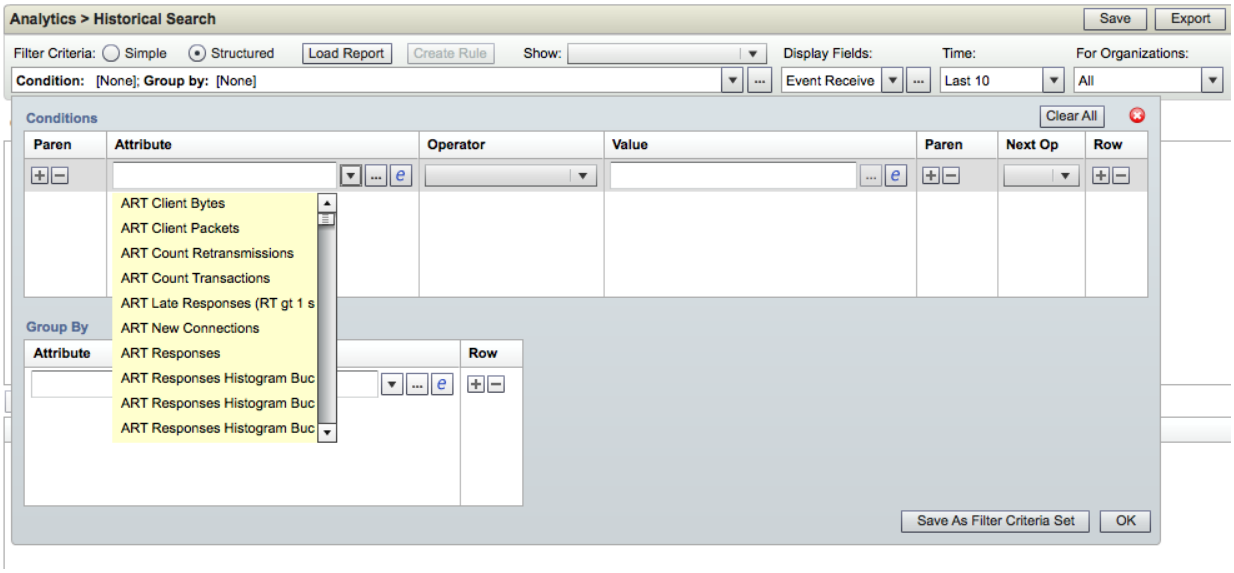

Selecting Event Attributes from the CMDB

You also have the option to browse all the attributes listed in the CMDB to find the one that you want. These two screenshots show the CMDB attribute browser, which you can access by clicking **...** next to the **Attribute** text

field.

The first screenshot illustrates browsing the CMDB attributes based on **Device Type** and **Feature Type**: **Availability**, **Change**, **Performance**, **Security**, and **All**. In this example, **Security** has been selected for **Feature Type**, and **Cisco IOS** has been selected for **Device Type**. This loads all the security attributes associated with the Cisco IOS into the Attribute List.

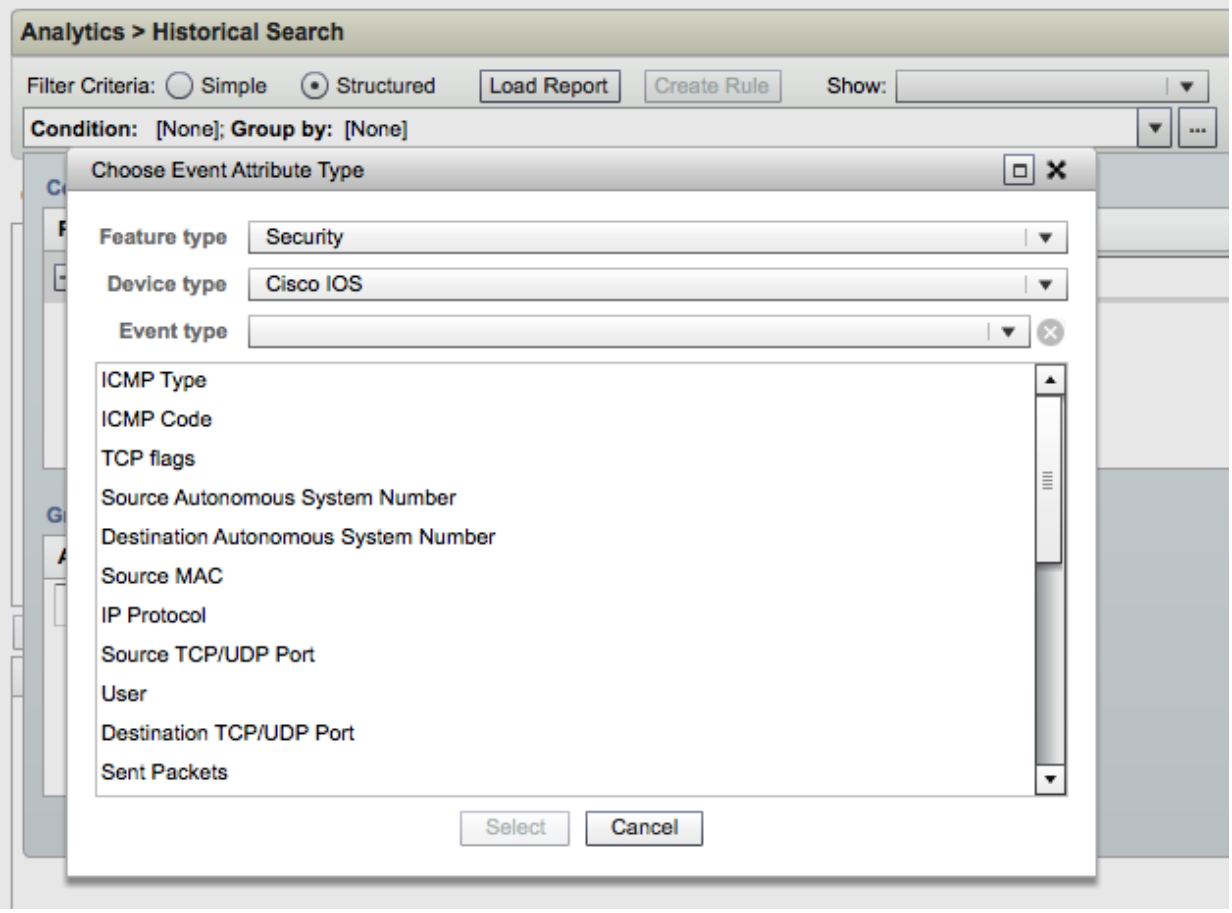

The second screenshot illustrates browsing the CMDB Event Types to find an event attribute. In this example, **Cisco ASA** is selected for **Device Type**. Clicking in the Event Type window opens an Event Browser for the CMDB. Select any group in the browser, and you will see the event types within that group that are applicable to the Device Type you selected.

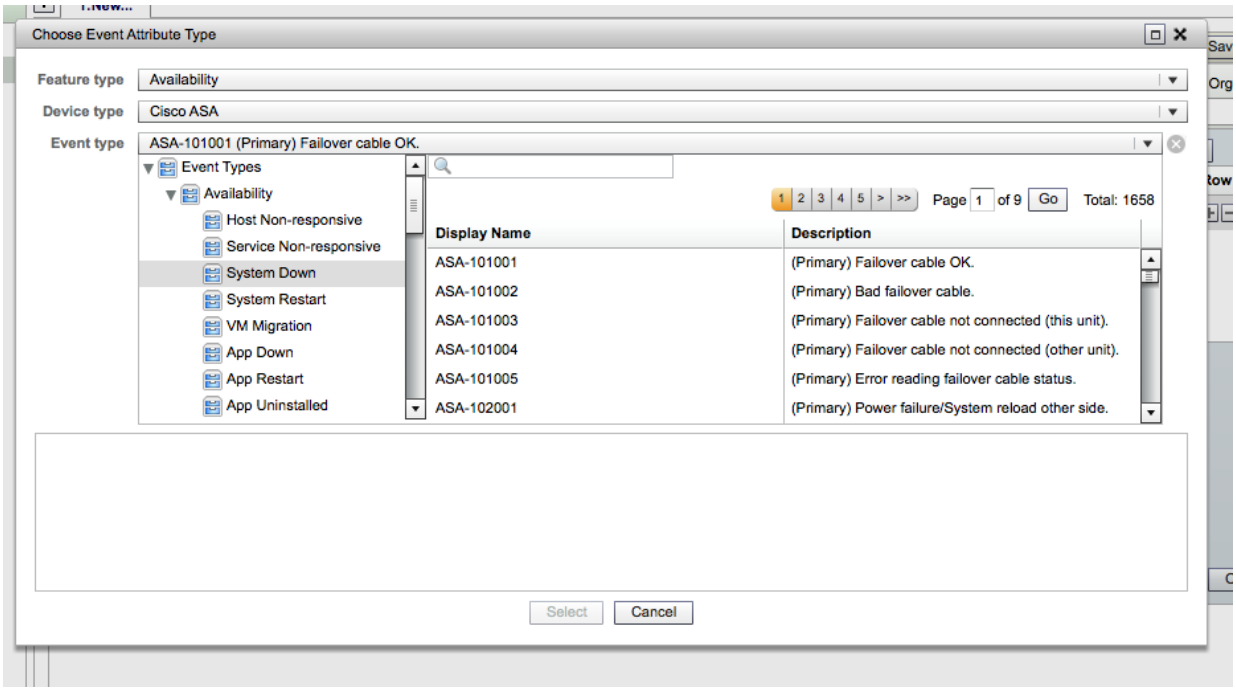

Selecting Attributes from the Raw Events Log Column of the Results Lists

All real time search results lists include a Raw Event Log column, and you can add a a Raw Event Log column to the list of results for historical searches. In addition to providing detailed information from the raw event logs, you can also use this column to view all the attributes associated with a reported event and add them to the display fields in your results list or to your filters for structured searches.

1. Cilck in the **Raw Event Log** column of your results list to collapse the view.
   The raw event log text will collapse into an information icon with a blue **+**.
2. Click on the blue **+** icon to open the **Event Details**.
   You will see the raw event log text and list of all the attributes associated with that event type.
3. Select **Filter** or **Display** to add an attribute to the search filters or display fields for that search.
4. Click **X** to close the Event Details window when you're done making your selections.

Using Expressions in Structured Searches and Rules

An expression can contain a single event attribute, multiple attributes, or functions that contain an event attribute as their argument. You can also use parentheses and arithmetic operators to form complex expressions.

You can enter an expression manually, paste it in, or build it dynamically using the Expression Builder. If you use the Expression Builder, you will have to enter parentheses or arithmetic operators in the expression.

- The Expression Builder
- Creating Expressions

The Expression Builder

You can access the Expression Builder by clicking the **e** icon next to the **Attribute** or **Value** field when creating a structured search or rule.

This screenshot shows the Expression Builder open for creating a rule.

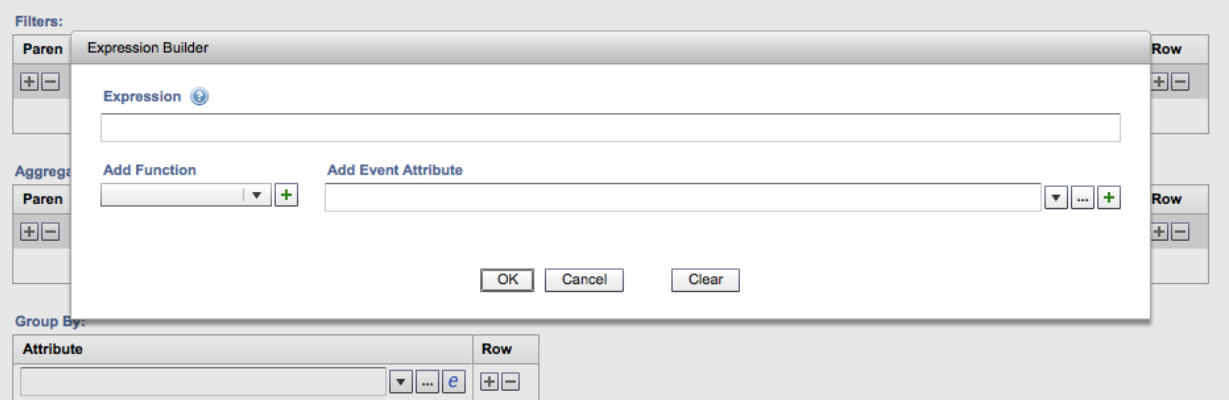

Creating Expressions

Adding a Function

To add a function to the expression, select it from the **Add Function** menu, and then click the **+** icon. The available functions depend on whether you are are creating an expression to use as part of a filter condition for a search or rule, or as part of the aggregation conditions for a rule.

Selecting Function-Specific Attributes

When you select any type of function, the function and a set of parentheses will be added to the expression. If you place your cursor within the parentheses and then open the **Event Attribute** menu, you will see event attributes that are relevant for that function. For example, if you select COUNT as the function, (MATCHED ITEMS) will automatically appear between the parentheses, and will be selected in the Event Attribute menu. If you select a function like **AVG** for an aggregation condition, you will see options such as **CPU UTIL** and **Apache Uptime**. If you select a function like **HourOfDay** for a filter condition, you will see options like **Access Time** and **Vulnerable Since**. You can search through the options in either situation by beginning to type a keyword in the Event Attribute menu. Selecting Attributes for Structured Searches, Display Fields, and Rules has more information about ways to search for and select event attributes.

Filter Condition Functions

If you select HourOfDay or DayOfWeek for the function, the **Event Attributes** menu will contain date and time-related event attributes, while if you select **DeviceToCMDBAttr**, it will contain device-related attributes.

Function	Description
HourOfDay	Specify an hour of the day in the condition
DayOfWeek	Specify a day of the week in the condition
DeviceToCMDBAttr	If you add the **DeviceToCMDBAttr()** function to the expression, the first argument must be an event attribute, and the second argument must be a CMDB attribute, which you can select using the **CMDB Attribute** menu. The DeviceToCMDBAttr function is used to create expressions for per-device thresholds.

This screenshot shows the beginning of creating an expression to use as the **Attribute** in a condition for an historical search. **HourOfDay** is selected as the **Function**, and **Access Time** is selected as the **Event Attribute**.

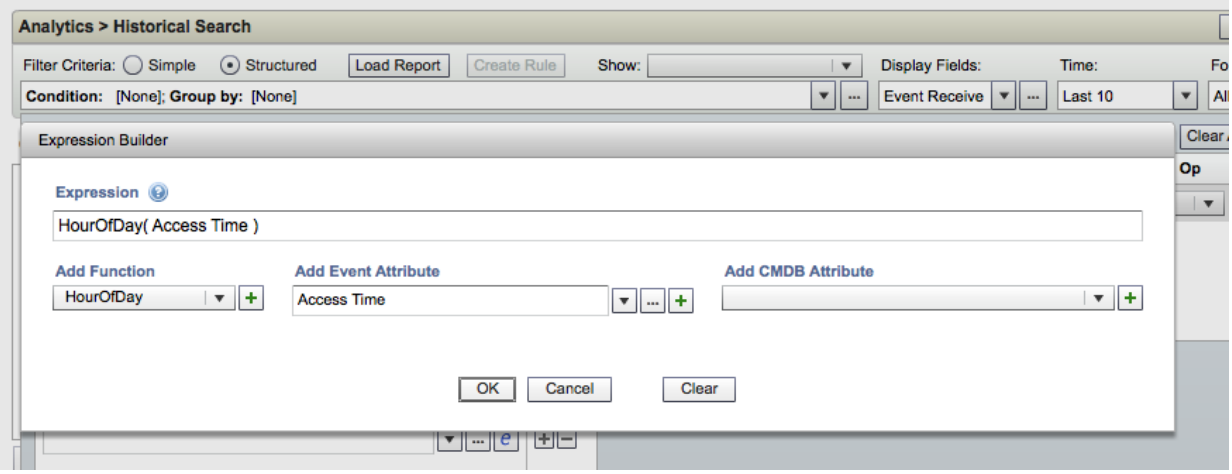

Aggregation Condition Functions

You use these functions to perform operations on numerical event attributes such as **Sent Bytes**, **Received Bytes**, **CPU Utilization**, or **Memory Utilization**.

Function	Description
Count	Count the number of items returned
Count Distinct	Count the number of distinct items returned
Sum	Add the numbers
Average	Average the numbers
Min	The lowest number
Max	The highest number
Last	The last number
First	The first number
Pctile95	The 95th percentile
PctChange	Percentage change
STAT_AVG	Statistical average. This function is used in conjunction with creating baseline reports.
STAT_ STDDEV	Statistical standard deviation. This function is used in conjunction with creating baseline reports .

This screenshot shows the beginning of creating an expression to use as an aggregation condition in rule. **Max** is selected as the **Function**, and **CPU Util** is selected as the **Event Attribute**.

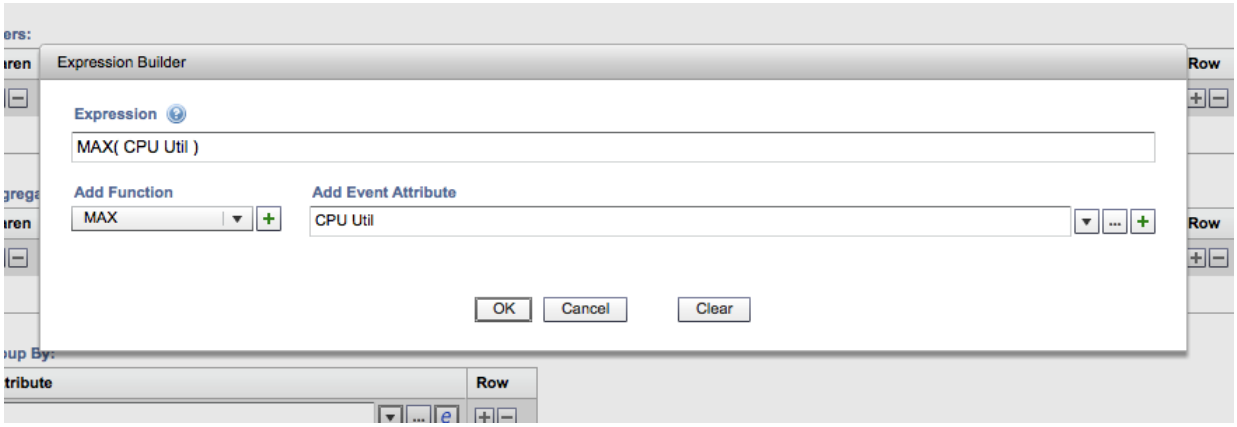

Keywords and Operators for Simple Searches

Both historical and real time searches have a simple search option that searches for keywords in the raw ASCII tex of event logs. You can use operators in your keyword searches to combine terms or create simple search filters.

- Keyword Operators
- Quotes and Backslash Characters in Search Terms

Keyword Operators

You can use the operators **AND**, **OR**, **AND NOT** between keywords. If you enter more than one keyword, then **AND** is assumed as the operator between them. You can also use parentheses **()** to change the precedence of the operators.

Examples of Using Keyword Search Operators

Search String	Results
TCP	Finds all events with **TCP** in the event logs
TCP 80	Finds all events with **TCP** and **80** in the event logs
TCP AND (80 OR 443)	Finds all events with **TCP** and **80** or **40** in the event logs
TCP AND NOT 80	Finds all events with **TCP** but not **80**

Quotes and Backslash Characters in Search Terms

If the search string contains quotation marks or back-slash characters, you must escape them by prefixing them with a backslash character. For example, if you wanted to search for `[location]="United States"` then you would need to enter `[location]=\"United States\"` as your search string.

Using Geolocation Attributes in Searches and Search Results

When you view the results of a search, you will see that IP address fields in the results, such as **Source IP** or Destination IP, often have a flag added to them to indicate the geolocation of that IP address. This topic describes the geolocation information that is associated with event attributes, and provides examples of how to use this information in searches and search results.

- Event and Geolocation Attributes
- Using Geolocation Attributes in Searches
- Viewing Geographic Locations from Search Results

Event and Geolocation Attributes

The event attributes **Source IP**, **Destination IP**, **Host IP**, and **Reporting IP** include geolocation attributes that you can use in search queries and as display fields in search results. In **Incident Reports** you may also see country flags included with IP addresses for **Incident Source** and **Incident Target**, which have the same geolocation attributes as **Source IP** and **Destination IP.**

Event Attribute	Geolocation Attributes
Source IP	• Source Country • Source City • Source State • Source Organization • Source Longitude • Source Latitude
Destination IP	• Destination Country • Destination City • Destination State • Destination Organization • Destination Longitude • Destination Latitude
Host IP	• Host Country • Host City • Host State • Host Organization • Host Longitude • Host Latitude
Reporting IP	• Reporting Country • Reporting City • Reporting State • Reporting Organization • Reporting Longitude • Reporting Latitude

Using Geolocation Attributes in Searches

You can use geolocation attributes in both real time and historical structured searches. For example, setting a search attribute to **Source Country != United States** will remove all Source IPs with a geolocation of United States from the search results.

This screenshot shows the results of using **Source Country != United States** and **Event Severity = 1** as the search criteria. The Source IP display field contains only IP addresses associated with countries other than the United States, as indicated by the national flags next to each IP address in the Source IP column.

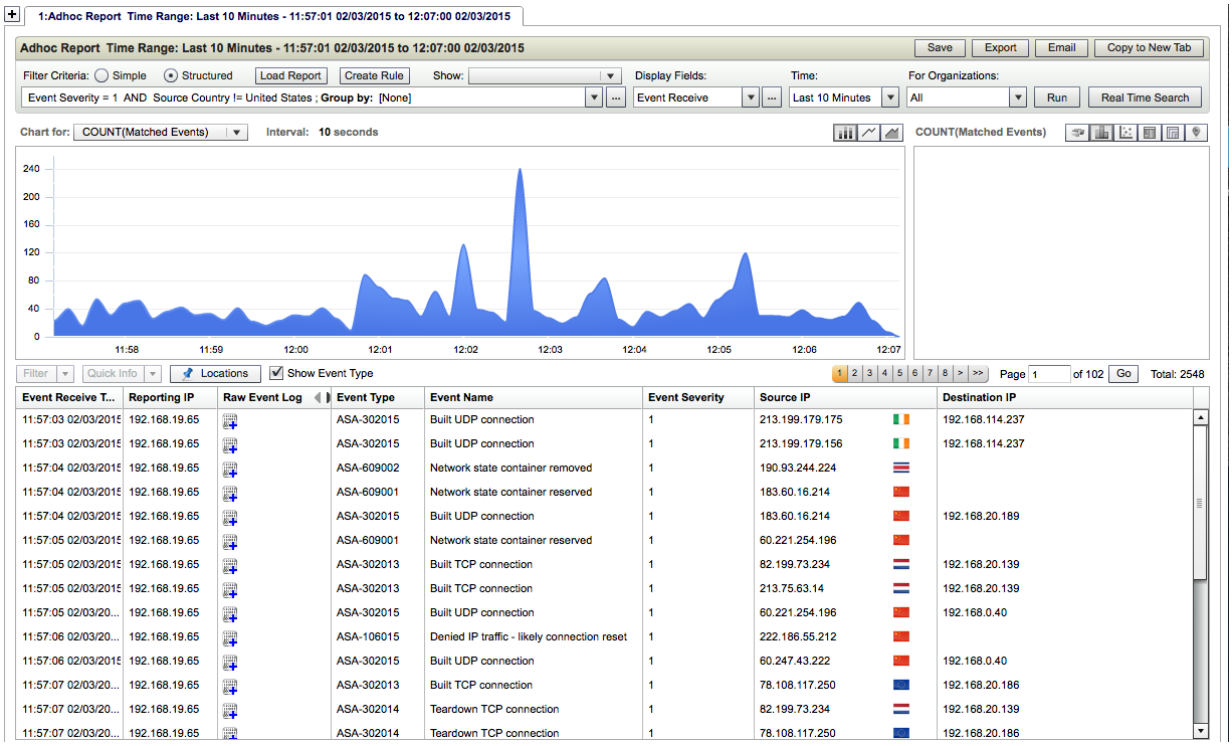

If you use a geolocation attribute such as Source Country as a **Display Field** or **Group By** condtion, then the results will include name information for that attribute, rather than a national flag.

This screenshot shows the results of the same query used previously, but with **Group By = Source Country.**

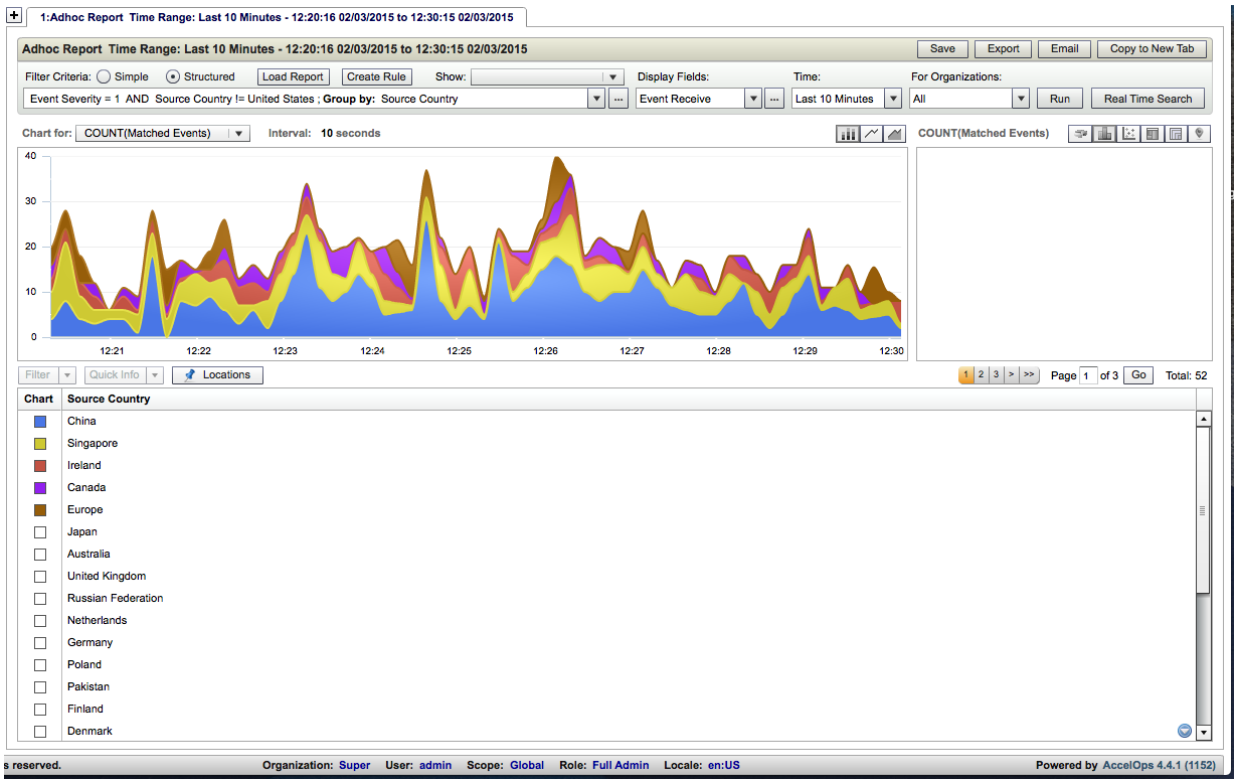

Viewing Geographic Locations from Search Results

If your search results contain geographic information, click the **Locations** button to view that information on a map.

This screenshot shows the results for the first example query presented in a map. Clicking on a number in the map will provide you with an overview of incidents for that location.

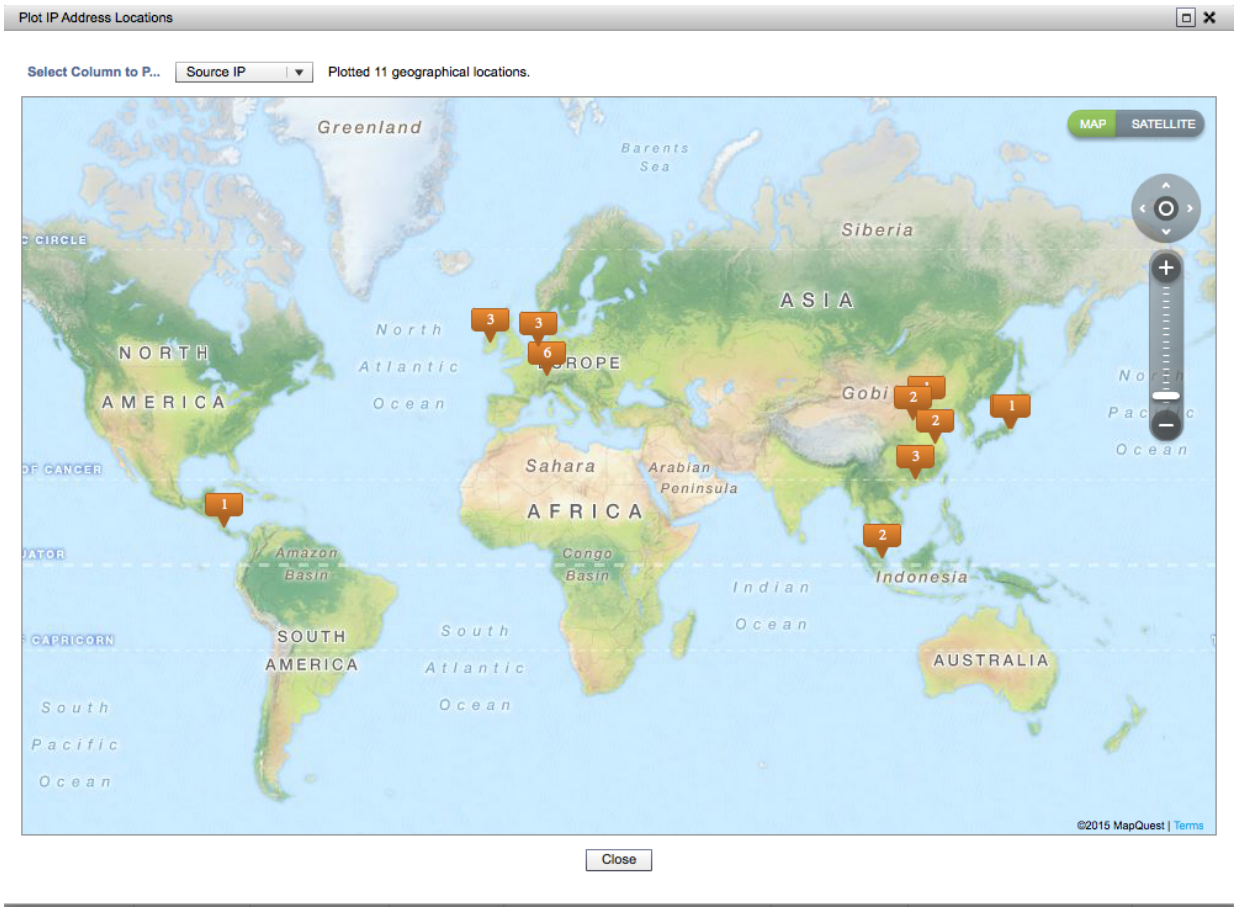
Plot IP Address Locations
Select Column to P...
Source IP
Plotted 11 geographical locations.
MAP
SATELLITE
Greenland
Barents Sea
Siberia
ASIA
NORTH AMERICA
North Atlantic Ocean
EUROPE
Gobi
Sahara
Arabian Peninsula
AFRICA
Amazon Basin
Congo Basin
Indian Ocean
Indonesia
SOUTH AMERICA
South Atlantic Ocean
AUSTRALIA
South Pacific Ocean
©2015 MapQuest | Terms
Close

Creating Filter Criteria and Display Column Sets

When you create searches, you have the option to select saved filter criteria and column sets to use. This topic describes how to create those sets.

1. Log in to your Supervisor node.
2. In the **Analytics** tab, select either **Display Column Sets** or **Filter Criteria Sets**, depending the type of set you want to create.
3. Click **New**.
4. Add the filter criteria or display columns that you want to the set.
   See Using Expressions in Structured Searches and Selecting Attributes for Structured Searches and Display Fields for more information about building searches and display columns.
5. Click **Save**.
   You set will be saved to the list of sets, and you will be able to use it in searches by clicking the **...** button next to the **Filter Criteria** text field in structured searches or the **Display Columns** menu for both structured and simple searches.

Saving Sets from Searches

Whenever you create a set of filter criteria for a structured search, or a set of display columns for both simple and structured search, you can save it by clicking the **Save as Filter Criteria Set** or **Save As Display Column Set** button.

Related Links

- Using Expressions in Structured Searches and Rules
- Selecting Attributes for Structured Searches, Display Fields, and Rules

Rules

FortiSIEM continuously monitors your IT infrastructure and provides you with information you can use to analyze performance, availability, and security. There may also be situations in which you want to receive alerts when exceptional, suspicious, or potential failure conditions arise. You can accomplish this by using rules that define the conditions to watch out for, and which trigger an incident when those conditions arise. This incident will appear on the Incident Summary dashboard, and you can also configure a notification policy that will send email and SNMP alerts that the incident has occurred. FortiSIEM includes over 500 system-defined rules, which you can see in **Analytics > Rules**, but you can also create your own rules as described in the topics in this section.

- Creating Rules
- Activating and Deactivating Rules
- Adding a Watch List to a Rule
- Cloning a Rule
- Running Historical Searches to Test Rule Sub Patterns
- Setting Rules for Event Dropping
- Setting Rules for Event Forwarding
- Setting Global and Per-Device Threshold Properties
- Using Geolocation Attributes in Rules
- Using Watch Lists as Conditions in Rules and Reports
- Viewing Rules

Creating Rules

FortiSIEM constantly monitors your IT infrastructure for events and collects information about them, but you can also set rules that will trigger incidents from events and send notifications when they occur. These topics describe the concepts and processes for creating rules.

- Creating a Rule
- Defining Rule Conditions
- Defining the Incident Generated by a Rule
- Defining Rule Exceptions
- Defining Clear Conditions
- Testing a Rule

Creating a Rule

Creating a new rule involves defining the attributes of the incident that is triggered by the rule, as well as the triggering conditions and any exceptions or clear conditions.

Creating New Rules from Clones

You can also create a rule by cloning an existing rule and editing it.

Restriction on names

Do not use certain keywords in subpattern names

- regexp

1. Log in to your Supervisor node.
2. Go to **Analytics > Rules**.
3. Select the group where you want to add the new rule.
4. Click **New**.
5. Enter a **Rule Name** and **Description**.
6. For **Status**, keep the rule **Inactive**.
   You can activate the rule after you're finished creating and testing it.
7. Select an **Incident Category** for the incident triggered by the rule.
   You can click **Add** and enter a custom incident category.
8. Select a **Severity** to associate with the incident triggered by the rule.
9. Select **Update the Perf Status column on summary dashboard** if you want the incident to display in the **Performance Status** column of the **Exec Summary** dashboard.
10. For **Attributes**, enter the functional area, such as **Security**, that you want to associate the rule with.
11. Enter a **Notification Frequency** for how often you want notifications to be sent when an incident is triggered by this rule.
12. Under **Conditions**, click **Add Subpattern** to create the rule conditions.
    See Defining Rule Conditions for detailed information on selecting event and aggregation attributes to use with rules. You can also see examples of rules with a single subpattern and multiple sub patterns.
13. Enter the time interval during which the rule conditions will apply.
    The minimal interval is **120** seconds.
14. Next to **Actions**, click **Edit** to define the incident that will be generated by this rule.
    See Defining the Incident Generated by a Rule for more information.
    One Incident Definition Required to Save: You must have at least one incident defined before you can save your rule.
15. Next to **Watch Lists**, click **Edit** to add a watch list to the rule.
    See Adding a Watch List to a Rule for more information.
16. If you want to define any **Exceptions** for the rule, click **Edit**.
    See Defining Rule Exceptions for more information.
17. If you want to define any **Clear Conditions** for the rule, click **Edit**.
    See Defining Clear Conditions for more information.
18. Click **Save**.
    Your new rule will be saved to the group you selected in an inactive state. Before you activate the rule, you should test it.

Defining Rule Conditions

Rule conditions define the event attributes and thresholds that will trigger an incident. Rule conditions are built from sub-patterns of event attribute filters and aggregation functions. You can specify more than one subpattern and the relationships and constraints between them.

- Specifying a Subpattern
- Setting the Relationship between Subpatterns
- Setting Inter-subpattern Constraints

Specifying a Subpattern

A subpattern defines the characteristics of events that will cause a rule to trigger an incident. A subpattern involves defining event attributes that will be monitored, and then defining the threshold values for aggregations of event attributes that will trigger an incident.

- Example of a rule with a single subpatten

This screenshot shows an example of a subpattern with a single event filter and a single event aggregation condition. Expressed as a sentence, this rule would be "When there are more than three events on a single Host IP where average CPU utlization is equal to 95%, trigger an incident."

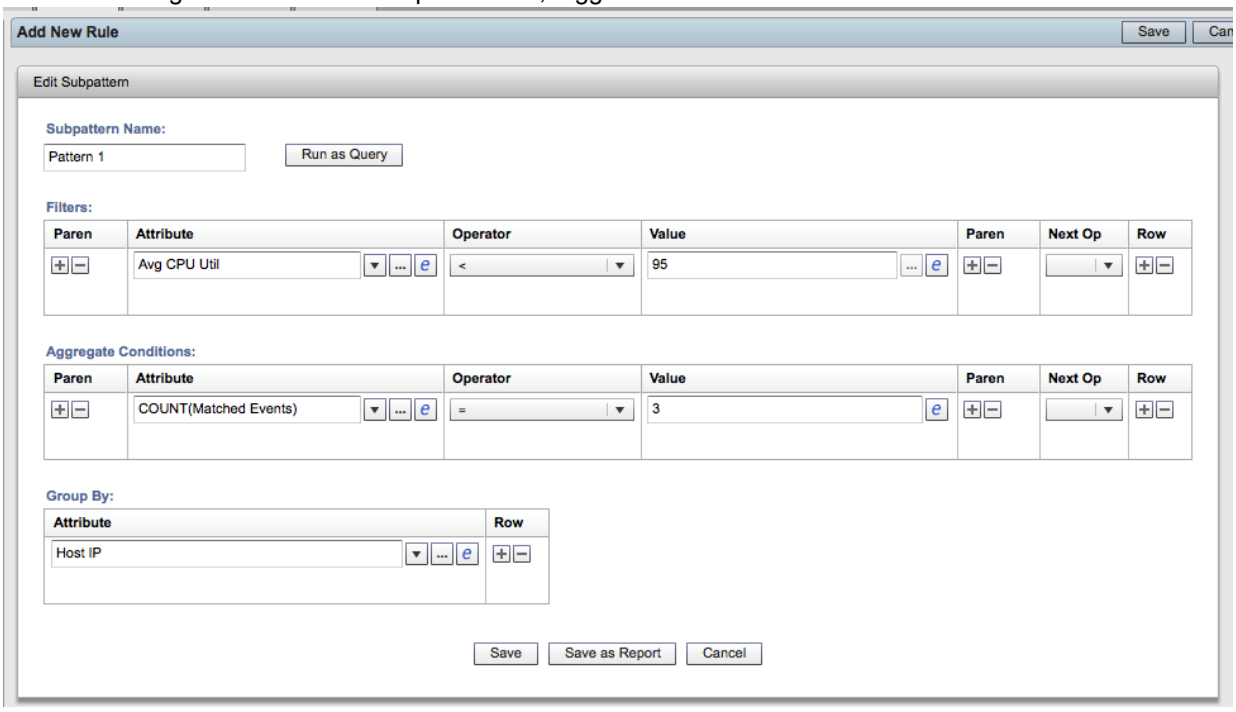

Event Filters

Event filter criteria determine which event attributes and values will be monitored by the rule, and are set in a way that is similar to the way you set event attributes for structured historical searches and real time searches. See Selecting Attributes for Structured Searches, Display Fields, and Rules for more information on finding attributes to use in your event filters.

Event Aggregation

While you could have a rule that triggers an incident on a single instance of a particular event, it is more likely that you will want your rule to trigger an incident when some number of events have been found that meet your event filter criteria.

Group By Attributes

This determines which event attributes will be used to group the events before the group constraints are applied, in a way that is similar to the way the Group By attribute is used to aggregate the results of structured searches.

Aggregate Conditions

The group aggregation conditions set the threshold at which some aggregation of events will trigger a rule to create an incident. You create an aggregation condition by using the **Expression Builder** to set a function, and then enter the **Operator** and **Value** for the aggregation condition.

Examples of Group By and Aggregate Conditions Settings

Scenario	Group By Attributes	Aggregate Conditions
10 or more events	none	**COUNT(Matched events) >= 10**
Connections to 100 or more distinct destination IPs from the same source IP	**Source IP**	**COUNT (DISTINCT Destination IP) >= 100**
Connections to 100 or more distinct destination IPs from the same source IP on the same destination port	**Source IP,Destination Port**	**COUNT (DISTINCT destination IP) >= 100**
Average CPU Utilization on the same server > 95% over 3 samples	**Host IP**	**COUNT (Matched Events) >= 3 AND AVG (CPU Util) > 95**
Logins from the same source workstation to 5 or more accounts on the same target server	**Source IP**, **Destination IP**	**COUNT(DISTINCT user) >= 5**

Setting the Relationship between Subpatterns

- Example of a rule with multiple subpatterns

If you have more than one sub-pattern, you must specify the relationship between them with these operators.

Operator	Meaning
AND	**Sub-pattern P1 AND Sub-pattern P2** means both sub-patterns P1 and P2 have to occur
OR	**Sub-pattern P1 OR Sub-pattern P2** means either P1 or P2 have to occur
FOLLOWED-BY	**Sub-pattern P1 FOLLOWED-BY Sub-pattern P2** means P1 has to be followed by P2 in time
AND-NOT	**Sub-pattern P1 AND-NOT Sub-pattern P2** means P1 must occur while P2 must not; the time order between P1 and P2 is not important
NOT-FOLLOWED-BY	**Sub-pattern P1 NOT-FOLLOWED-BY P2** means P1 must occur and P2 must not occur after P1

Setting Inter-subpattern Constraints

You may want to relate attributes of a sub-pattern to the corresponding attributes of another sub-pattern, in a way that is similar to a JOIN operation in an SQL, by using the relationship operators **<**, **>**, **<=**, **>=**, **=**, **!=**.

Examples of inter-subpattern relationships and constraints

Scenario	Sub-pattern P1 - filter	P1 - Group-by attribute set	P1 Group constraint	Sub-pattern P2 filter	P2- group-by attribute	P2 group constraint	Inter-P1-P2 relationships	Inter-P1-P2 constraints
5 login failures from the same source to a server not followed by a successful logon from the same source to the same server	Event type = Login Success	Source IP, Destination IP	COUNT (Matched Event) >= 5	Event type = Login failure	Source IP, Destination IP	COUNT (Matched Event) > 0	P1 NOT_ FOLLOWED_BY P2	P1's Source IP = P2's Source IP
An security attack to a server followed by the server scanning the network, that is, attempting to communicate to 100 distinct destination IP addresses in 5 minute time windows	Event type = Attack	Destination IP	COUNT (Matched Event) > 0	Event Type = Connection Attempted	Source IP	COUNT (DISTINCT Destination IP) > 100	P1 FOLLOWED_BY P2	P1's Destination IP = P2's Source IP
Average CPU > 95% over 3 sample on a server AND Ping loss > 75%	Event Type = CPU_ Stat	Host IP	COUNT (Matched Event) >= 3 AND AVG (cpuUtil) > 95	Event Type = PING Stat	Host IP	pingLossPct > 75	P1 AND P2	P1's Host IP = P2's Host IP

Example of a Rule with a Single Condition Sub-Pattern

This topic shows an example of how to create a rule with a single sub-pattern based on the condition that Average CPU on a server is more than 95% over 3 sample measurements.

Attribute	Group By Attribute	Aggregate Conditions
Avg CPU Util	**Host IP**	**COUNT (Matched Event) >= 3**

1. For **Rule Name**, enter **Hi Avg CPU**.
2. For **Description** enter **Average CPU on a server is more than 95% over 3 sample measurements**.
3. For **Severity**, select **9 - High**.
4. For **Attributes**, select **All**.
5. Set the **Notification Frequency** for **1 Hour**.
6. Next to **Conditions**, click **AddSubpattern**.
7. For **Subpattern Name**, enter **Pattern 1**.
8. Under **Filters**, set these options:

Option	Setting
Attribute	Avg CPU Util
Operator	>=
Value	95

9. Under **Aggregate Conditions**, click the **Expression Builder** icon next to the **Attribute** field, select **COUNT (Matched Events)** from the **Add Function** menu, and then click **OK**.
10. Under **Aggregate Conditions**, select **=** for **Operator** and enter **3** for **Value**.
11. Under **Group By**, select **Host IP**.
12. Click **Save**.
13. Enter **5** for the time interval during which the conditions will apply.
14. You would now complete the rule by Defining the Incident Generated by a Rule, and any exceptions or clear conditions. You could also associate it with a notification policy.

This screenshot shows the subpattern settings for this example.

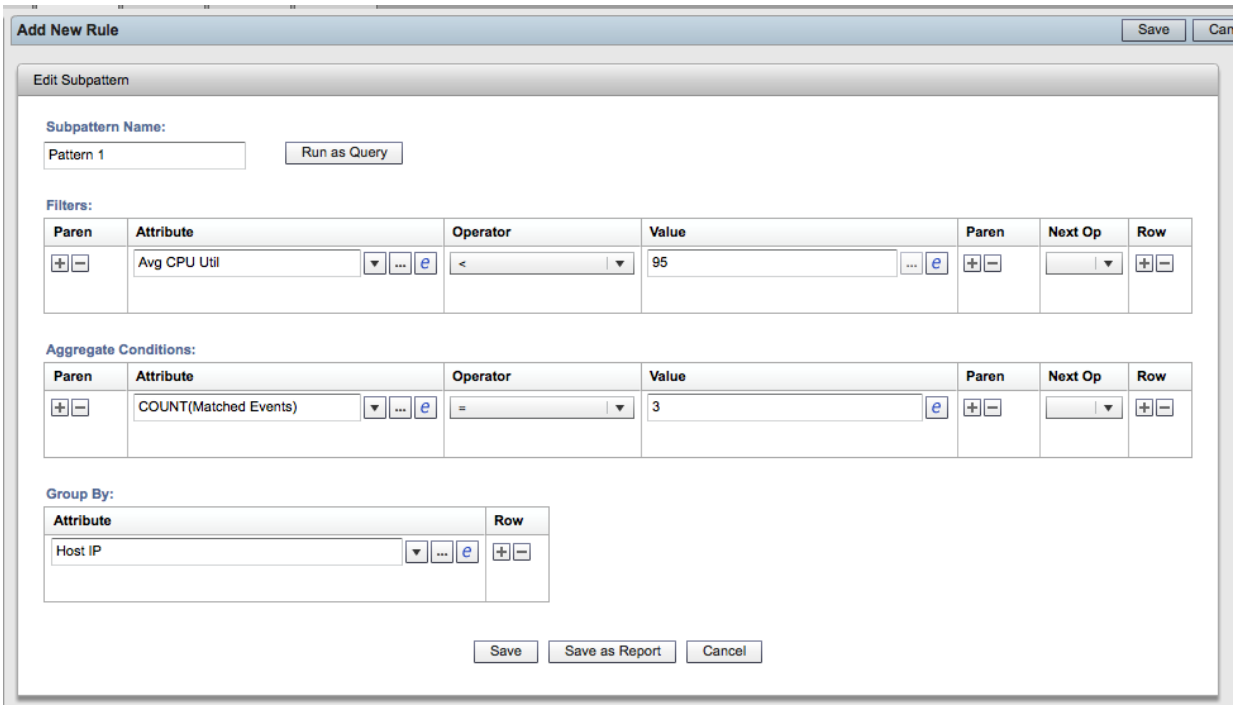

The following steps describe how to create a rule that matches the above example 1:

1. Enter a name for the rule in the 'Rule Name' text box.
2. Enter a description for the rule in the 'Description' text box.
3. Use the drop down menu to choose a 'Severity' for the rule.
4. Click on the '+ Add Condition' button.
    a. Chose the 'Function' for the rule. In this case 'AVG' is chosen.
    b. Choose the 'Attribute' for the rule. In this case 'CPU Util' is chosen.
    c. Chose the 'Operator' for the rule. In this case '>=' is chosen.
    d. Enter the 'Value' for the rule. In this case '95' is entered.
5. Select the devices to apply the rule to.
6. Enter the number of events that must occur for the rule to fire. In this case '3' is used.
7. Enter the time frame for the rule. In this case '600' seconds is used.

Example of a Rule with Multiple Sub-Patterns

This topic provides an example of a rule with two sub-patterns, and also how to use the **Event Type** attribute as a filter.

Rule Conditions

The purpose of this rule is to trigger an incident when five login failures from the same source to a server are not followed by a successful login from the same source to the same server within one hour. This requires two sub-patterns, the first one to detect "five login failures from the same source to a server," and a second one to detect "a successful logon from the same source to the same server." The two sub-patterns need to be interrelated to make the complete rule.

Sub-pattern 1 (P1)

Event Filter Attribute	Group By Attributes	Aggregate Conditions
Event type = **Logon Failure**	Source IP, Destination IP	COUNT (Matched Event) >= 5

Sub-pattern 2 (P2)

Event Filter Attribute	Group By Attributes	Aggregate Conditions
Event type = **Logon Success**	Source IP, Destination IP	COUNT(Matched Event) > 0

P1/P2 Interrelationships and Constraints

Interrelationships	Constraints
P1 NOT_ FOLLOWED_BY P2	P1's Source IP = P2's Source IP, P1's Destination IP = P2's Destination IP

Creating Sub-Pattern P1

The following steps describe how to create a rule that matches the above example 2:

1. Log in to your Supervisor node.
2. Go to **Analytics > Rules**.
3. Click **New**.
4. For **Rule Name**, enter **Suspicious Login Failure**.
5. For **Description**, enter the rule conditions stated in the introduction to this topic.
6. For **Severity**, select **10 - High**.
7. For **Attributes**, select **All**.
8. Next to **Conditions**, click **Add Subpattern**. You will now create the first subpattern for "five login failures from the same source to a server.".
9. For **Subpattern Name**, enter **LogonFailures**.
   To create this sub pattern you will want to specify that all types of logon failures should be monitored. For this reason, you will want to specify an entire folder of event types as the rule condition, rather than a single attribute of a event.
10. For **Attribute**, select **Event Type**.
11. For **Operator**, select **IN**.
12. For **Value**, click **...** to open the **CMDB Browser**.
13. In the CMDB Browser, go to **Event Types > Security > Logon Failure**, and click **Folder >>** to select the **Logon Failure** events group.
    Your filter condition, as shown in the screenshot, can be read as "For any type of event in the Logon Failure event group . . ."
14. Under **Aggregate Conditions**, click the Expression Builder icon next to **Attribute** and select **COUNT(Matched Events)**.
15. For **Operator**, enter **>=**.
16. For **Value**, enter **5**.
17. Under **Group By**, enter **Source IP** for **Attribute**, and then click **+** to add another **Group By** attribute.
18. Enter **Destination IP**.
19. Click **Save**.

This screenshot shows the complete entry for sub-pattern P1.

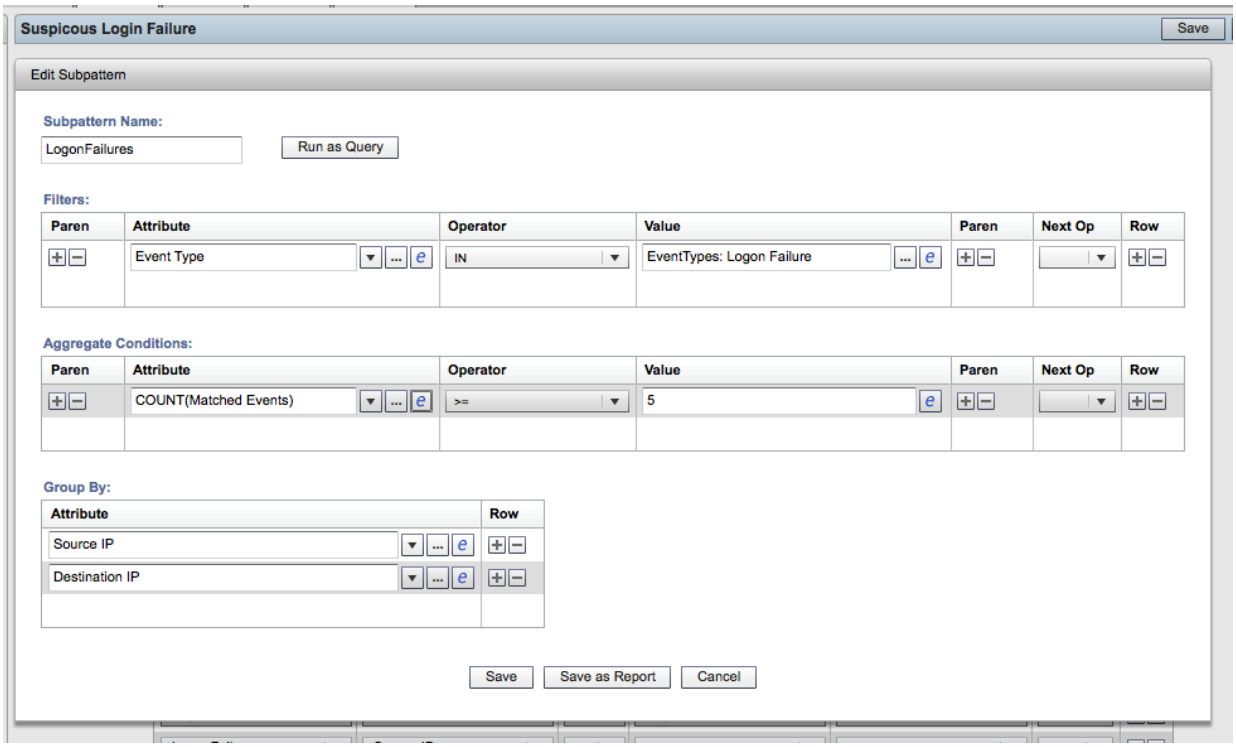

Creating Sub-Pattern P2

1. In your rule, next to **Conditions**, click **Add Subpattern**.
2. For **Subpattern Name**, enter **LogonSuccess**.
3. For **Attribute**, select **Event Type**.
4. For **Operator**, select **IN**.
5. For **Value**, click **...** to open the **CMDB Browser**.
   This button only becomes active if you select **Event Type** as an attribute.
6. In the CMDB Browser, go to **Event Types > Security > Logon Failure**, and click **Folder >>** to select the **Logon Failure** events group.
   Your filter condition, as shown in the screenshot, can be read as "For any type of event in the Logon Failure event group . . ."
7. Under **Aggregate Conditions**, click the Expression Builder icon next to **Attribute** and select **COUNT(Matched Events)**.
8. For **Operator**, enter **>**.
9. For **Value**, enter **0**.
10. Under **Group By**, enter **Source IP** for **Attribute**, and then click **+** to add another **Group By** attribute.
11. Enter **Destination IP**.
12. Click **Save**.

This screenshot shows the complete entry for sub-pattern P2.

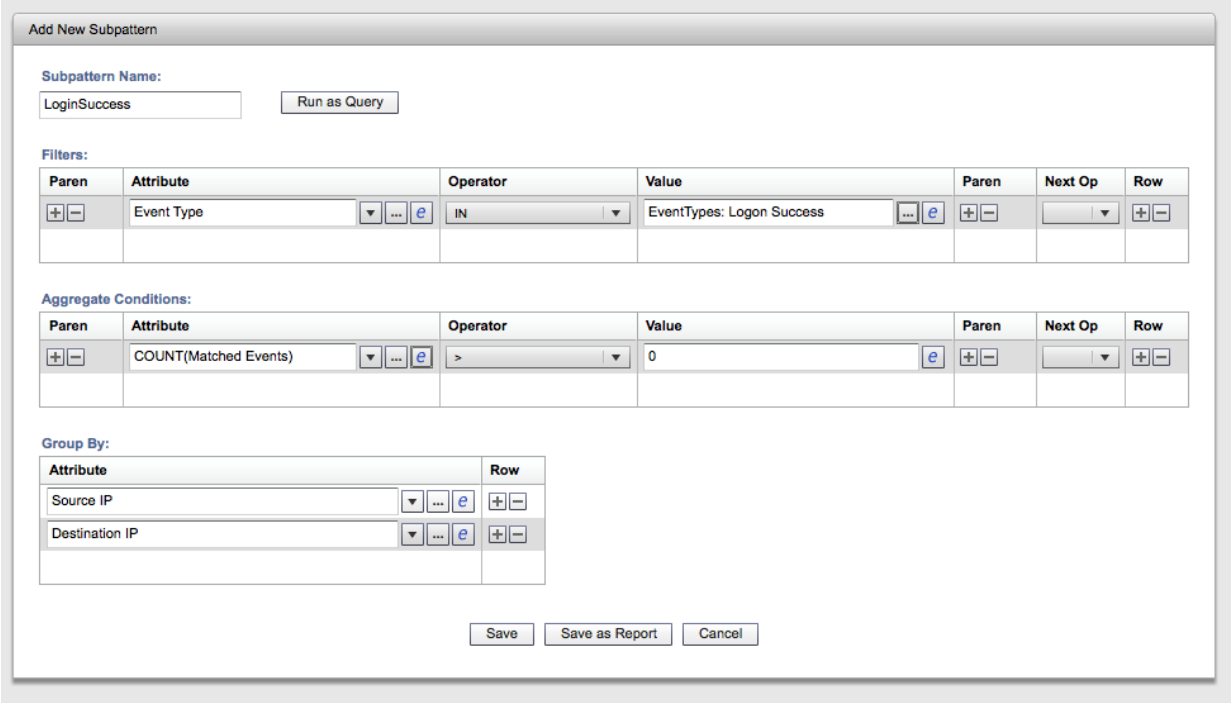

Defining the Relationship Between Patterns

You will now see both of your sub-patterns listed under the **Conditions** for your rule definition.

1. Makes sure that **LogonFailures** is selected as the first pattern under **If this Pattern occurs**, and under **Next Op**, select **NOT_FOLLOWED_BY**.
2. Select **LoginSuccess** as the second subpattern.
3. Click **AddSubpattern Relationship**.
4. For the first relationship definition, select **LogonFailures** for **Subpattern**, **Source IP** for **Attribute**, and **=** for **Operator**.
5. For the second subpattern, select **LogonSuccess** for **Subpattern**, **Source IP** for **Attribute**, and **AND** for **Next Op.**
6. Under **Row**, click **+**.
7. For the second relationship definition, for the first subpattern, select **LogonFailure** for **Subpattern**, **Destination IP** for **Attribute**, and **=** for **Operator**.
8. For the second subpattern, select **LogonSuccess** for **Subpattern**, and **Destination IP** for **Attribute.**

Defining the Incident to be Generated by the Rule

1. In your rule definition, click **Edit** next to **Generate Incident**.
2. For **Incident Name**, enter **Suspicious_Login_Failure**.
3. Under **Incident Attributes**, select **Source IP** for **Event Attribute**, **LoginFailures** for **Subpattern**, and **Source IP** for **Filter Attribute**.
4. Under **Row**, click **+**.
5. For the second incident attribute, select **Destination IP** for **Event Attribute**, **LoginFailures** for **Subpattern**, and **Destination IP** for **Filter Attribute**.

6. Under **Triggered Event Attributes**, make sure that **Event Receive Time**, **Event Type**, **Reporting IP**, and **Raw Event Log** are listed in the **Selected Attributes**.
7. Click **OK**.

Defining the Incident Generated by a Rule

Defining an incident involves setting attributes for the incident based on the subpatterns you created as conditions for the rule, and then setting attributes for the incident that will be used in analytics and reports.

One Incident Definition Required to Save

You must have at least one incident defined before you can save your rule.

1. In the rule you want to define an incident for, click **Edit** next to **Actions: Generate Incident**.
2. Enter an **Incident Name, Display Name**, and **Description**.
3. Under **Incident Attributes**, you will define attributes for the incident based on the **Group By** and **Aggregate Conditions** attributes you set for your sub patterns. Typically you will set the Incident attributes to be the same as the Group by attributes in the subpattern.
    a. Select the **Event Attribute** you want to add to Incident.
    b. Select a **Subpattern**.
    c. This will populate values from the **Group By** attributes in the subpattern to the **Filter Attribute** menu.
    d. In the **Filter** menu, select the attribute you want to set as equivalent to the **Event Attribute**.

Incident Definition Settings for the Single Subpattern Example

In the single sub pattern example, **Pattern1** has the **Group By** attribute set to **Host IP**, and the **Aggregate Conditions** attribute set to **COUNT(Matched Events)**. You can then select these to set as the incident attributes as shown in this screenshot.

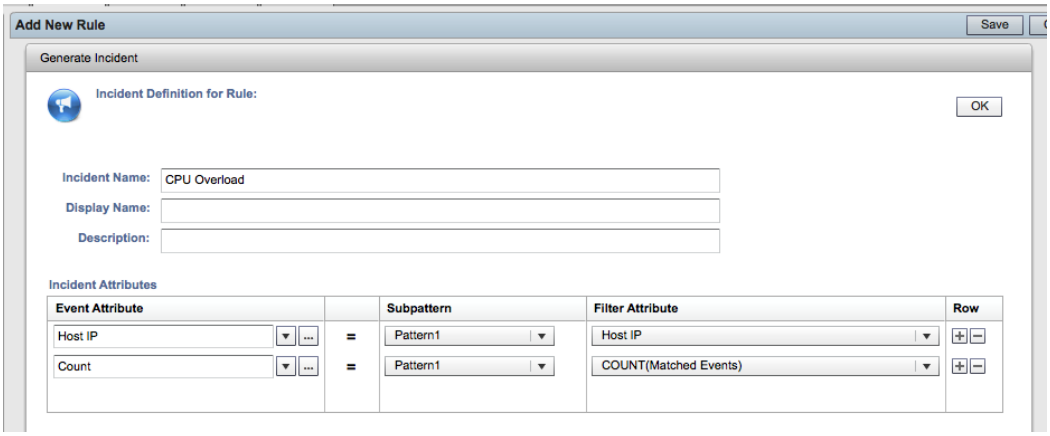

4. Under **Triggered Event Attributes**, select the attributes from the triggering events that you want to include in dashboards and analytics for this event.
   This is pre-populated with typical attributes you would want included in an incident report.
5. Click **OK**.

Defining Rule Exceptions

Once you activate a rule, it continuously monitors your IT infrastructure for conditions that would trigger an event. However, you may also want to define exceptions to those conditions. For example, you may know that a server will be going down for maintenance during a specific time period and you don't want your **Server Down - No Ping Response** rule to trigger an incident for it.

1. In **Analytics > Rules**, select the rule you want to add the exception to, and click **Edit**.
2. Next to **Exceptions**, click **Edit**.
3. Select an **Attribute** and **Operator**, and enter a **Value**, for the conditions that will prevent an incident from being generated.
   The values in the Attribute menu are from the **Event Attributes** associated with the incident definition.
4. Click the + icon to set an effective time period for the exception.
   You can set effective time periods for single and recurring events, and for durations of time from hours to days.
5. Enter any **Notes** about the exception.
6. Click **OK**.

Exception Condition for the Single Subpattern Example

This screenshot shows the exception conditions for the single sub-pattern example, where a specific server is set as an exception to the Host IPs that will trigger incidents during the maintenance period from March 9 to March 16 2015, starting at 14:00 Pacific Time for every day during that period, and lasting for 8 hours each day. The two **Attribute** options are populated from the attributes associated with the incident definition for the example.

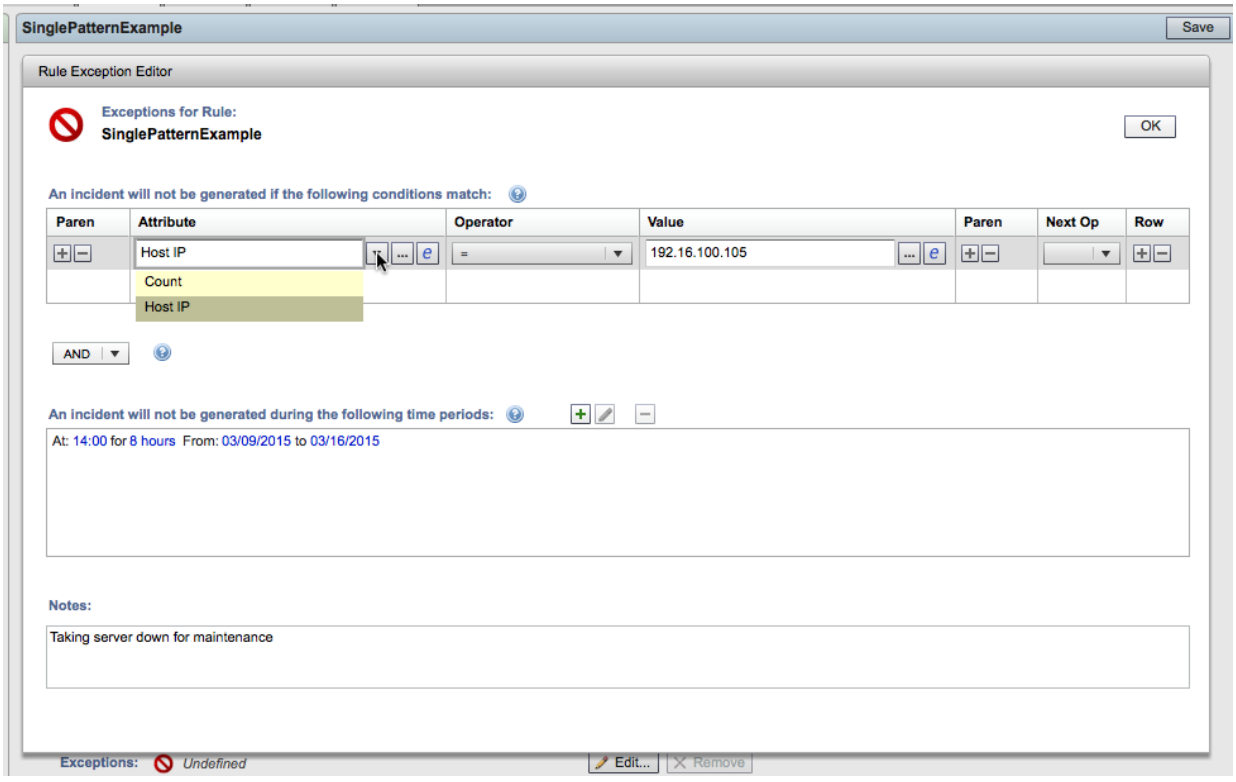

Defining Clear Conditions

Clear conditions specify conditions in which incidents will have their status changed from **Active** to **Cleared**. You can set the time period that must elapse for the clear condition to occur, and then set the conditions based on the triggering of the original rule, or on a sub pattern based on the Incident Attributes.

1. In **Analytics > Rules**, select the rule you want to add the clear condition to, and click **Edit**.
2. Next to **Clear Condition**, click **Edit**.
3. Set the **Time Period** that should elapse for the clear condition to go into effect.
4. If you want the clear condition to go into effect based on the firing of the original rule, select t**he Original Rule Does Not Trigger**.
   For example, if you wanted the clear condition to change the status of **Active** incidents to **Cleared** after the original rule had not been triggered for ten minutes, you would set **Cleared Within** to **10 Minutes** and select this option.
5. If you want to base the clear condition on a sub-pattern of the incident attributes, select **the following conditions are met**.
   The incident attributes from your rule will load and the clear condition attributes will be set to match.
6. Define the pattern to use by clicking the **Edit** icon next to the clear sub pattern.
7. Click **Save**.

Clear Condition for the Single Subpattern Example

This screenshot shows the exception condition settings for the example of a rule with a single subpattern. In the original rule, an incident was generated if there were three events over 10 seconds where **Avg CPU Util** exceeded 95% on a single host. In this example, those incidents will change status from from **Active** to **Cleared** if there are three events over 10 seconds where **Avg CPU Util** is under 100%.

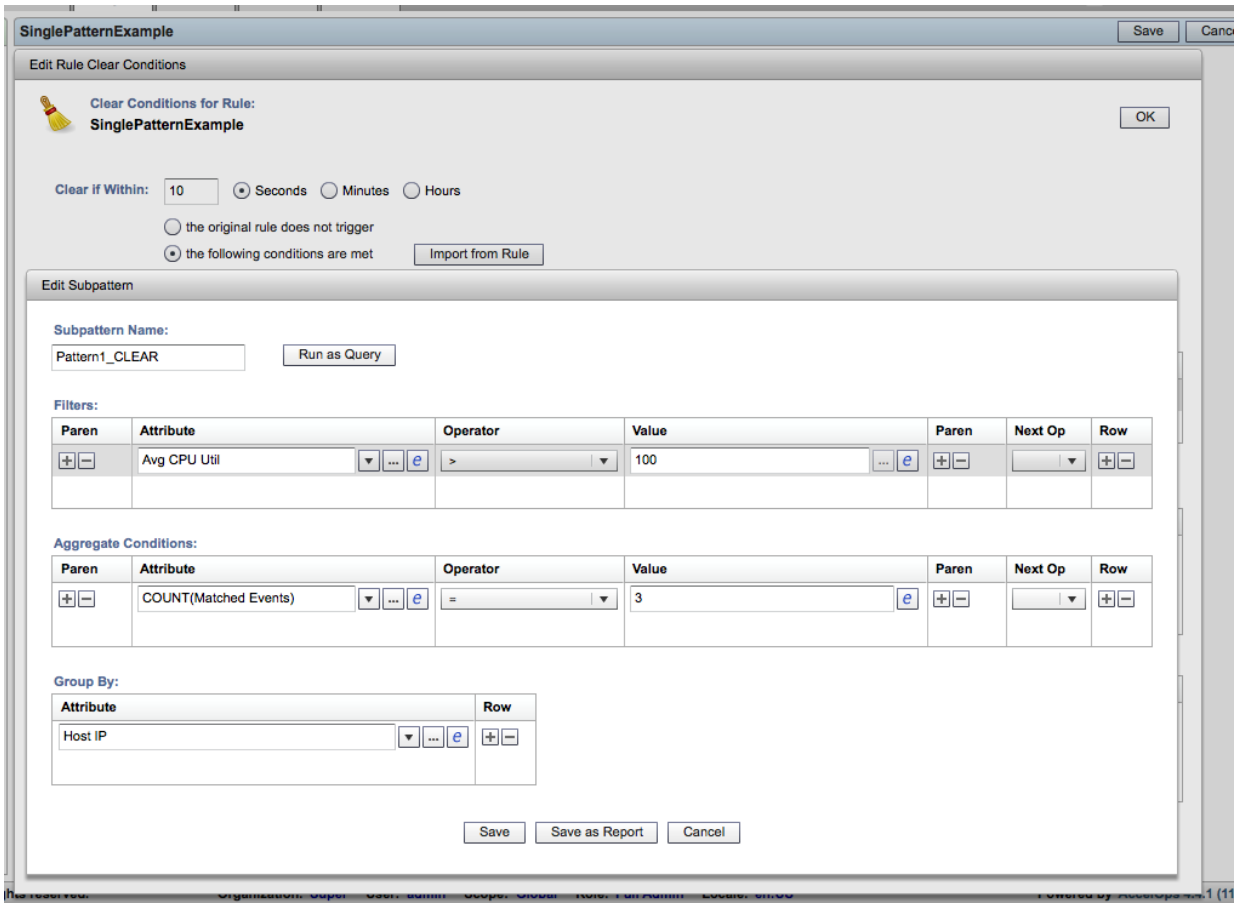
SinglePatternExample
Save
Edit Rule Clear Conditions
Clear Conditions for Rule:
SinglePatternExample
OK
Clear if Within: 10 Seconds Minutes Hours
the original rule does not trigger
the following conditions are met
Import from Rule
Edit Subpattern
Subpattern Name:
Pattern1_CLEAR
Run as Query
Filters:
Paren Attribute Operator Value Paren Next Op Row
Avg CPU Util > 100
Aggregate Conditions:
Paren Attribute Operator Value Paren Next Op Row
COUNT(Matched Events) = 3
Group By:
Attribute Row
Host IP
Save
Save as Report
Cancel

Testing a Rule

After you've created or a edited a rule, you should test it to see if behave as expected before you activate it. This topic describes how to test a rule using synthetic events.

Procedure

1. Go to **Analytics > Rules**, and deactivate the rule you want to test.
   Cloning Active Rules for Testing: You cannot test an active rule. If you can't deactivate a rule for testing, you can clone an inactive version of it.
2. Select the rule, and then click **Test Rule**.
   This will open the **Rule Debugger**.
3. Enter a **Reporting IP** where the synthetic event should originate from.
   Reporting IP Group Membership: If the rule you're testing specifies that the **Reporting IP** should be a member of a group, you should make sure that the Reporting IP you enter here is in that group.
4. Under **Raw Event**, enter the raw event log text that contains the triggering conditions for the rule.
5. Under **Pause**, enter the number of seconds before the next test event will be sent, and then click **+** under **Action** to enter additional test events.
   You will need to create as many events as are necessary to trigger the rule conditions.
6. Click **Run Test**.
   If the test succeeds you are now ready to activate the rule.

Test Results

The test will run through a four stage process, which you can observe in the **Test Results** tab of the rule. A yellow icon will also appear in the **Status** column for the rule to indicate that the test is running.

1. Rules are checked for syntax errors.
2. Events are parsed and sent to Rule Workers.
   If there are errors in the rule syntax or event parsing errors, see the examples under **Troubleshooting for Rule Testing** for suggestions on how to correct them. As events are being parsed, you can view their **Event Details** by clicking on the **Raw Event Log** icon next to the event.
3. Rule Worker nodes evaluate the events against the rule conditions, and if they match, they are sent to the Rule Master.
4. The Rule Master creates incidents, which then appear in the Incidents dashboards.

When the test successfully completes, a green icon will appear in the **Status** column next to the rule name.

Test Example

This screenshot shows the example of a test for the rule **Multiple Admin Login Failures: Net Devices**. The conditions for this rule are that the the **Reporting IP** must belong to a network device, and there must be 3 login failure events from the same IP and user.

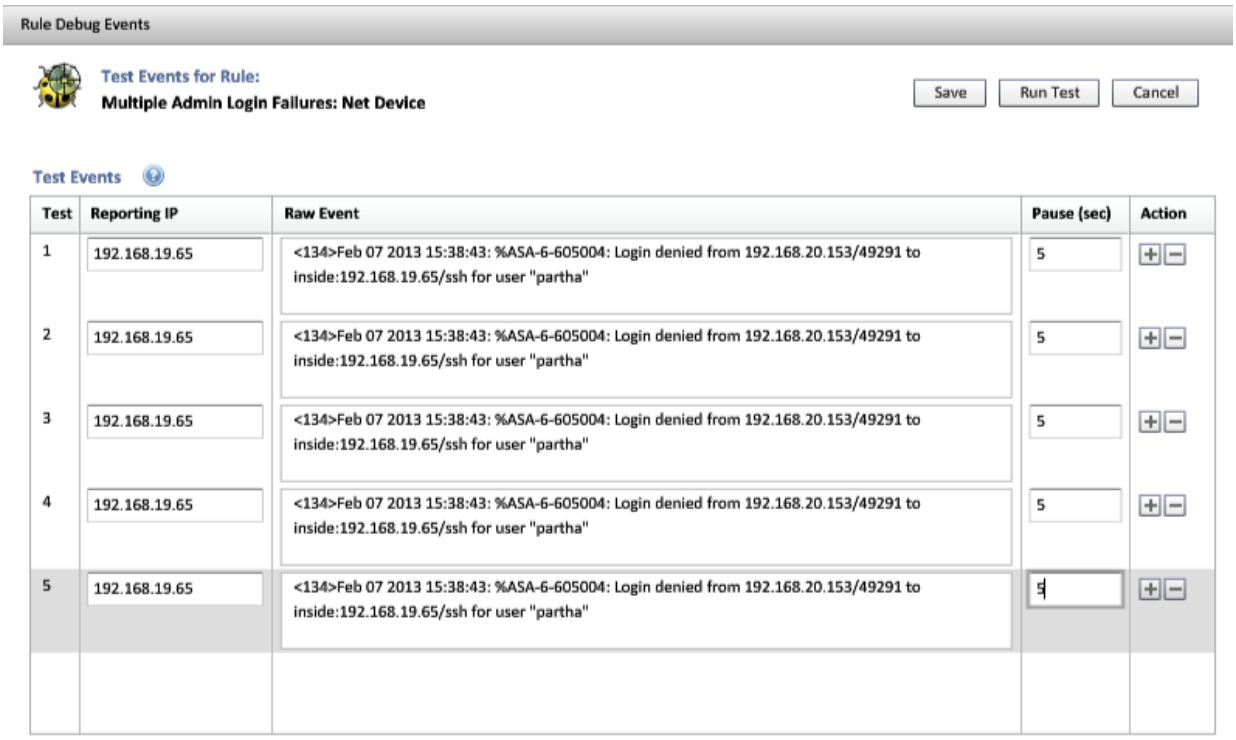

Troubleshooting for Rule Testing

If the test fails, a red icon will appear under the **Status** column next to the rule name, and you will see the error message in the **Test Results** tab for the rule.

Rule Syntax Error

The rule is changed to introduce the following error


```
Rule testing fails as shown below:
```

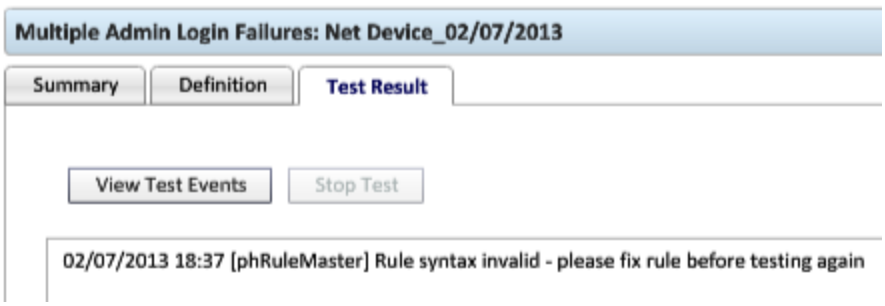

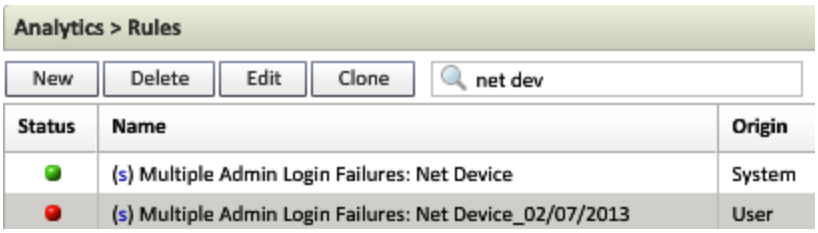

Rule Semantics Error

This means that the conditions of the rule were not met by the event. For example, if five events were required to meet the condition, but only one was sent.

Event Parsing Error

This means that some text in the raw event log did not pass the event parser. For example, if "denied" is the term expected by the parser in the test example, but the raw event log contains the term "deny," then the event will not pass the parser.

Activating and Deactivating Rules

When you create a new rule, you must activate it before it will start to monitor events. You may also want to deactivate a rule, for example to test it, instead of deleting it from the system.

1. Log in to your Supervisor node.
2. Go to **Analytics > Rules**.
3. Browse or search to find the rule that you want to activate or deactivate.
4. Select **Active** for the rule you want to activate, or clear the Active option if you want to deactivate a rule.

Activating Rules for Multi-Tenant Deployments

For Service Provider deployments you can activate or deactivate rules for individual organizations, and also set default rules for all organizations.

1. Navigate to the rule that you want to activate, deactivate, or set as default.
2. In the **Summary** tab of the rule, click **Edit**.
3. Next to **Status**, click **Edit**.
4. To set this rule as a default rule for all organizations, select **Activation Default**.
5. Select an organization to activate the rule for, or clear its selection to deactivate the rule.
6. Click **OK**.
7. Click **Save**.

Adding a Watch List to a Rule

1. Go to **Analytics > Rules**.
2. Select the rule you want to add the watch list to, and then click **Edit**.
3. Next to **Watch Lists**, click **Edit**.
4. Select the watch list you want to add, and use the **Add >>** button to add it to the rule.
5. For **Incident Attribute**, select the incident information you want to add to the watch list.
   Watch List Attribute Type Must Match Incident Attribute: The Type that you set for the watch list must match the Incident Attribute Types for the rule. For example, if your watch list Type is IP, and the Incident Attribute Type for the rule is string, you will not be able to associate the watch list to the rule.
6. Click **OK**.
   Next to **Watch Lists**, you will see **Watch List has been defined**.

Cloning a Rule

You can clone a rule to use it as the basis for creating another rule, or to use in testing.

1. Log in to your Supervisor node.
2. Go to **Analytics > Rules**.
3. Search or browse to select the rule you want to clone.
4. Click **Clone**.
5. Enter a new name for the cloned rule and click **OK**.
   The cloned rule will be added to the same group as the original rule but will be inactive.

Running Historical Searches to Test Rule Sub Patterns

If you are trying to analyze why a rule is triggering an excessive number of incidents, or why it isn't triggering any, you can run an historical search with the rule sub patterns to see how the sub pattern behaves in relation to past events. If the search has interesting results, you can then generate a report for further investigation. This is a way that you can test rules without having to deactivate them.

1. Go to **Analytics > Rules**.
2. Select a rule and then click **Edit**.
3. Click **Edit** next to the sub pattern you want to use in the search.
4. Click **Run as Query**.
5. Enter information for the time period you want to search.
6. Click OK.
   An historical search will run based on the sub pattern filters, aggregate conditions, and group by conditions.

Using a Sub Pattern in a Report

If the search includes results that you want to share or investigate further, you can save the rule as a report.

1. In the sub pattern you want to save, click **Save as Report**.
   The report will be saved in **Analytics > Reports**, and will have the phrase **From Rule** in the report name.
2. Select the report and click **Run Now** to generate a report from the sub pattern.

Setting Rules for Event Dropping

Some devices and applications generate a significant number of logs, which may be very verbose, contain little valuable information, and consume storage resources. You can configure Event Dropping rules that will drop events just after they have been received by FortiSIEM, preventing these event logs from being collected and processed. Implementing these rules may require some thought to accurately set the event type, reporting device type, and event regular expression match, for example. However, dropped events do not count towards licensed Events per Second (EPS), and are not stored in the Event database. Dropped event also do not appear in reports, and do not trigger rules. You can also specify that events should be dropped but stored, so event information will be available for searches and reports, but will not trigger rules. And example of an event type that you might want to store but not have trigger any rules would be an IPS event that is a false positive.

Procedure

1. Log in to your Supervisor node.
   For multi-tenant deployments you should log in to the Super/Global account if you want to set a system-wide event dropping rule. If you want to set an event-dropping rule for a specific organization, either log in as an administrator for that organization, or or log in using the Super/Global Account and then select the organization to which the rule should apply when you are creating it.
2. Go to **Admin > General Settings > Event Handling**.
3. Under **Event Dropping Rule**, click **Add**.
4. Next to **Reporting Device**, click **Edit**, and use the CMDB Browser to find device group or individual device that you want to create the rule for.
5. Next to **Event Type**, click **Edit**, and use the Event Type Browser to find the group of event types, or a specific event type, that you want to create the rule for.
6. If the event type you select has an **Source IP** or **Destination IP** attribute, you can enter specific IP addresses to which the rule should apply.
7. For **Regex Filter**, enter any regular expressions you want to use to filter the log files.
   If any matches are made against your regular expression, then the event will be dropped.
8. For multi-tenant deployments, select the **Organization** to which the rule should apply.
9. Select the **Action** that should be taken when the event dropping rule is triggered.
10. Enter any **Description** for the rule.
11. Click **Save**.

Notes

- All matching rules are implemented by FortiSIEM, and inter-rule order is not important. If you create a duplicate of an event dropping rule, the first rule is in effect.
- If you leave a rule definition field blank, then that field is not evaluated. For example, leaving **Event Type** left blank is the same as selecting **All Event Types**.
- FortiSIEM drops the event at the first entry point. If your deployment uses Collectors, events are dropped by the Collectors. If your deployment doesn't use Collectors, then the event will be droppedby the Worker or Supervisor where the event is received.
- You can use the report System Event Processing Statistics to view the statistics for dropped events. When you run the report, select AVG(Policy Dropped Event Rate(/sec) as one of the dimensions for Chart For to see events that have been dropped to this policy.

Setting Rules for Event Forwarding

In systems management, many servers may need access to forward logs, traps and Netflows from network devices and servers, but it is often resource intensive for network devices and servers to forward logs, traps and netflows to multiple destinations. For example, most Cisco routers can forward Netflow to two locations at most. However, FortiSIEM can forward/relay specific logs, traps and Netflows to one or more destinations. If you want to send a log to multiple destinations, you can send it to FortiSIEM, which will use an event forwarding rule to send it to the desired locations.

1. Log in to your Supervisor node.
2. Go to **Admin > General Settings > Event Handling**.
3. Under **Event Forwarding Rule**, for multi-tenant deployments, select the organization for which the rule will apply.
4. Click **Add**.
5. For **Sender IP**, enter the IP address of the device that will be sending the logs.
6. For **Severity**, select an operator and enter a severity level that must match for the log to be forwarded.
7. Select the **Traffic Type** to which the rule should apply.
   The **Forward To > Port** field will be populated based on your selection here.
8. For **Forward to > IP**, enter the IP address to which the event should be forwarded.
9. Click **OK**.

Multiple Destinations from the Same Sender IP

If you want the same sender IP to forward events to multiple destinations, create a rule for each destination.

Duplicate Rules Create Duplicate Logs

FortiSIEM will implement all rules that you create and enable, so if you create a duplicate of an event forwarding rule, two copies of the same log will be sent to the destination IP.

Setting Global and Per-Device Threshold Properties

Overview

In many cases when you create a rule, you set values for device thresholds that should trigger an incident. The example of a rule with a single sub-pattern, for example, contains a condition where if the average CPU utilization of a server exceeds 95% over 3 samples, an incident should be triggered. This is an example of setting an **absolute value** for the threshold in the rule itself.

Instead of setting an absolute value for the threshold, you can define global threshold properties that you can use as functions within a rule, and also define these threshold properties on a per-device basis. The advantage of this approach is that if you want to change the threshold values in a rule, you can edit the threshold property, rather than having to edit the rule. This is accomplished by using the **DeviceToCMDBAttr** function to return the value set for that device in the rule.

This table illustrates the difference between using an absolute value, shown in the first column, and threshold property, shown in the second column, in the aggregation conditions for a rule. For the threshold property, the function takes the form of **DeviceToCMDBAttr(Host IP, Threshold Property)**, while it takes the form of **DeviceToCMDBAttr(Host IP, Component, Threshold)** for devices with components as shown in the second example.

Rule Name	Aggregate Condition based on Absolute Value	Aggregate Condition based on Threshold Property Value
Server CPU Critical	AVG(CPU Utilization) > 95	AVG(CPU Utilization) > DeviceToCMDBAttr (Host IP,Server CPU Util Critical Threshold)
Server Disk Space Critical	AVG(Disk Utilization) > 99	AVG(Disk Utilization) > DeviceToCMDBAttr(Host IP,Disk Name,Disk Space Util Critical Threshold)

In the first example, when the rule evaluates the function, the **Server CPU Critical** rule will return the value of **Server CPU Util Critical Threshold** for the host IP if that has been defined for the reporting device, otherwise the global threshold value will return. In the second example, i f the **Disk Space Util Critical Threshold** is defined for a (Host IP,Disk Name) tuple, then the function returns that value, otherwise the global threshold value returns. This is an example of a **Map** threshold, in which there is one threshold value for each component, and which apply only to disk and interface components.

Defining a Global Threshold Property

FortiSIEM includes over 30+ pre-defined global threshold properties that you can edit and use in rules, but you can also create custom threshold properties.

1. Go to **Admin > Device Support.**
2. Click the **Custom Properties** tab.

3. Click **Add**.
4. Enter a **Name** and **Display Name** for the new threshold property.
5. Enter the **Default Value** for the threshold.
6. Select the **Type** of threshold value.
   For most global threshold values you will select **Double**. For **Map** thresholds, which apply to disks and interfaces, select the **Item Type** for the threshold value, and then select the **Component Type** to which it applies.
7. Click **Save**.

Defining Per-Device Threshold Properties

1. Go to **CDMB > Devices**.
2. Select a device.
3. In the **Device Details** pane, click **Edit**.
4. Click the **Properties** tab.
5. For any of the threshold properties, enter a value.
   If you want to edit a **Map** property, click **Edit** next to the property name, and then enter the value. If that device does not have any components to which that property could apply, you will see an error message.
6. Click **OK**.

Using the DeviceToCMDBAttr Function in a Rule

Using the example of the Server CPU Critical rule, you would use the DeviceToCMDB function to set a threshold for the aggregation conditions of the rule in this way:

1. In the sub pattern of the rule, under **Aggregation Conditions**, click the expression builder icon next to the **Attribute** field.
2. In the expression builder, under **Add Function**, select **AVG**.
3. In the **Add Event Attribute** field, select **CPU Utilization**.
4. Click **OK**.
   The expression builder will close, and you will see the function and event attribute you selected listed as the Attribute for the Aggregate Conditions.
5. For **Operator**, select **=**.
6. Click the expression builder icon next to the **Value** field.
7. In the **Add Function** menu, select **DeviceToCMDBAttr**.
8. In the **Select Function Pattern** dialog, select **DeviceToCMDBAttr(EventAttr,CMDBAttr)**.
9. Under **Add Event Attribute**, select **Host IP**.
10. Under **Add CMDB Attribute**, select **Server CPU Util Critical Threshold**.
11. Click **OK**.
12. Click **Save**.

Using Geolocation Attributes in Rules

In the same way that you can use geolocation attributes in searches and search results, you can also use them in creating rules. FortiSIEM includes four system-level rules based on geolocation attributes:

- Failed VPN Logon from Outside My Country
- Successful VPN Logon from Outside My Country
- Large Inbound Transfer From Outside My Country
- Large Outbound Transfer To Outside My Country

This screenshot shows the sub pattern for **Failed VPN Logon from Outside My Country** as an illustration of the way you can use geolocation attributes in a rule.

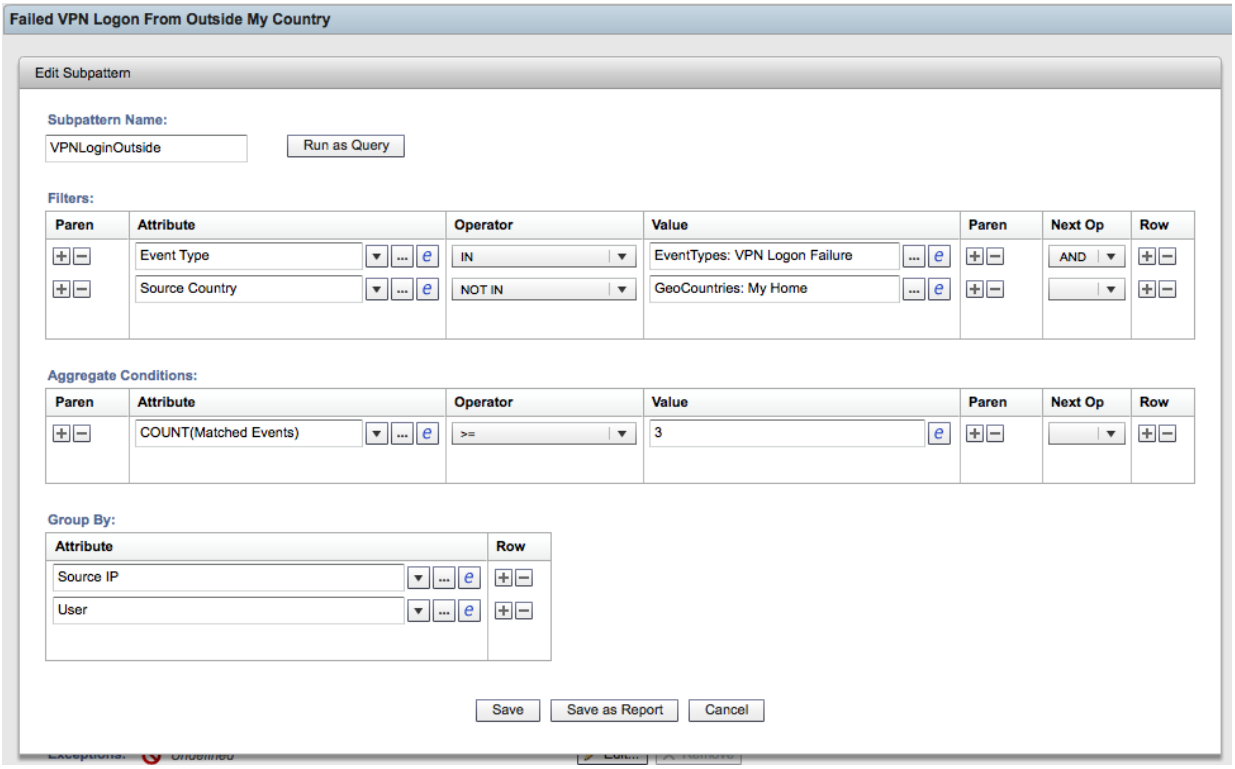

Using Watch Lists as Conditions in Rules and Reports

You may want to create a rule that refers to the attributes in a watch list, for example if you want to create a condition in which a **Source IP** listed in your **DNS Violators** watch list will trigger an incident.

1. Go to the rule or report where you want to use the watch list.
2. Under **Conditions** for the report, or under **Filters** in your rule subpattern, enter the watch list attribute you want to filter for in the **Attribute** field.
   For example, **Source IP**.
3. For **Operator**, select **IN**.
4. Click **...** next to **Value**, and use the CMDB Browser to find and select the watch list you want to use.
   For example, **DNS Violators**.
5. Click **Folder >>** to select the watch list, and then click OK.
6. Continue with creating your search criteria or rule sub pattern as you normally would.

Viewing Rules

FortiSIEM includes a large set of rules for Availability, Performance, Change, and Security incidents in addition to the rules that you can define for your system.

1. To view all system and user-defined rules, **go to Analytics > Rules**.
2. For multi-tenant deployments, use the **Organizations** menu in the upper-right corner of the Rules List pane to filter rules by organization.
3. Select any rule in the Rules List to view information about it.

All rules have three information tabs:

Tab	Description
Summary	This tab provides an overview of the rule's logic, its status, and its notification settings.
Definition	An XML definition of the rule. This is what will be copied to your clipboard if you **Export** a rule.
Test Results	If you are testing a rule, you can view the results here.

Reports

You can think of reports as saved or pre-defined versions of searches that you can load and run at any time. FortiSIEM includes over 2000 pre-defined reports that you can access in **Analytics > Reports**. Topics in this section describe how to access and view information about reports, how to create baseline reports, and how to use specialized reports like the **Identity and Location** report. You can refine the results of your reports in the same way that you would refine the results of an historical search or a real time search.

- Baseline Reports
- Creating a Report or Baseline Report
- Identity and Location Report
- Report Bundles
- Running System and User-Defined Reports and Baseline Reports
- Scheduling Reports
- Viewing Available Reports

Baseline Reports

- How FortiSIEM Sets Baselines
- Evaluating Rules and Detecting Deviations

When you are setting up FortiSIEM to monitor your IT infrastructure, you may want to define what is "normal" activity within your systems, and have incidents triggered when a a deviation from that normal activity occurs. For example, you can always assume that there will be some logon failures to a server on a daily basis. Rather than creating a rule that will trigger an incident when a certain hard-coded number of failures occurs, you can set up baseline reports that will trigger an incident when the total number of logon failures over a time period is twice the average over the same time period, or when the deviation from the average is threee times the standard deviation over a specific time period.

By creating a baseline report, you can set mean and standard deviations for any metric and use them in rule, and FortiSIEM will evaluate the current monitored values against the mean and standard deviation for that time period.

How FortiSIEM Sets Baselines

Establishing a baseline means recognizing that data center resource usage is time dependent:

- Usage is different during weekdays and weekends, and may also be different depending on the day of the week or month
- Usage is dramatically higher during business hours, typically 8am-5pm

FortiSIEM maintains distinct baselines for weekdays, weekends and for each hour of day - a total of 24*2 = 48 buckets. Baselines for days of the week or month are not maintained to save memory usage, as this would require 31*24 = 1764 buckets, a 15 fold-increase of memory.

A baseline report is a set of **Keys** that represent the baselined metrics, and a collection of **Values**. You can see examples of these Keys and Values in the System-Defined Baseline Reports. These are then used in this process to build the report:

1. During the current hour, the Supervisor and any Worker nodes operate in parallel to save a baseline report in memory by analyzing the report events as a stream.
2. When the hour finishes:
    1. The report is written to disk (on NFS for FortiSIEM cluster).
    2. The Supervisor module summarizes individual baseline reports from all nodes and forms the baseline for the current hour.
    3. The baselines are stored in a SQLite database on a local Supervisor.
    4. The Supervisor module reads the previous baseline for the current time interval from the SQLite database**.** Then it combines the previous values with the current values to create a new baseline.
    5. The new baseline is then stored in SQLite database.
3. For the new hour, a new baseline is created following this process

As this process illustrates, baselining is continuous in FortiSIEM, and new baseline values are learned adaptively.

Evaluating Rules and Detecting Deviations

A baseline rule contains expressions that involve using the functions **STAT_AVG()** and **STAT_STDDEV()** to set dynamic thresholds.

These examples show how **STAT_AVG()** and **STAT_STDDEV()** would be used to evaluate the conditions for the example of logon failures in the introduction to this topic.

Condition Statement	How the Baseline is Evaluated
Current value of X is more than 2 times the statistical average of X for the current hour	Baseline evaluated using Baseline Report with ID X > 2 * STAT_AVG (X:ID)
Deviation of X from its statistical average is more than 3 times its standard deviation for the current hour	All baselines evaluated using Baseline Report with ID ABS(X - STAT_AVG(X:ID) > 3 * STAT_ STDDEV(X:ID)

When FortiSIEM processes these rules:

1. Rule engine computes the current values in memory.
2. Every 5 minutes:
    1. It looks for **STAT_AVG(X:ID)** and **STAT_STDDEV(X:ID)** in memory
    2. If it fails, it retrieves them from the SQLite database and caches them for future use during the hour.
    3. Evaluates the rule conditions

A sample rule condition involving statistical functions is shown below with (X = AVG(fwConnCount); ID = 112).

```
<PatternClause window="1800"> <SubPattern id="3238092" name="StatHighConn">
<SingleEvtConstr>eventType = "PH_DEV_MON_FW_CONN_UTIL"</SingleEvtConstr>
<GroupEvtConstr> (AVG(fwConnCount)-STAT_AVG(AVG(fwConnCount):112))/STAT_
STDDEV(AVG(fwConnCount):112) >= 3 AND STAT_STDDEV(AVG(fwConnCount):112) > 0
</GroupEvtConstr> <GroupByAttr>hostName,hostIpAddr</GroupByAttr>
</SubPattern></PatternClause>
```

Setting Sample Points for Baselines

Two sample points are needed to avoid premature triggering of a rule before a baseline is set and becomes active.

- If the first data is received for a subject on Monday, then the rules will start triggering for that subject for that baseline starting Wednesday
- If the first data is received for a subject on Saturday, then the rules will start triggering for that subject for that baseline starting next Saturday

System-Defined Baseline Reports

- Network Traffic Analysis
- Performance / Availability Monitoring

- Logon Activity

Network Traffic Analysis

Report	Description	ID	Fields
DNS Request Profile	This report baselines DNS requests on a per client basis: the number of requests and distinct destinations it attempted to resolve	113	**Key**: Source IP **Values**: Number of Requests, Distinct Destination Count - means and standard deviation for each
DNS Traffic Profile	This report baselines DNS traffic characteristics on a per client basis: sent and receive bytes and packets.	113	**Key**: Source IP **Values**: Sent Bytes, Received Bytes, Total Bytes - mean and standard deviation for each
Destination Traffic Profile	This report baselines traffic destined to a server. The data is reported by network flow (Netflow, Sflow) and firewall logs. For each destination IP, the number of distinct peers, the number of distinct ports opened on the server and the total number of flows are tracked.	126	**Key**: Destination IP **Values**: Distinct Source IP, Distinct Destination Ports, Total Flows - mean and standard deviation for each
Source Traffic Profile	This report baselines traffic generated by a source. The data is reported by network flow (Netflow, Sflow) and firewall logs. For each source IP, the number of distinct peers, the number of distinct ports opened by the source, the total number of flows and total bytes exchanged are tracked.	125	**Key**: Source IP **Values**: Distinct Destination IP, Distinct Destination Ports, Total Flows, Total Bytes - mean and standard deviation for each
Firewall Connection Count Profile	This report provides baseline of permitted firewall connection count typically gathered by SNMP.	112	**Key**: Firewall Name, Firewall IP **Values**: Firewall Connection Count - mean and standard deviation for each
Firewall Denied Aggregate Traffic Profile	This profile baselines denied firewall traffic from firewall logs - volume of denied traffic, distinct attacker count, distinct target IP and port.	108	**Key**: Firewall Name, Firewall IP **Values**: Denied Flows, Distinct Denied Source IP, Distinct Denied Destination IP, Distinct Denied Destination Port - mean and standard deviation for each

Report	Description	ID	Fields
ICMP Traffic Profile	This report baselines generated ICMP traffic by each source: number of ICMP packets and number of distinct destinations	114	**Key**: Source IP **Values**: Distinct Destinations, Total Flows, Total Bytes - mean and standard deviation for each
Inbound Firewall DeniedTCP/UDP Port Profile	This report provides baseline of denied inbound TCP/UDP port usage as reported by firewall logs. For every port, the number of denied attempts and the number of distinct source are profiled.	106	**Key**: Destination Protocol, Port **Values**: Distinct Source IP, Total Flows - mean and standard deviation for each
Inbound Firewall PermittedTCP/UDP Port Usage Profile	This report provides baseline of permitted inbound TCP/UDP port usage. The data is reported by firewall logs. For every inbound destination port and protocol combination, the total number of unique sources, destinations and the total bytes and flows are profiled	104	**Key**: Destination Protocol, Port **Values**: Distinct Source IP, Distinct Destination IP, Total Flows, Total Bytes - mean and standard deviation for each
Outbound Firewall DeniedTCP/UDP Port Profile	This report provides baseline of denied outbound TCP/UDP port usage as reported by firewall logs. For every port, the number of denied attempts and the number of distinct destinations are profiled.	107	**Key**: Destination Protocol, Port **Values**: Distinct Destination IP, Total Flows - mean and standard deviation for each
Outbound Firewall PermittedTCP/UDP Port Usage Profile	This report provides baseline of permitted inbound TCP/UDP port usage. The data is reported by firewall logs. For every inbound destination port and protocol combination, the total number of unique sources, destinations and the total bytes and flows are profiled	105	**Key**: Destination Protocol, Port **Values**: Distinct Source IP, Distinct Destination IP, Total Flows, Total Bytes - mean and standard deviation for each

Performance / Availability Monitoring

Report	Description	ID	Fields
Device CPU, Memory Usage Profile	This report provides baselines cpu, memory usage - the data is collected by SNMP or WMI. For every host, CPU, real and virtual memory utilization are profiled	109	**Key**: Host Name **Values**: CPU Utilization, Memory Utilization, Virtual Memory Utilization - mean and standard deviation for each
Device Disk I/O Profile	This report provides baselines disk I/O usage for servers, VMs and ESX - the data is collected by SNMP or WMI or VCenter API. For every host and disk combination, read and write volumes are profiled	121	**Key**: Host Name, Datastore Name, Disk Name **Values**: Disk Read KBps, Disk Write KBps - mean and standard deviation for each
Network Interface Traffic Profile	This report provides baselines network interface traffic. The data is collected by SNMP. For each network interface, the total sent and received bytes are profiled.	110	**Key**: Host Name, Interface name **Values**: Sent Bytes, Received Bytes - mean and standard deviation for each
Network Interface Error Profile	This report provides baselines network interface errors and discards. The data is collected by SNMP. For each network interface, the total errors and discards are profiled.	111	**Key**: Host Name, Interface name **Values**: Errors, Discards - inbound and outbound - mean for each
Server Process Count profile	This report baselines the number of processes running at a server. The data is collected by SNMP.	123	**Key**: Host name **Values**: Process Count - mean and standard deviation
Reporting EPS Profile	This report baselines the rate at which devices sends events to FortiSIEM.	116	**Key**: Host Name, Host IP **Values**: Events/sec - mean and standard deviation
Reported Event Type Profile	This report provides baselines for distinct event types reported by a device.	119	**Key**: Host Name, Host IP **Values**: Distinct Event Type - mean and standard deviation
Reported Error Log Profile	This report baselines the number of system errors reported in logs on a per device basis.	120	**Key**: Host Name, Host IP **Values**: Number of events classified as system errors - mean
STM Response Time Profile	This report baselines Synthetic Transaction Monitoring response times	123	**Key**: Host Name, Monitor Name **Values**: Response Time - mean and standard deviation

Logon Activity

Report	Description	ID	Fields
Successful Logon Profile	This report baseline successful log on activity at a host. The data is collected from logs.	115	**Key**: Host Name, Host IP **Values**: Successful Logons, Distinct Source IP, Distinct Users - mean and standard deviation
Failed Logon Profile	This report baseline failed log on activity at a host. The data is collected from logs.		**Key**: Host Name, Host IP **Values**: Failed Logons, Distinct Source IP, Distinct Users - mean and standard deviation
Privileged Logon Profile	This report baseline successful log on activity at a host. The data is collected from logs.	118	**Key**: Host Name, Host IP **Values**: Privileged Logons - mean and standard deviation

Creating a Report or Baseline Report

Creating a report or baseline report is like creating a structured historical search, because you set the **Conditions** and **Group By** attributes that will be used to process the report data, and specify **Display Fields** to use in the report summary.

Cloning an Existing Rule: You can clone an existing rule to use as the basis for a new rule by selecting the existing rule, and then click Clone.

1. Log in to your Supervisor node.
2. Go to **Analytics > Reports**, and select the category for your new report.
   Select **Baseline** for baseline reports.
3. Click **New**.
4. Enter a report **Name** and **Description**.
5. For baseline reports, select **Anomaly Detection Baseline**.
6. Enter the **Conditions** to use in your report.
   See Selecting Attributes for Structured Searches, Display Fields, and Rules and Using Expressions in Structured Searches and Rules for more information on setting conditions. For creating baseline reports, see Baseline Reports for information on how to use the **STAT_AVG** and **STAT_STDDEV** functions in creating expressions for baseline reports.
7. Select the **Group By** attribute to use in processing the search results.
   The topic Example of How a Structured Historical Search is Processed explains how the Group By attribute is used in search results.
8. Set the **Display Fields** to use in your search results.
   See Selecting Attributes for Structured Searches, Display Fields, and Rules for more information on using event attributes in display fields.
9. Click **Save**.
   Your report will be saved into the selected category, and you can now run it or schedule it to run later.

Related Links

- Creating a Structured Historical Search
- Selecting Attributes for Structured Searches, Display Fields, and Rules
- Example of How a Structured Historical Search is Processed
- Using Expressions in Structured Searches and Rules
- Baseline Reports

Identity and Location Report

Overview

The Identity and Location report is constructed by associating a network identity like an IP address, or MAC address, to a user identity like a user name, computer name, or domain, and tying that to a location, like a wired switch port, a wireless LAN controller, or VPN gateway. When any element of these associations changes, a new entry is created in the report.

The associations between IP addresses, users, and locations are obtained by combining Windows Active Directory events, DHCP events, and WLAN and VPN logon events, with discovery results to produce a report combining all of this information into a comprehensive listing of users and machines by their identity and location.

The Identity and Location Report Display Fields

The Identity and Location Report contains these display fields:

Display Field	Description
IP Address	IP adress of a host whose identity and location is recorded in this result. You can view IP addresses with country flags in a map by clicking **Locations**.
MAC Address	MAC address of the host
User	User associated with this IP Address. Obtained from one of these event types: **Windows Domain Logon**, **WLAN Login**, **VPN Logon**, **AAA Authentication**. See the section on Report Information and Event Types on this topic for more information.
Host Name	Obtained from the **Windows Domain Logon** and **WLAN Authentication** event types.
Domain	Information displayed here depends on the logon event type it was obtained from: • **Windows Domain Logon**: the Domain name • **VPN Logon**: the reporting IP address of the VPN gateway • **WLAN Logon**: the reporting IP address of the WLAN controller • **AAA Logon**: the reporting IP of the AAA server
VLAN ID	For hosts directly attached to a switch, this is the VLAN ID of the switch port
Location	For h osts attached to a switch port, this is the switch name, reporting IP address, and interface name
First Seen	The time at which this entry was first created in the FortiSIEM Identity and Location table
Last Seen	The time at which some attribute of this entry was last updated. If there is a conflict, for example a host acquiring a new IP address because of DHCP, then the original entry is closed and a new entry is created. A closed entry will never be updated.

Report Information and Event Types

This table lists the events and event types that contribute to information in the Identity and Location Report, as well as what information is collected for each type of event.

	IP	MAC	Host Name	User	Domain	VLAN	Loca-tion	Contributing event type
DHCP Renew Events	x	x						• WIN-DHCP-IP-LEASE-RENEW • WIN-DHCP-IP-ASSIGN • Linux_DHCPACK • Generic_DHCPACK
AD Successful Login Events	x		x (resolvable by DNS or in FortiSIEM CMDB)	x (if in Event)	x			• Win-Security-540 • Win-Security-4624
AAA Successful Login Events	x			x			x	• Win-IAS-PassedAuth • CisACS_01_PassedAuth
VPN Successful Login Events	x			x			x	• Cisco-VPN3K-IKE/25 • ASA-722022 • ASA-713228 • ASA-713049-Client-VPN-Logon-success
WLAN Successful Login Events	x (if in Event)	x		x (if in Event)			x	• Cisco-WLC-53-bsnDot11StationAssociate
WLAN Discovery Events	x (if in Event)	x		x (if in Event)			x	• PH_DISCOV_CISCO_WLAN_HOST_LOCATION • PH_DISCOV_ARUBA_WLAN_HOST_LOCATION
VoIP Call Manager Discovery Events	x	x	x		x			• PH_DISCOV_VOIP_PHONE_ID

	IP	MAC	Host Name	User	Domain	VLAN	Loca-tion	Contributing event type
FortiSIEM L2 discovery Events	x	x	x (if resolvable by DNS or in FortiSIEM CMDB)			x	x	• PH_DISCOV_HOST_LOCATION

Creating New Identity Events

There may be a situation in which a new event type is added to FortiSIEM, and you want to use the parsed attributes of that event in the Identity and Location report. Once you have made sure that the event will parse correctly, you will need to edit the identityDef.xml file for your Supervisor and any Worker nodes in your deployment.

1. Log in to your Supervisor host machine as admin.
2. Change the directory to /opt/phoenix/config/xml.
3. Logon to FortiSIEM Super as admin
4. Edit the **identityDef.xml** file:
    1. Create a new `<identityEvent>`.
    2. For `<eventType>`, enter the ID of the event containing the identity attribute.
    3. For `<eventAttributes>`, enter the name of the event attribute and its corresponding identity attribute. For `reqd`, enter `yes` if the event must have this event attribute for use in the identity and location report.
    Possible location attributes include:
        - ipAddr
        - macAddr
        - computerName
        - domain
        - domainUser
        - aaaUser
        - vpnUser
        - geoCountry
        - geoState
        - geoCity
        - geoLatitude
        - vlanId
        - netEntryPt
        - netEntryPort
2. Restart identityMaster and identityWorker
3. Repeat for any Worker nodes.

This code sample is an example of a new **<identityEvent>** entry in the **identityDef.xml** file

```
<identityEvent>     <eventType>PH_DISCOV_CISCO_WLAN_HOST_LOCATION,PH_DISCOV_ARUBA_
WLAN_HOST_LOCATION</eventType>     <eventAttributes>        <eventAttribute name-
e="hostIpAddr" identityAttrib="ipAddr" reqd="no"/>        <eventAttribute name-
e="hostMACAddr" identityAttrib="macAddr" reqd="no"/>        <eventAttribute
name="user" identityAttrib="domainUser" reqd="no"/>        <eventAttribute name-
e="domain" identityAttrib="domain" reqd="no"/>        <eventAttribute name-
e="nepDevName" identityAttrib="netEntryPtName" reqd="yes"/>        <eventAttribute
name="nepDevIpAddr" identityAttrib="netEntryPt" reqd="yes"/>        <eventAt-
tribute name="nepDevPort" identityAttrib="netEntryPort" reqd="yes"/>
<eventAttribute name="wlanContrIpAddr" identityAttrib="wlanContrIpAddr" reqd-
d="yes"/>        <eventAttribute name="wlanContrHostName" iden-
tityAttrib="wlanContrHostName" reqd="yes"/>        <eventAttribute
name="hostGeoCountry" identityAttrib="geoCountry" reqd="no"/>        <eventAt-
tribute name="hostGeoState" identityAttrib="geoState" reqd="no"/>        <eventAt-
tribute name="hostGeoCity" identityAttrib="geoCity" reqd="no"/>
<eventAttribute name="hostGeoLatitude" identityAttrib="geoLatitude" reqd="no"/>
    <eventAttribute name="hostGeoLongitude" identityAttrib="geoLongitude" reqd-
d="no"/>     </eventAttributes>  </identityEvent>
```

Report Bundles

Report bundles are groups of reports for common IT infrastructure analytics, such as **Windows Server Health**. Be defining a bundle and placing reports into it, you can run all the reports at the same time, and apply the same filter conditions to all reports. You can view system and user-defined report bundles under **Analytics > Report Bundles**.

- Creating a Report Bundle
- Running a Report Bundle

Creating a Report Bundle

Creating a report bundle involves naming and describing the bundle, adding reports to the bundle, and then setting what you want to include in the report results.

1. Log in to your Supervisor node.
2. Go to **Analytics > Reports > Report Bundles**.
3. Click the **+** icon at the top of the Analytics navigation pane.
4. For **Group**, enter the name of the bundle, and then enter a **Description**.
5. Under **Select Group Members**, select the report category that contains the report you want to add to the bundle. When you select a category, all the reports in that category will be added to the selection window.
6. Select a report and use the **>>** button to add it to the bundle.
7. **Select Show Table** if you want all reports to include tables by default.
   You can set individual reports to show tables by selecting the report under **Show Reports**, clicking **Edit,** and then selecting **Show Table**.
8. Enter the number of **Rows per Table**.
9. Click **OK**.

Running a Report Bundle

1. Log in to your Supervisor node.
2. Go to **Analytics > Reports > Report Bundles**.
3. Select a report bundle to run.
4. At the top of the Analytics navigation pane, click the blue **Arrow** icon.
5. For multi-tenant deployments, select the **Organization** for which the reports should apply.
6. Select the **Time Range** for the results.
7. Set any **Data Conditions** to use in filtering the results.
   The most common use cases for setting data conditions involves imposing additional restrictions on the reporting devices, for example reporting devices IN a particular device group. These conditions are AND-ed to the filter conditions in every report of the bundle.
8. Click **Export**.
   The reports will run in the background, and when ready, you will see a dialog to save or download the PDF files.

Scheduling Report Bundles to Run

You can also schedule report bundles to run once or on recurring occasions in the future. Select a report bundle as you would to run it, and then click the **Clock** icon in the top-right corner of the Analytics navigation pane. Follow the steps described in Scheduling Reports to schedule the report bundles.

This screenshot shows the UI controls for working with report bundles.

Running System and User-Defined Reports and Baseline Reports

FortiSIEM includes a number of baseline reports for common data center analytics, as well as over 300 reports relating to IT infrastructure. You can also create your own reports. This topic describes how to run a system-generated or user-defined baseline report.

1. Log in to your Supervisor node.
2. Go to **Analytics > Reports** and select the subcategory containing the report you want to run**.**
   For baseline reports, select **Baseline**.
3. Select the report to run.
4. Click **Run Now** to run the report immediately, or **Run Later** to schedule the report.
5. If you chose **Run Now** and have a multi-tenant deployment, select the **Organization** for which you want to run the baseline report, and then click **OK**.
   The report will run and the results will be displayed.
   - For baseline results, the values in the **Profile Date Type** column indicate whether the baseline date type is a **Weekend (Saturday and Sunday) - 0** or **Workday - 1**. The values in **Hour of Day**, **1 - 24**, column indicate the time on which the baseline is based.
   - You can further refine the results of reports and baseline reports as described in Using Search Results to Refine Historical Searches.
   - For baseline reports, you can create scatter plots of the report results, use the **Quick Info** menu to get more information about items in the report results, and also view geolocation information about the results. For other types of reports you can use all the charts and other methods of refining results that are related to historical search.

Related Links

- Scheduling Reports
- Using Search Results to Refine Historical Searches
- System-Defined Baseline Reports
- Overview of Historical Search Results and Charts
- Using the Analysis Menu
- Using Geolocation Attributes in Rules
- Refining the Results from Historical Search

Scheduling Reports

You can schedule reports to run once or on recurring periods in the future. When the test runs, the results will be saved to the **Results** tab for the report, and in **Analytics > Generated Reports**.

Prerequisites

- When you schedule a report, you can specify notifications that should be sent for that report. In addition, you should make sure that the default settings for notifications for all scheduled reports have been set up.

Procedure

1. Log in to your Supervisor node.
2. Go to **Analytics > Reports**.
3. Select the report you want to schedule.
4. Click **Run Later**.
   The Schedule Tab: You can also schedule a report by going to the **Schedule** tab for the report. Click the **+** icon and follow the rest of the steps in this topic.
5. Select **Schedule this report for**:
6. For multi-tenant deployments, select the **Organization** for which this report should apply.
7. Select the **Report Time Range**.
8. Select the **Schedule Settings**.
9. Select the **Output Format,** whether you want to include the **Chart** in the output, and the **Maximum Rows to Display**.
10. Specify the **Notifications** that should be sent when the report runs.
    Click **Specify custom notifications** if you want to send notifications to specific email addresses.
    To copy the report to a remote directory, first define the remote location in **Admin > General Settings > Analytics > Report to be copied to this remote location when scheduler runs any report**. and then select **Copy to a remote directory** option.
11. Specify the amount of time the report should be retained after it has run.
12. Click **OK**.
    The report will run at the time you scheduled.

Related Links

- Setting Up Email Alert Routing for Scheduled Reports

Viewing Available Reports

The Synced Reports Group

The Synced Reports group contains the reports that can be synched with Report Server for use in Visual Analytics.

1. Log in to your Supervisor node.
2. Go to **Analytics > Reports**.
3. For multi-tenant deployments, select the **Organization** for which you want to view the available reports.
4. Expand the **Reports** list, and select the subcategory of report you want to view.
5. Select the report you want to view information about.
   Each report has four information tabs:

Report Tab	Description
Summary	Includes the **Filter** and **Group By** conditions for the report, and the report's **Display** attributes
Schedule	Information about when the report is scheduled to run. See Scheduling Reports for more information. You can click the + icon to set a schedule for the report to run.
Results	The results from any scheduled runs of the report, or results you have saved from running the report.
Definition	The XML definition of the report.

Audit

Audit Reports can be used to determine if a device is running the recommended OS and installed software versions, performance metrics are within bounds and harmful events have not triggered.

- Creating Audit Report
- Running an Audit
- Exporting Audit Results
- Scheduling an Audit

Creating Audit Report

To create an Audit Report

1. Go to **Analytics** tab
2. Expand **Audit** node on the left tree and go to the folder to which the new report will belong. You can also create a new folder first by clicking on the **+** on top of the left tree.
3. Click **New**.
4. Enter the following information for an Audit Report
    1. **Name**: Name of the Audit Report
    2. **Description**: Description of the Audit Report
    3. **Vendor**: Select a specific device vendor from the drop down list. The Audit Report will be specific to the chosen device vendor and model
    4. **Model**: Select a vendor specific model from the drop down list. The Audit Report will be specific to the chosen device vendor and model
    5. Specify **Failed Criteria** for the Audit Report. A device will fail the audit if **any** of the specified criteria is matched.
        1. **OS Version** Condition:
            1. Choose an **operator**: possible choices are IN, NOT IN, CONTAINS, NOT CONTAINS
            2. Specify **value** to be matched: this can be a comma separated list
        2. **Install Software** Condition:
            1. Specify Condition **name**. This is just for reference purposes.
            2. Specify **Install software name** - the name has to be exactly identical to the discovered installed software in CMDB > Devices > Installed Software > Name
            3. Choose an **operator**: possible choices are IN, NOT IN, CONTAINS, NOT CONTAINS
            4. Specify **value** to be matched: this can be comma separated list
        3. **Rules** Condition:
            a. Click **...** and the Rule selector dialog appears
            b. Select the appropriate Rule folder from the left most tree. If you do not know the specific folder, then choose the top level Rules folder.
            c. Select the rules from the middle section. You can also type a search string. You can expand the window and shrink the left most section to see more of the rule descriptions. The rules in the selected folder will appear in the middle section.
            d. Click **Items >>** to place the selected rules on the rightmost section
            e. Click **OK.**
        4. **Report** Condition:
            a. Click **...** and the Report selector dialog appears
            b. Select the appropriate Report folder from the left most tree. If you do not know the specific folder, then choose the top level Reports folder. The reports in the selected folder will appear in the middle section.
            c. Select the reports from the middle section. You can also type a search string. You can expand the window and shrink the left most section to see more of the report descriptions.
            d. Click **Items >>** to place the selected reports on the rightmost section
            e. Click **OK**.

Audit Policy Criteria Matching Notes

1. For each criteria, only devices in CMDB with vendor and model specified in the Audit Report is considered
2. If any of the criteria matches, then the device fails the audit
3. IN and NOT IN are exact match while CONTAINS and NOT CONTAINS are case insensitive sub-string match
4. For OS Version match, the entered value is compared with the Version column in CMDB > Device.
5. For Installed Software Version match, the entered value is compared with the Version column in CMDB > Device > Installed Software
6. For Rule match, the specified rule must trigger during the time interval specified in the Audit Report. Organization id and access IP of the device is compared to the Organization Id and Host IP in an incident.
7. For Report match, the specified reports run for the time duration specified in Audit Report must have data.

Running an Audit

To run an Audit,

1. Select an Audit Policy
2. Click **Run Now**
3. In the follow up dialog,
    a. Select the organizations for which to run the audit (meaningful for Service Provider version)
    b. Choose a time window - absolute or relative
    c. Click **OK**

The Audit Policy check results are displayed in the right bottom pane.

Summary tab shows a high level overview of the Audit Policy check.

- **Audit Result Distribution** chart shows the device pass/fail distribution for every selected organization.
- **Failed Criteria distribution** chart shows the contribution of each audit criteria to the devices that failed the audit.
- **Detail tab** shows the Audit Policy check for each device matching the vendor, model specified in the policy.
- **Organization** specifies the entity to which the device belongs
- **Device Name** is the host name of the device in CMDB
- **Audit Status** is the Pass/Fail flag
- **Details** specifes the reasons for Audit Policy check failure

Exporting Audit Results

To export an Audit Report,

1. Select an Audit Policy
2. Run the Audit Policy Check. The results will be shown in the bottom right pane.
3. Click **Export**
    a. Add **User Notes**
    b. Choose **Output Format** - PDF or CSV
    c. Click **Generate Report** - the PDF file will be stored in local disk

Scheduling an Audit

To schedule a report to run at a later time

1. Choose between one of two options
   - **Run this report for** - If the 'Run this report for' button is selected, a report will be scheduled for the super user, containing data from the organizations selected. The super user will be the owner of the report. The recipients of the report may be defined in the 'Send Notifications' section below or in Admin -> General Settings -> Analytics.
   - **Schedule this report for** - If the 'Schedule this report for' button is selected, multiple reports will be scheduled -- one for each selected organization -- and containing only that organization's data. The reports will be owned by the respective organizations. The recipients of the report are taken from Admin -> General Settings -> Analytics. When multiple reports are run in this way the notification recipients cannot be indicated in the 'Send Notifications' section below.
2. Select all the **Organizations** for which to run the Audit Report
3. Select the **Report time range**
4. Specify **Schedule settings** - when to run this report
5. Choose **Output Format** - PDF or CSV
6. Select **notification** - report recipients and method
   - If you choose **Send default notification**, then the settings in **Admin > General Settings > Analytics > Alerts to be sent when scheduler runs any REPORT,** is used.
   - If you choose **Specify custom notifications**, then you can specify email addresses.
   - If you choose **Copy to a remote directory**, then the settings in **Admin > General Settings > Analytics > Reports to be copied to this remote location when scheduler runs any REPORT,** is used.

Visual Analytics

Visual Analytics is an add-on for FortiSIEM that lets you create custome visualizations of FortiSIEM report data, as well as dashboards containing multiple visualization charts. FortiSIEM Visual Analytics has three components:

- The FortiSIEM Report Server, which syncs with and replicates FortiSIEM reports in near-real time.
- Tableau Server from Tableau Software, which enables the publication and distribution of your visualizations.
- Tableau Desktop, also from Tabeleau Software, which is your primary tool for creating visualizations.

See Installation and Configuration of FortiSIEM Visual Analytics for information about setting up FortiSIEM Report Server. For more detailed information about Tableau Server and Desktop, including installation, configuration, and examples of creating sheets and workbooks, you should consult the Product Support section of the Tableau Software website.

- FortiSIEM Visual Analytics Architecture
- Installation and Configuration of FortiSIEM Visual Analytics
- Working with the Report Server
- Installing and Configuring Tableau Server
- Creating and Managing Workbooks

Visual Analytics Architecture

Overview and Report Server Architecture

With FortiSIEM Visual Analytics, you can now create visual representations of the data that is stored in FortiSIEM. This includes:

- **Structured data** stored in the FortiSIEM CMDB relational PostgreSQL database, such as:
  - Discovered information about devices, systems, applications and users
  - Identity and location information
  - Incidents and notifications
- **Unstructured data** such as logs, events, performance metrics etc. that are monitored by FortiSIEM and stored in the EventDB NoSQL database, which is accessible by Supervisors and Workers over NFS.

In order to provide near real-time visual analytics without compromising the performance of your FortiSIEM deployment, both structured and unstructured data is exported to a separate virtual machine, the **FortiSIEM Report Server**, running PostgreSQL. The Report Server contains two databases that are queried by FortiSIEM Visual Analytics:

- **phoenixdb**
  This database contains the entire FortiSIEM CMDB and is populated via asynchronous PostgreSQL replication (slony) in near-real time.
- **reportdb**
  This database contains the results of event queries

You can find more information about FortiSIEM Report Server in the topic Report Server Architecture: phoenixdb and reportdb and its related topics.

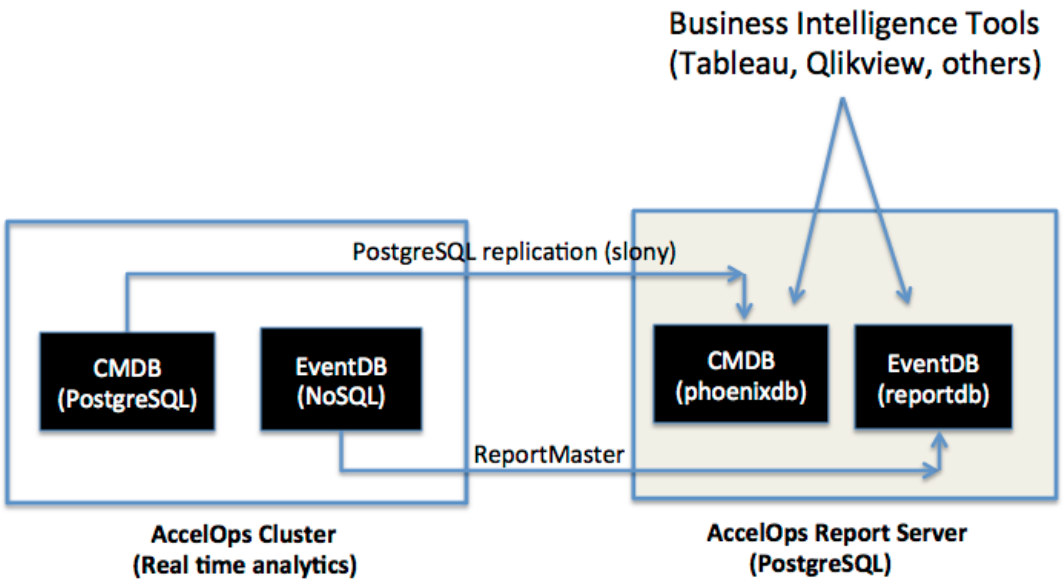

Using FortiSIEM Report Server with Tableau Software

FortiSIEM Report Server integrates with Tableau Software to provide the interface for creating and publishing your data visualizations. Workbooks containing visualizations based on FortiSIEM data are created using Tableau Desktop, and then are published to Tableau server, where they can be accessed on any Windows or OS X device by users how have been granted permission for viewing or editing them. FortiSIEM provides some workbooks for visualizations, but you can construct others for custom analytics. You can find more information about workbooks in the section Creating and Managing Workbooks.

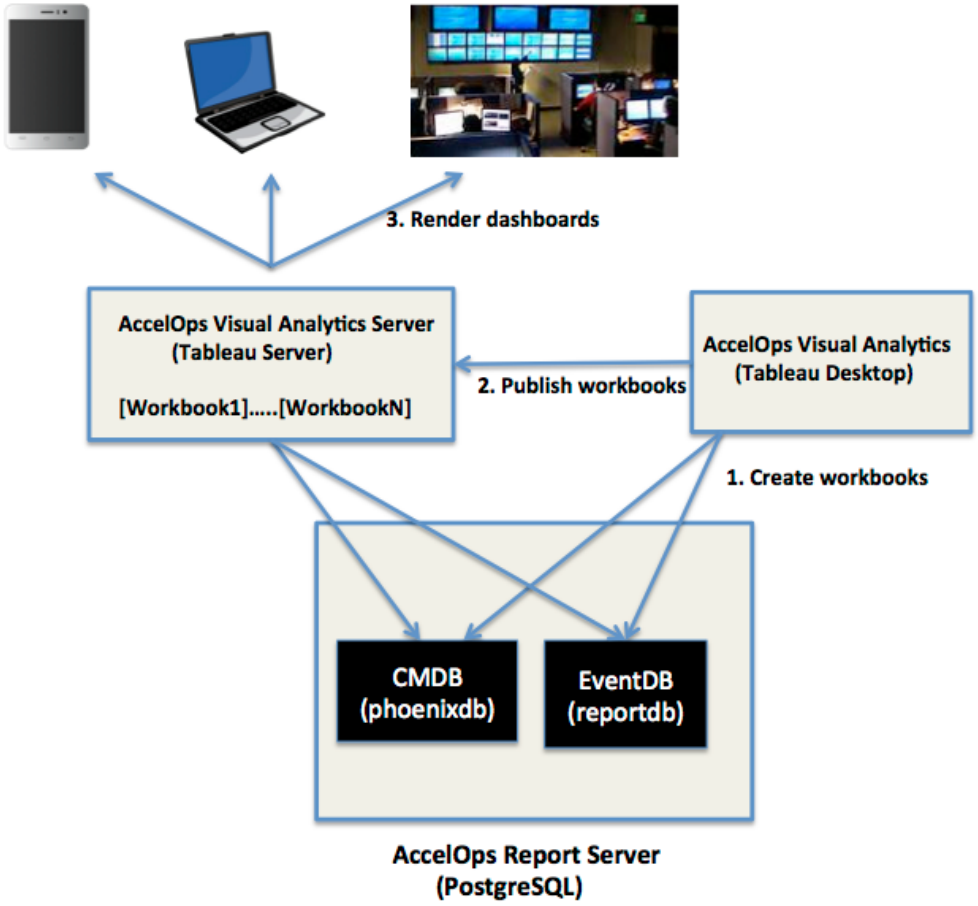

Installation and Configuration of FortiSIEM Visual Analytics

Installation and configuration of FortiSIEM Visual Analytics involves setting up FortiSIEM Report Server, and then integrating it with Tableau Server and Desktop from Tableau Software. Topics in this section contain setup and configuration instructions for Report Server. For information on setting up and configuring Tableau Server and Desktop, see the online Tableau Software documentation.

- Requirements for Visual Analytics Report Server
- Setting Up Visual Analytics
- Hypervisor Installations for Report Server
- Syncing with the Report Server

Requirements for Visual Analytics Report Server

You install Visual Analytics Report Server as FortiSIEM node, and these requirements assume that you have already set up and installed FortiSIEM. If you are working with a fresh install of FortiSIEM that includes Report Server, see the topics under Installation for complete requirements and installation instructions for the FortiSIEM Virtual Appliance.

Dedicated Machine for Report Server: You must install Visual Analytics Report Server on a dedicated machine.

Hardware Requirements for Report Server Nodes

Component	Quantity	Host SW	Processor	Memory	OS/App Storage	Reports Data Storage (1 year)
Report Server	1	ESX	8 Core 3 GHz, 64 bit	16 GB	200GB (80GB OS/App, 60GB CMDB, 60GB SVN)	See recommendations under Hardware Requirements for Supervisor and Worker nodes

Setting Up Visual Analytics

There are three components to FortiSIEM Visual Analytics:

- FortISIEM Report Server
- Tableau Server
- Tableau Desktop

Setting up Visual Analytics involves setting up each of those components in order, and establishing the relationship between them.

1. You must first install Report Server as described in Installing and Registering FortiSIEM Report Server in VMware ESX.
2. After installing Tableau Server on a Windows server, and installing Tableau Desktop on a Windows or Mac OS X device, you then connect the two systems as described in the Tableau Software product documentation.
3. When this connection is established, it automatically triggers the remote registration and configuration of the FortiSIEM Report Server, including replication of the CMDB and EventDB data from the FortiSIEM Cluster to the FortiSIEM Report Server, as well as the user account required for access to the original databases.

Registration of the Report Server and replication of the FortiSIEM database data may take some time depending on the size of the original CMDB. Registration is complete when the replication process catches up with the latest data in the system. From that point on, replication from the CMDB to FortiSIEM Report Server takes place in near real time, letting you run Visual Analytics queries against CMDB data that has been replicated to the Report Server's phoenixdb.

You can find full information about setting up all components of FortiSIEM Visual Analytics in the section Installation and Configuration of FortiSIEM Visual Analytics

Hypervisor Installations for Report Server

These topics cover the installation of Report Server in various hypervisor enviroments.

- Installing and Registering FortiSIEM Report Server in Amazon Web Services
- Installing and Registering FortiSIEM Report Server in KVM
- Installing and Registering FortiSIEM Report Server in Microsoft Hyper-V
- Installing and Registering FortiSIEM Report Server in VMware ESX

Installing and Registering FortiSIEM Report Server in Amazon Web Services

Follow the instructions for setting up FortiSIEM virtual appliance as described in Setting Up Supervisor, Worker and Collector Nodes in AWS section in the Installation and Upgrade Guide, and then register the Report Server to the Supervisor as described in Installing and Registering FortiSIEM Report Server in VMware ESX.

Turn on archive mode for Report server CMDB replication

1. Mount a NFS shared directory on both Super and report server and make sure that this mount can survive system reboot. For example:
2. Make this shared directory own by postgres.postgres:
3. On Super, edit postgresql.conf under /cmdb/data to turn on archive mode by uncommenting (removing # in the first column) the following lines and make sure archive_command points to the correct directory which is created in step 1.

```
archive_mode = on # allows archiving to be done
# (change requires restart)
archive_command = 'cp %p /data/replication/archive/%f'
```

4. On Report Server, edit /cmdb/data/recovery.conf and uncomment the following lines and make sure restore_ command and

```
restore_command = 'cp /data/replication/archive/%f %p'
archive_cleanup_command = 'pg_archivecleanup
/data/replication/archive %r
```

5. On Super, restart postgresql DB 'service postgresql-9.1 restart'
6. On Super, restart App Server (Glassfish)
7. On Report Server, restart postgresql DB 'service postgresql-9.1 restart'

Registering Report Server

1. In the Admin tab, select License Management.
2. Under Report Server Information, click Add.
3. Enter the Report Server IP Address, and the Database Username and Password you want to use to administer Report Server.
4. These are also the credentials that you will use when you set up the Visual Analytics Server to read data from Report Server.
5. Click Run in Background if you want Report Server registration to run in the background for larger installations.
6. When CMDB size is under 1GB, registration takes approximately 3 minutes to complete.
7. When registration completes, click OK in the confirmation dialog.
8. Under the Admin tab, select Cloud Health and make sure Report Server is up and running.

Installing and Registering FortiSIEM Report Server in KVM

Follow the instructions for installing FortiSIEM virtual appliance as described in Importing a Supervisor, Collector, or Worker Image into KVM, and then register the Report Server with the Supervisor as described in Installing and Registering a Report Server Node in ESX.

Turn on archive mode for Report server CMDB replication

1. Mount a NFS shared directory on both Super and report server and make sure that this mount can survive system reboot. For example: /data/replication/archive
2. Make this shared directory own by postgres.postgres
3. On Super, edit postgresql.conf under /cmdb/data to turn on archive mode by uncommenting (removing # in the first column) the following lines and make sure archive_command points to the correct directory which is created in step 1.

```
archive_mode = on               # allows archiving to be done
                                # (change requires restart)
archive_command = 'cp %p /data/replication/archive/%f'
```

4. On Report Server, edit /cmdb/data/recovery.conf and uncomment the following lines and make sure restore_ command and archive_cleanup_command are pointing to the directory created in step 1:

```
restore_command = 'cp /data/replication/archive/%f %p'
archive_cleanup_command = 'pg_archivecleanup /data/replication/archive %r
```

5. On Super, restart postgresql DB 'service postgresql-9.1 restart'
6. On Super, restart App Server (Glassfish)
7. On Report Server, restart postgresql DB 'service postgresql-9.1 restart'.

Registering Report Server

1. In the **Admin** tab, select **License Management**.
2. Under Report Server Information, click **Add**.
3. Enter the **Report Server IP Address**, and the **Database Username** and **Password** you want to use to administer Report Server.
   These are also the credentials that you will use when you set up the Visual Analytics Server to read data from Report Server.
4. Click **Run in Background** if you want Report Server registration to run in the background for larger installations. When CMDB size is under 1GB, registration takes approximately 3 minutes to complete.
5. When registration completes, click **OK** in the confirmation dialog.
6. Under the **Admin** tab, select **Cloud Health** and make sure Report Server is up and running.

Installing and Registering FortiSIEM Report Server in Microsoft Hyper-V

Follow the virtual appliance installing instructions in Installing in Microsoft Hyper-V, and then register the Report Server node with the Supervisor as described in Installing and Registering FortiSIEM Report Server in VMware ESX.

Turn on archive mode for Report server CMDB replication

1. Mount a NFS shared directory on both Super and report server and make sure that this mount can survive system reboot. For example: /data/replication/archive
2. Make this shared directory own by postgres.postgres
3. On Super, edit postgresql.conf under /cmdb/data to turn on archive mode by uncommenting (removing # in the first column) the following lines and make sure archive_command points to the correct directory which is created in step 1.

```
archive_mode = on               # allows archiving to be done
                                # (change requires restart)
archive_command = 'cp %p /data/replication/archive/%f'
```

4. On Report Server, edit /cmdb/data/recovery.conf and uncomment the following lines and make sure restore_ command and archive_cleanup_command are pointing to the directory created in step 1:

```
restore_command = 'cp /data/replication/archive/%f %p'
archive_cleanup_command = 'pg_archivecleanup /data/replication/archive %r
```

5. On Super, restart postgresql DB 'service postgresql-9.1 restart'
6. On Super, restart App Server (Glassfish)
7. On Report Server, restart postgresql DB 'service postgresql-9.1 restart'.

Registering Report Server

1. In the **Admin** tab, select **License Management**.
2. Under Report Server Information, click **Add**.
3. Enter the **Report Server IP Address**, and the **Database Username** and **Password** you want to use to administer Report Server.
   These are also the credentials that you will use when you set up the Visual Analytics Server to read data from Report Server.
4. Click **Run in Background** if you want Report Server registration to run in the background for larger installations. When CMDB size is under 1GB, registration takes approximately 3 minutes to complete.
5. When registration completes, click **OK** in the confirmation dialog.
6. Under the **Admin** tab, select **Cloud Health** and make sure Report Server is up and running.

Installing and Registering FortiSIEM Report Server in VMware ESX

These instructions are for installing Report Server on VMWare ESX, and assume that you have already installed and configured FortiSIEM environment. For instructions for a complete FortiSIEM install, see the topics under Installation.

Installing Report Server

Follow the instructions for installing FortiSIEM virtual appliance as described in the Installation and Upgrade Guide under topics 'Installing a Supervisor, Worker, or Collector Node in ESX' and 'Configuring the Supervisor, Worker, or Collector from the VM Console'.

Turn on archive mode for Report server CMDB replication

1. Mount a NFS shared directory on both Super and report server and make sure that this mount can survive system reboot. For example: /data/replication/archive
2. Make this shared directory own by postgres.postgres
3. On Super, edit postgresql.conf under /cmdb/data to turn on archive mode by uncommenting (removing # in the first column) the following lines and make sure archive_command points to the correct directory which is created in step 1.

```
archive_mode = on

# allows archiving to be done

# (change requires restart)
archive_command = 'cp %p /data/replication/archive/%f'
```

4. On Report Server, edit /cmdb/data/recovery.conf and uncomment the following lines and make sure restore_ command and archive_cleanup_command are pointing to the directory created in step 1:

```
restore_command = 'cp /data/replication/archive/%f %p'
archive_cleanup_command = 'pg_archivecleanup /data/replication/archive %r
```

5. On Super, restart postgresql DB 'service postgresql-9.1 restart'
6. On Super, restart App Server (Glassfish)
7. On Report Server, restart postgresql DB 'service postgresql-9.1 restart'

Registering Report Server

1. In the **Admin** tab, select **License Management**.
2. Under Report Server Information, click **Add**.
3. Enter the **Report Server IP Address**, and the **Database Username** and **Password** you want to use to administer Report Server.
   These are also the credentials that you will use when you set up the Visual Analytics Server to read data from Report Server.
4. Click **Run in Background** if you want Report Server registration to run in the background for larger installations. When CMDB size is under 1GB, registration takes approximately 3 minutes to complete.
5. When registration completes, click **OK** in the confirmation dialog.
6. Under the **Admin** tab, select **Cloud Health** and make sure Report Server is up and running.

Syncing with the Report Server

Using FortiSIEM Visual Analytics involves first syncing reports contained in the primary FortiSIEM application to the FortiSIEM Report Server.

1. Log in to your Supervisor node.
2. Go to **Analytics > Reports > Synced Reports**.
3. Select a report.
   Currently only reports that contain a **Group By** condition can be synced. Both system and user-created reports can be synched as long as the contain a Group By condition.
4. Select **Sync**.

When the sync process initiates, the Supervisor node dynamically creates a table within the Report Server reportdb database. When the sync is established, it will run every five minutes, and the last five minutes of data in the synced report will be pushed to the corresponding table. This lets you run Visual Analytics on event data stored in the Report Server reportdb database.

Working with the Report Server

This section contains information on FortiSIEM Report Server architecture, viewing and querying CMDB and Event data in contained in the Report Server databases, and database maintenance.

- Report Server Architecture: phoenixdb and reportdb
- Working with CMDB Data in FortiSIEM Report Server
- Working with Event Data in FortiSIEM Report Server

Report Server Architecture: phoenixdb and reportdb

FortiSIEM Report Server contains two databases:

- **phoenixdb**
  This database contains the entire FortiSIEM CMDB and is populated via asynchronous PostgreSQL replication (slony) in near-real time.
- **reportdb**
  This database contains the results of event queries.

Topics in this section describe how to view the tables in these databases, and how those tables are organized. For viewing the tables, we recommend using the pgAdmin PostgreSQL database utility, which you can download from the pgAdmin website.

Working with CMDB Data in FortiSIEM Report Server

Data from the FortiSIEM CMDB database is populated to the FortiSIEM Report Server and stored in the Report Server phoenixdb. This section contains information on how to view the organization of phoenixdb, and write queries against the data it contains.

- Viewing phoenixdb Organization
- Querying Incident Data in FortiSIEM Report Server
- Querying Other CMDB Tables in FortiSIEM Report Server

Viewing phoenixdb Organization

This database contains the contents of the entire FortiSIEM CMDB database, including incidents.

1. In the pgAdmin utility, go to **File > Add Server**.
2. In the **New Server Registration** dialog, enter connection details for FortiSIEM Report Server.
   For **Maintenance DB**, select **phoenixdb**.
   For **Username** and **Password**, use the read-only user name and password that you created when you provisioned the Report Server.
3. Click **OK**.
   When the connection to the FortiSIEM Report Server is established, phoenixdb will load in the Object browser. There are approximately 197 tables in phoenixdb, which are replicated from the FortiSIEM cluster.
4. Select a table to view, then right-click to open the **Options** menu.
5. In the **Options** menu, select **View Data**, and then select an option for which rows you want to view.
   For example, to view the contents of the `ph_device` table, which contains CMDB information about discovered devices, you would select and then right click on `ph_device`, then select **View Data > View All Rows**.

You can also use this method to examine Views and other objects in the phoenixdb database.

Querying Incident Data in FortiSIEM Report Server

There are two ways to look at the incident data inside FortiSIEM Report Server:

- **Incident Tables** (`ph_incident` and `ph_incident_detail`)
  Contains the incidents
- **Incident View** (`ph_incident_view`)
  This is a database view that adds other context to the incident tables by joining with other tables in the database. Added information includes location and business service. Some information is parsed out for easier query, such as host names and IP address fields from `incident_source`, and `incident_target` fields in `ph_incident` are parsed out as separate fields in `ph_incident_view`.

This topic describes how to view the data contained in **Incident View**.

1. Follow the instructions in Viewing phoenixdb Organization to access the phoenixdb database in FortiSIEM Report Server.
2. Go to **Views > ph_incident_view > Columns** to view the table columns.
3. Go to **Views > ph_incident_view > View Data >** View Last 100 Rows to view the incidents.

Reference: Attribute Columns in the ph_incident_view Table

Column Name	Format	Description
incident_id	integer	Unique id for an incident
cust_org_id	integer	Customer Id (for AO-SP)
first_seen_time	integer	The time when the incident was first seen. The format is UNIX time but with milliseconds granularity. It is defined as the number of milliseconds that have elapsed since 00:00:00 Coordinated Universal Time (UTC), Thursday, 1 January 1970
last_seen_time	integer	The time when the incident was last seen. The format is UNIX time but with milliseconds granularity. It is defined as the number of milliseconds that have elapsed since 00:00:00 Coordinated Universal Time (UTC), Thursday, 1 January 1970
incident_et	string	Incident event type id e.g. PH_RULE_SERVER_HW_CRITICAL
incident_status	integer	0: Active 1: Auto Cleared 2: Manually Cleared 3: System Cleared
incident_count	integer	The number of times this exact incident (with the same parameters: source, destination etc has happened)
biz_name	string	Associated business service name
severity	integer	Numerical severity of the incident - range 0-10
severity_cat	string	Incident severity category: 0-4: LOW, 5-8: MEDIUM and 9-10: HIGH
orig_device_ip	string	IP address of the device that reported the incident
ph_incident_category	string	Category of infrastructure affected by this incident: possible valies: Network, Server, Storage, Virtualization, Application, Internal
incident_src	string	Incident Source string formatted as a list of <Attribute>:Value; e.g. srcIpAddr:10.1.1.1,srcName:JoeLaptop
src_ip_addr	string	Source IP parsed out from incident_src field
src_name	string	Source Name parsed out from incident_src field

Column Name	Format	Description
src_device_location	string	(Geo) Location display name string for the object specified in incident_src
src_country	string	(Geo) Country name string for the object specified in incident_src
src_state	string	(Geo) State name for the object specified in incident_src
src_building	string	(Geo) Building name for the object specified in incident_src
src_floor	string	(Geo) Floor for the object specified in incident_src
src_latitude	double	(Geo) Latitude for the object specified in incident_src
src_longitude	double	(Geo) Latitude for the object specified in incident_src
incident_target	string	Incident Destination string formatted as a list of <Attribute>:Value; e.g. "destIpAddr:10.1.1.1,destName:JoeLaptop" or "hostIpAddr:10.1.1.1,hostName:JoeLaptop"
dest_ip_addr	string	Destination IP parsed out from incident_target field
dest_name	string	Destination Name parsed out from incident_target field
dest_device_location	string	(Geo) Location display name string for the object specified in incident_target
dest_country	string	(Geo) Country name string for the object specified in incident_target
dest_state	string	(Geo) State name for the object specified in incident_target
dest_building	string	(Geo) Building name for the object specified in incident_target
dest_floor	string	(Geo) Floor for the object specified in incident_target
dest_latitude	double	(Geo) Latitude for the object specified in incident_target
dest_longitude	double	(Geo) Longitude for the object specified in incident_target
host_ip_addr	string	Host IP address parsed out from incident_target field
host_name	string	Host Name parsed out from incident_target field

Column Name	Format	Description
host_device_location	string	(Geo) Location display name string for the object specified in incident_target - populated if incident_target contains hostIpAddr
host_country	string	(Geo) Country name string for the object specified in incident_target - populated if incident_target contains hostIpAddr
host_state	string	(Geo) State name for the object specified in incident_target - populated if incident_target contains hostIpAddr
host_building	string	(Geo) Building name for the object specified in incident_target - populated if incident_target contains hostIpAddr
host_floor	string	(Geo) Floor for the object specified in incident_target - populated if incident_target contains hostIpAddr
host_latitude	double	(Geo) Latitude for the object specified in incident_target - populated if incident_target contains hostIpAddr
host_longitude	double	(Geo) Longitude for the object specified in incident_target - populated if incident_target contains hostIpAddr
vm_name	string	VM Name if incident involves a Virtual machine - populated if incident_target contains vmName
user_attr	string	User name if incident involves user, i.e. incident_target contains user
target_user_attr	string	Target user name if incident involves user, i.e. incident_target contains targetUser
ldap_domain	string	Domain if incident involves user, i.e. incident_target contains domain
computer	string	Computer name incident_target contains computer
target_computer	string	Target Computer name incident_target contains targetComputer
incident_details	string	Incident Details containing evidence on why the incident triggered e.g. Triggered Event Count = 90 or AVG(CPUUtil) = 90 etc

Sample Incident Queries

- Show Incident Categories with Severity and Frequency Occurrence
- Show Incident Location

Show Incident Categories with Severity and Frequency Occurrence

This query will show which parts of the infrastructure are triggering events.

1. Follow the instructions in Viewing phoenixdb Organization to access the phoenixdb in FortiSIEM Report Server.
2. Under **Views**, select `ph_incident_view`.
3. In pgAdmin, click on the **SQL** icon in the menu bar to open the SQL query window.
4. Enter this SQL query:

```
SELECTph_incident category, incident_et, severity_cat, src_ip_addr,
host_name, COUNT(*)FROMph_incident_viewGROUPBYph_incident category,
incident_et, severity_cat, src_ip_addr, host_nameORDERBYCOUNT
(*) DESC;
```

5.
6. When the query executes, you will see a list of matching incidents in the Output Pane.

Show Incident Location

1. Follow the instructions in Viewing phoenixdb Organization to access the phoenixdb in FortiSIEM Report Server.
2. Under **Views**, select `ph_incident_view`.
3. In pgAdmin, click on the **SQL** icon in the menu bar to open the SQL query window.
4. Enter this SQL query:

```
SELECT host_device_location, severity_cat, ph_incident_category,
COUNT(*)FROM ph_incident_viewGROUP BY host_device_location, ph_incident_
category, severity_catORDER BY host_device_location ASC, severity_cat ASC
COUNT(*) DESC;
```

5. When the query executes, you will see a list of incidents and their locations in the Output Pane.

Working with Event Data in FortiSIEM Report Server

Data from the FortiSIEM EventDB database is populated to the FortiSIEM Report Server and stored in the Report Server reportdb. This section contains information on how to view the organization of reportdb, and write queries against the data it contains.

- Viewing reportdb Organization
- Syncing FortiSIEM Report with Report Server
- Deleting a Report from FortiSIEM Report Server
- Modifying an Existing Report in FortiSIEM Report Server

Viewing reportdb Organization

This database contains the reports that are synched from the FortiSIEM cluster.

1. In the pgAdmin utility, go to **File > Add Server**.
2. In the **New Server Registration** dialog, enter connection details for FortiSIEM Report Server.
   For **Maintenance DB**, select **reportdb**.
   For the **Port** enter **30000** (default port used for the reported).
   For **Username** and **Password**, use the read-only user name and password that you created when you provisioned the Report Server.
3. Click **OK**.
   When the connection to the Report Server is established, reports will load in the Object browser.
4. Select a table to view, then right-click to open the **Options** menu.
5. In the **Options** menu, select **View Data** , and then select an option for which rows you want to view.

Syncing FortiSIEM Report with Report Server

1. Log in to FortiSIEM.
2. Go to **Analytics > Reports**.
3. Select a report.
   Any reports with a **Sync** checkbox can be synced. Run the report to make sure it contains some data.
4. For each report you want to sync, select the **Sync** checkbox.
   AO-SP: In the **Sync Details** dialog, select the organizations whose data needs to be synced.
5. Click **OK**.
6. After several minutes, follow the instructions in Viewing reportdb Organization to view the reportdb database.
7. Under **Tables**, you should now see the synced reports.

Table Structure for Synced Reports

When you sync FortiSIEM report to FortiSIEM Report Server, two pairs of tables are created in reportdb, one pair for each organization in the case of AO-SP. For each organization, multiple tables are created:

1. A **parent table** containing data for all months: the table name is of the form `<Report Name>_<ID>_<custId>`
2. A **child table** for the current month: `<Report Name>_<ID>_<custId>_<yYYYYmMM>` where `YYYY` is the year and `MM` is the month.

Queries should be written using the parent table. To see data in the parent table, follow the instructions in Viewing reportdb Organization . The reportdb database fields are generated from the display fields in FortiSIEM report definitions. Only the field `report_time` is added to the Report Server table definitions to capture the time when the particular report is generated. For example, if you synced the report **Network Devices by CPU, Memory**, you would see these fields:

Field	Description
`report_time`	UNIX time at which the report is generated. **Unix time** (or **POSIX time** or **Epoch time**) is a system for describing instants in time, defined as the number of seconds that have elapsed since 00:00:00 Coordinated Universal Time (UTC), Thursday, 1 January 1970 not counting leap seconds.
`hostName`	Host Name of the device for which CPU and memory are being measured
`hostIpAddr`	Access IP of the device for which CPU and memory are being measured
`AVG(cpuUtil)`	Average of all the CPU utilization metrics within the last 5 minutes ending with report_time
`AVG(memUtil)`	Average of all the CPU utilization metrics within the last 5 minutes ending with report_time

Deleting a Report from FortiSIEM Report Server

1. Log in to FortiSIEM.
2. In **Analytics > Reports > Synced Reports**, select the report you want to delete.
3. In the **Sync Details** dialog, clear the **Sync** option for the report, and then click **OK**.
   The report will no longer be synced with Report Server. You can verify this by making sure the **Sync** option is not selected for the report on the **Analytics > Reports > Synced Reports** page. You can now delete the report from FortiSIEM Report Server.
4. Log in to FortiSIEM Report Server via SSH and navigate to the directory `/opt/phoenix/deployment/jumpbox`.
5. Run the `phreportdbmanager.py` command, along with the table name and date as arguments, to delete the report.

```
phreportdbmanager.py --remove tablenames='"Network Devices By CPU, Memory_
1278492569_1"'reporttimes=2014-10
```

Viewing the Names of Reports in Report Sever: Use the pgAdmin utility to view the names of all tables and reports in Report Server, as described in

Viewing reportdb Organization.

When the deletion process completes, you will see a command line output like this:

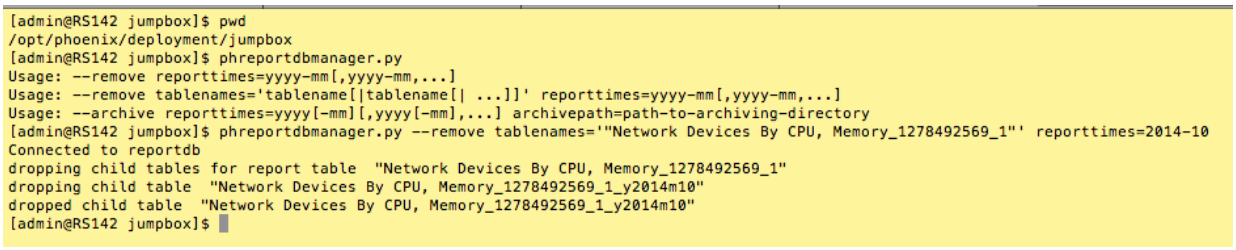

6. After you have deleted the table containing the report information, you will need to delete the parent table, which will now be empty of content, using the same `phreportdbmanager.py` command.

Modifying an Existing Report in FortiSIEM Report Server

Suppose a system report is synced and exported to FortiSIEM Report Server. When you modify that report in FortiSIEM, you must rename it, at which point it becomes a user report. When you then sync that report for FortiSIEM Report Server, a new table is created on the FortiSIEM Report Server.

Suppose now that you have a user-defined report that is already synced to the FortiSIEM Report Server, but you modify it inline in FortiSIEM, which means that you have changed the report conditions without changing the report name. This will cause a change in the table, but a new table will not be created. Here are some examples of inline modifications, and how they affect the structure of the table as well as the data collected in the table:

Modification	Effect
GROUP BY field added	The corresponding table has the new GROUP BY field, but only newer data populates the field
GROUP BY field removed	There is no change in the corresponding table, and newer data does not populate the field
GROUP BY field changed	For example, the field `srclpAddr` is changed to `destlpAddr`. Both fields are retained, but newer data populates `destlpAddr`.
Aggregated fields added	The corresponding table has the new field, but only newer data populate that field
Aggregated field removed	There is no change in the corresponding table, and newer data does not populate the field
Aggregated Field Changed	For example, `AVG(cpUutil)` is changed to `MAX(cpuUtil)`. Both fields are retained, but newer data populates `MAX (cpuUtil)`.

Installing and Configuring Tableau Server

Prerequisites

Before you begin installing Tableau Server, make sure you have read the section on **Tableau Server** in Requirements for Visual Analytics Report Server. This contains information on the Administrator Account and Ports that you will need during the configuration process. You may want to also consult the Tableau Server Administration Guide before you begin the installation process.

Installation

1. Download the installation file from Tableau Software.
2. Double-click the installation file to launch the Setup Wizard.
3. When the Setup Wizard launches, click **Next** to begin the installation process.
4. Enter a **Destination Location** where you want to install the server files, and then click **Next**.
5. When the system verification process completes, click **Next**.
6. Enter a location for the Start Menu folder, or use the default location, and then click **Next**.
7. Click **Install** to complete the installation process.
8. Click **Next** to begin the server activation process.

Activation

1. If you are evaluating Tableau Server, click **Start trial now**. Otherwise, click **Activate the product** to enter a license key.
2. If you enter a license key, click **Activate**.
3. Click **Continue** to launch the Tableau Server configuration process.

Configuration

1. In the Configuration dialog, enter a **User Name** and **Password** for the domain admin account that you will use to administer the Tableau Server.
2. If necessary, enter a **Gateway** port through which you will connect to the server over HTTP.
3. Click **OK**.
   The initialization process will launch and complete within several minutes.
4. Click **Finish** to complete the configuration process.
5. Launch the Tableau Server user interface by entering the URI for the server in a browser window.
   The URI will be be in the format of `http://<Windows_Server_IP_Address>:<Port_Number_Used_In_Step_2>`
6. Sign in to the server by entering the credentials for the domain admin account that you created in Step 1, and then click **Sign In**.

7. Click the **Admin** tab and select **Maintenance**.
8. Under **Status**, check to make sure that all systems are up and running.

You are now ready to install Tableau Desktop. After you have completed the Desktop installation process and connect to Report Server for the first time to create a sheet, as described in Creating a Single Sheet Workbook, you will also establish the connection between FortiSIEM Report Server and Tableau Server.

Creating and Managing Workbooks

This section contains information on using Visual Analytics Desktop to create sheets and workbooks that are based on FortiSIEM reports, and then publishing them for others to use.

- Viewing Workbooks
- Creating and Publishing Workbooks
- Adding Users to Workbooks

Viewing Workbooks

1. Log in to Visual Analytics Server.
2. Click the **Content** tab and select **Workbooks**.
3. Click on a workbook.
   The workbook along with the various worksheets are displayed.
4. Select a workbook or worksheet.
5. You will be prompted for credentials that will allow the workbook or worksheet to access database information. Enter the Admin credential that you used to set up FortiSIEM Report Server and click **OK**.
6. When your credential is accepted, the chart associated with the selected workbook or worksheet will be displayed.

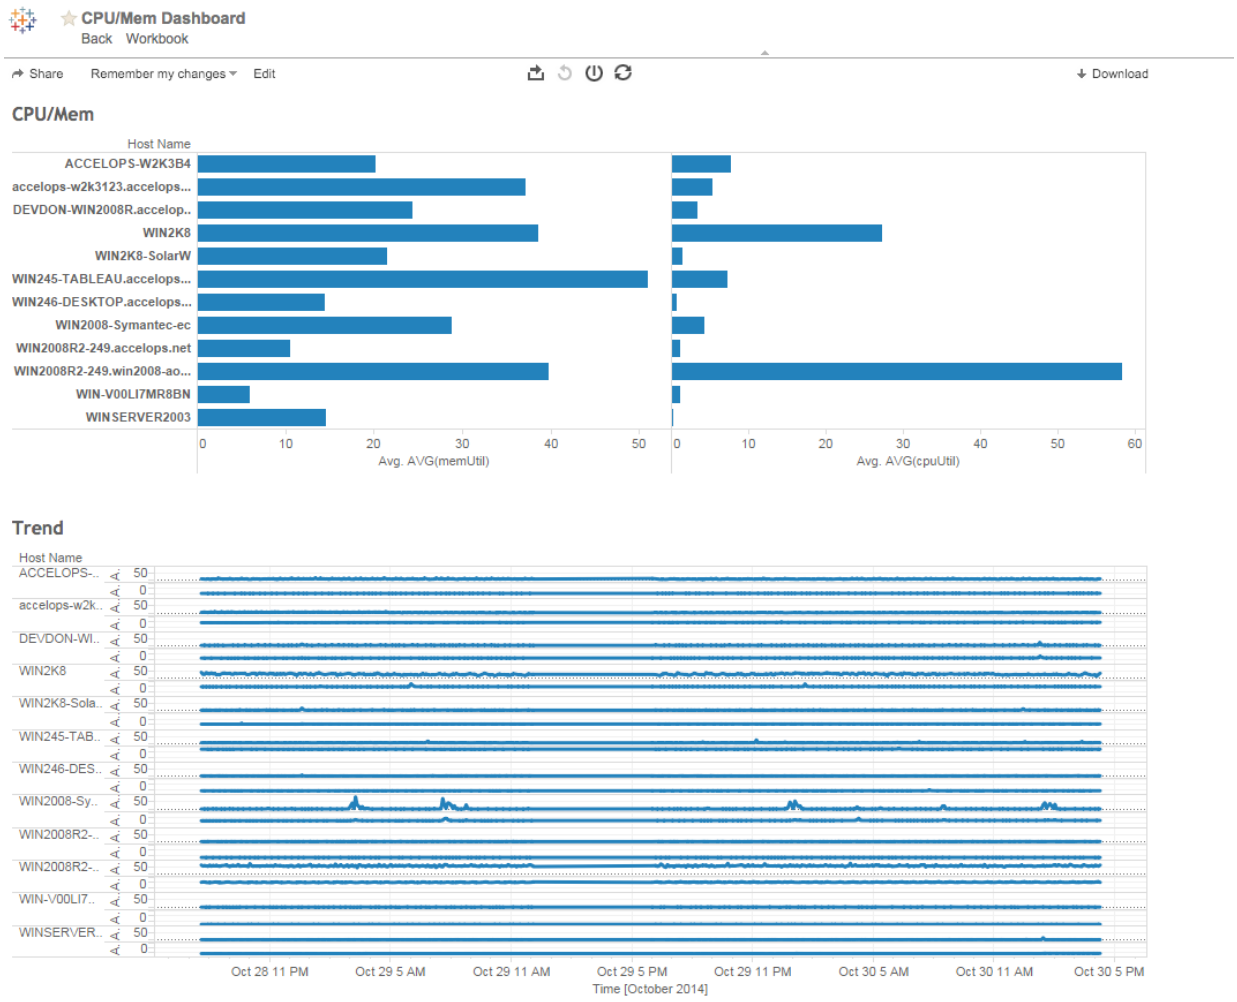

Creating and Publishing Workbooks

Workbooks are collections of FortiSIEM reports that have been synced to FortiSIEM Report Server, and which are then the basis for charts and dashboards that can be published to Visual Analytics Server for access by other users. Information in this section describes how to create single and multiple sheets of report information, and then make them accessible to other users.

- Creating a Single Sheet Workbook
- Creating a Multiple Sheet Workbook
- Using FortiSIEM Workbooks with Tableau Visual Analytics Desktop and Server

Creating a Single Sheet Workbook

These instructions demonstrate how to create a single-sheet workbook that will chart the CPU and memory utilization trend for various servers. This example uses the **Servers by CPU, Memory** report and its associated table, but any report with a table in the reportdb database can also be used. The Tableau Desktop online Help also contains extensive information about building sheets and workbooks with the Tableau Desktop editor, which powers the FortiSIEM Visual Analytics Desktop.

- Prerequisites
- Procedure

Prerequisites

- Follow the instructions in Syncing FortiSIEM Report with Report Server to sync the report you want to use for your worksheet.
- You will need to know the name of the parent table for your synced report. Follow the instructions in Viewing reportdb Organization to find the table that corresponds to your report.

Procedure

Create the Sheet

1. Launch Tableau Visual Analytics Desktop.
2. Connect to FortiSIEM Report Server with the **Username** and **Password** that you used during Report Server installation.
   - For **Database**, enter **reportdb**
   - For **Port**, enter **30000**

     Connecting to Port 30000: It's important to make sure you enter the correct port to connect to the reportdb database. If you leave this option blank you will connect to the default PostgreSQL database port, which will connect you with phoenixdb instead of reportdb. For more information about the databases contained in Report Server, see Report Server Architecture: phoenixdb and reportdb.
3. Under **Tables**, select the parent table for your report.
   For the steps following, we will use the **Servers by CPU, Memory** table and its associated columns.
4. Drag the table to the **View** pane and click **Update Now**.
   The data in the table will load into the pane below. Note that the table columns match closely to the Report Display Columns in FortiSIEM.

5. For **Connection**, select **Live**.
6. Click **Go to Worksheet**.
   In the worksheet view you will see that a set of **Dimensions** and **Measures** are populated for the table.
7. Under **Measures**, select **Report Time** and drag it to the **Dimensions** section to create Report Time as a calculated measurement.
8. Under **Dimensions**, right-click on **Report Time** to edit the calculation formula and convert it to a human-readable format from UNIX time.
   The formula should look like `DATEADD('second',INT([Report Time]),#1969-12-31 16:00:00#)`
   You may also want to rename **Report Time** to **Time** to make it easier to read on the resulting chart.
9. Drag **Report Time** from **Dimensions** to **Columns**.
10. Under **Columns**, right-click on **Report Time** and select **Exact Date**.
    You should now see dates and time increments in your chart as the X-axis.
11. Under **Measures**, select and drag **AVG(cpuUtil)** and **AVG(memUtil)** to **Rows**.
12. Set the aggregation of both **AVG(cpuUtil)** and **AVG(memUtil)** to **AVG**.
    For example, **AVG(AVG(cpuUtil))** and **AVG(AVG(memUtil)).**
    You should now see both measures on the Y-axis of your chart.
13. Under **Dimensions**, drag **Host Name** to the **Color** section under **Marks**.
    Each host will be assigned a color and added to the chart.
14. Change the chart display name for AVG(cpuUtil) and AVG(memUtil) by clicking on each in the Y-axis to launch the Edit Y-Axis dialog.
    You can now edit the **Title** and **Range**, as well as other attributes, for each measure.
15. Under **Data**, click on the data source to open the Options menu, then click **Refresh**.
16. Rename the sheet by clicking on the data source to open the Options menu, then select **Rename** and enter a new name.

Your sheet is now complete. Hover your mouse over a trend line to view information about a specific host.

Create the Workbook

1. Click the **Dashboard** tab on the bottom of the Sheet editor to open the **Dashboard** editor.
2. Under **Dashboard**, select an appropriate **Size** and screen resolution.
3. Under **Dashboard**, select the sheet and drag it into the display pane.
4. Open the **Dashboard** options menu and select **Rename**.
   Change the name of the dashboard from **Server CPU/Memory Trend** to **Server Performance**.
5. In the **File** menu, select **Save**.

Publish the Workbook

1. In the **Server** menu, select **Sign In...**
2. Enter the IP address and port number for the Visual Analytics Server.
3. Enter the **Username** and **Password** for the Visual Analytics Server admin user, and then click **Sign In**.
4. In the **Server** menu, select **Publish Workbook**.
5. Enter attributes for the workbook, such the associated **Project**, **Name**, **View Permissions**, and **Views to Share**.
   See Adding Users to Workbooks for more information about user permissions for workbooks.
6. Click **Publish**.

Creating a Multiple Sheet Workbook

These instructions demonstrate how to create a multiple-sheet workbook that will contain a set of charts related to Network Health. This example uses the **Network Devices by Ping RTT**, **Network Interfaces By Utilization**, and **Network Devices By CPU, Memory** reports, but any report with an associated table and views in the reportdb database could be used. The Tableau Desktop online Help also contains extensive information about building sheets and workbooks with the Tableau Desktop editor, which powers the FortiSIEM Visual Analytics Desktop.

- Prerequisites
- Procedure

Prerequisites

- Follow the instructions in Syncing FortiSIEM Report with Report Server to sync the reports you want to use for your worksheet.
- You will need to know the name of the parent table for your synced reports. Follow the instructions in Viewing reportdb Organization to find the table that corresponds to your report.

Procedure

Create a View

Each report you want to include in your workbook corresponds to a table in the FortiSIEM reportdb. These tables need to be joined to cross-link the information that will appear in your workbook. In the case of a Network Health workbook that includes the sheets Network Devices by Ping RTT, Network Interfaces By Utilization, and Network Devices By CPU, Memory, the joining keys are **host name** and **time**.

1. Follow the instructions in Viewing reportdb Organization to find the parent tables for the reports you want to join. For each report there is one parent table and multiple child tables containing data for a particular month.
2. Create a SQL statement in pgAdmin to join the tables.
   In this example data is captured for one day. This enables quick generation of the data visualization.

```
 SELECT cpu.report_time, cpu."hostName", cpu."hostIpAddr", cpu."AVG
(cpuUtil)", cpu."AVG(memUtil)",
       uptime."SUM(sysDownTime)", uptime."AVG(avgDurationMSec)",
uptime."LAST(sysUpTime)",
       uptime."SUM(pollIntv)", util."intfName", util."intfAlias",
       util."AVG(inIntfUtil)" AS "totalAvgInIntfUtil", util."AVG
(outIntfUtil)" AS "totalAvgOutIntfUtil",
       util."AVG(recvBitsPerSec)" AS "totalAvgRecvBitsPerSec",
       util."AVG(sentBitsPerSec)" AS "totalAvgSentBitsPerSec",
       util."AVG(outQLen)", util."AVG(intfSpeed64)"
FROM "Network Devices By CPU, Memory_1278492569_1" cpu,
     "Network Devices by Ping RTT_2021056235_1" uptime,
     "Network Interfaces By Utilization_382117475_1" util

WHERE ((cpu.report_time * 1000)::double precision *
'00:00:00.001'::interval + '1969-12-31 16:00:00-08'::timestamp with time
zone) >= (now() - 1::double precision * '1 day'::interval)
```

```
     AND ((uptime.report_time * 1000)::double precision *
'00:00:00.001'::interval + '1969-12-31 16:00:00-08'::timestamp with time
zone) >= (now() - 1::double precision * '1 day'::interval)
     AND ((util.report_time * 1000)::double precision *
'00:00:00.001'::interval + '1969-12-31 16:00:00-08'::timestamp with time
zone) >= (now() - 1::double precision * '1 day'::interval)
     AND cpu.report_time = uptime.report_time AND cpu."hostName" =
uptime."hostName" AND uptime.report_time = util.report_time AND
uptime."hostName" = util."hostName";
```

3. Click the **Play** icon in pgAdmin to execute the query.
   Make sure the output pane contains data that is the result of the query execution.
4. Modify the SQL statement to create a view.
   Add this command at the top of the SQL statement:

   ```
    CREATE OR REPLACE VIEW ph_network_health_view AS
   ```

   Add this command at the bottom of the SQL statement:

   ```
    grant select on ph_network_health_view TO public;
   ```

   Your complete SQL statement should look like this:

   ```
   CREATE OR REPLACE VIEW ph_network_health_view AS

   SELECT cpu.report_time, cpu."hostName", cpu."hostIpAddr", cpu."AVG
   (cpuUtil)", cpu."AVG(memUtil)",
          uptime."SUM(sysDownTime)", uptime."AVG(avgDurationMSec)",
   uptime."LAST(sysUpTime)",
          uptime."SUM(pollIntv)", util."intfName", util."intfAlias",
          util."AVG(inIntfUtil)" AS "totalAvgInIntfUtil", util."AVG
   (outIntfUtil)" AS "totalAvgOutIntfUtil",
          util."AVG(recvBitsPerSec)" AS "totalAvgRecvBitsPerSec",
          util."AVG(sentBitsPerSec)" AS "totalAvgSentBitsPerSec",
          util."AVG(outQLen)", util."AVG(intfSpeed64)"
   FROM "Network Devices By CPU, Memory_1278492569_1" cpu,
        "Network Devices by Ping RTT_2021056235_1" uptime,
        "Network Interfaces By Utilization_382117475_1" util

   WHERE ((cpu.report_time * 1000)::double precision *
   '00:00:00.001'::interval + '1969-12-31 16:00:00-08'::timestamp with time
   zone) >= (now() - 1::double precision * '1 day'::interval)
        AND ((uptime.report_time * 1000)::double precision *
   '00:00:00.001'::interval + '1969-12-31 16:00:00-08'::timestamp with time
   zone) >= (now() - 1::double precision * '1 day'::interval)
        AND ((util.report_time * 1000)::double precision *
   '00:00:00.001'::interval + '1969-12-31 16:00:00-08'::timestamp with time
   zone) >= (now() - 1::double precision * '1 day'::interval)
        AND cpu.report_time = uptime.report_time AND cpu."hostName" =
   uptime."hostName" AND uptime.report_time = util.report_time AND
   uptime."hostName" = util."hostName";
   ```

```
grant select on ph_network_health_view TO public;
```

5. In pgAdmin, click the **Play** icon to execute the statement.
6. Using pgAdmin, navigate to the **Views** and make sure the `ph_network_health_view` has been created.
7. Right-click on `ph_network_health_view` to open the **Options** menu, then select **View Data > View Last 100 Rows** to make sure the view contains data.

Create a Workbook that Uses the View

1. Launch FortiSIEM Visual Analytics Desktop.
2. Connect to FortiSIEM Report Server with the **Username** and **Password** that you used during Report Server installation.
   For **Database**, enter **reportdb**.
   For **Port**, enter **30000**.

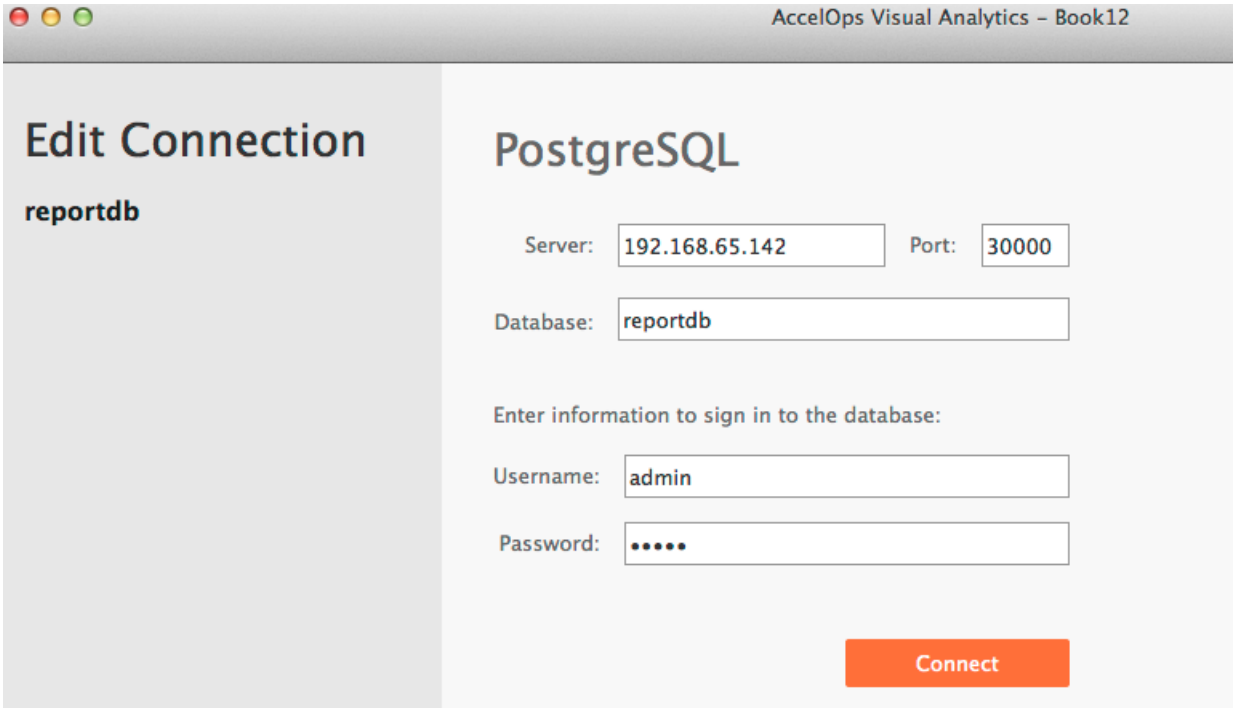

3. Under **Tables**, enter the name of the view you created in the search box to locate the view.

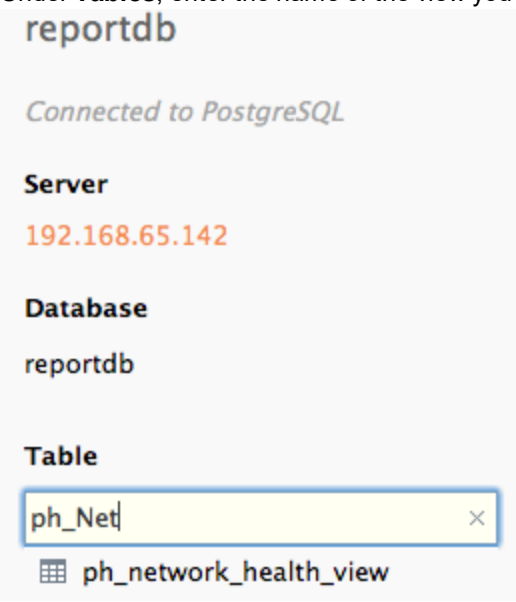

4. Drag the view into the **Join** pane and click **Update Now**.
   The data in the view will load into the pane below.

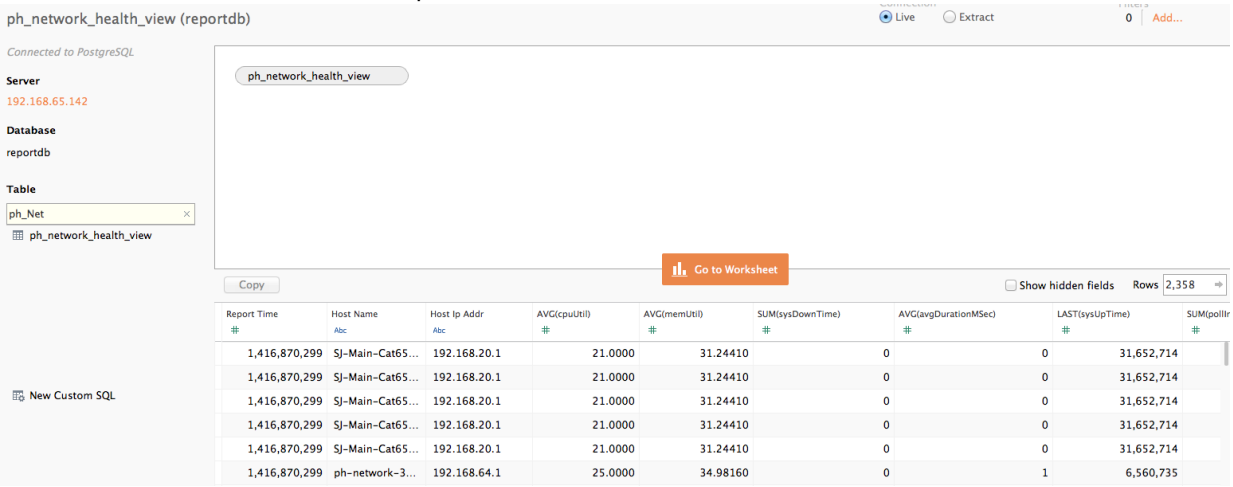

5. For **Connection**, select **Live**.
6. Click **Go to Worksheet**.
   In the worksheet view you will see that a set of **Dimensions** and **Measures** are populated for the view.
   An example worksheet showing CPU and Memory Utilization with several Dimensions and Measures populated from the original table.

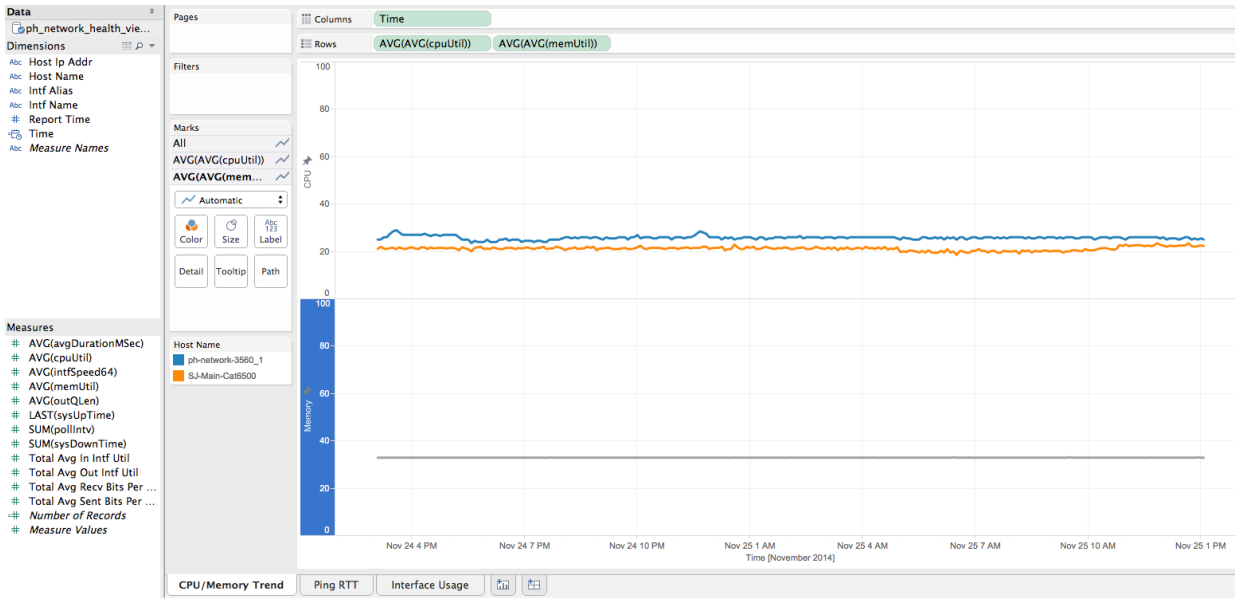

7. For each report in your workbook you can now create an individual sheet, as described in Creating a Single Sheet Workbook.

Create the Workbook

1. Click the **Dashboard** tab on the bottom of the Sheet editor to open the **Dashboard** editor.

2. Drag each sheet you've created into the **Join** pane.
   An example of three worksheets loaded into the Dashboard Join pane.

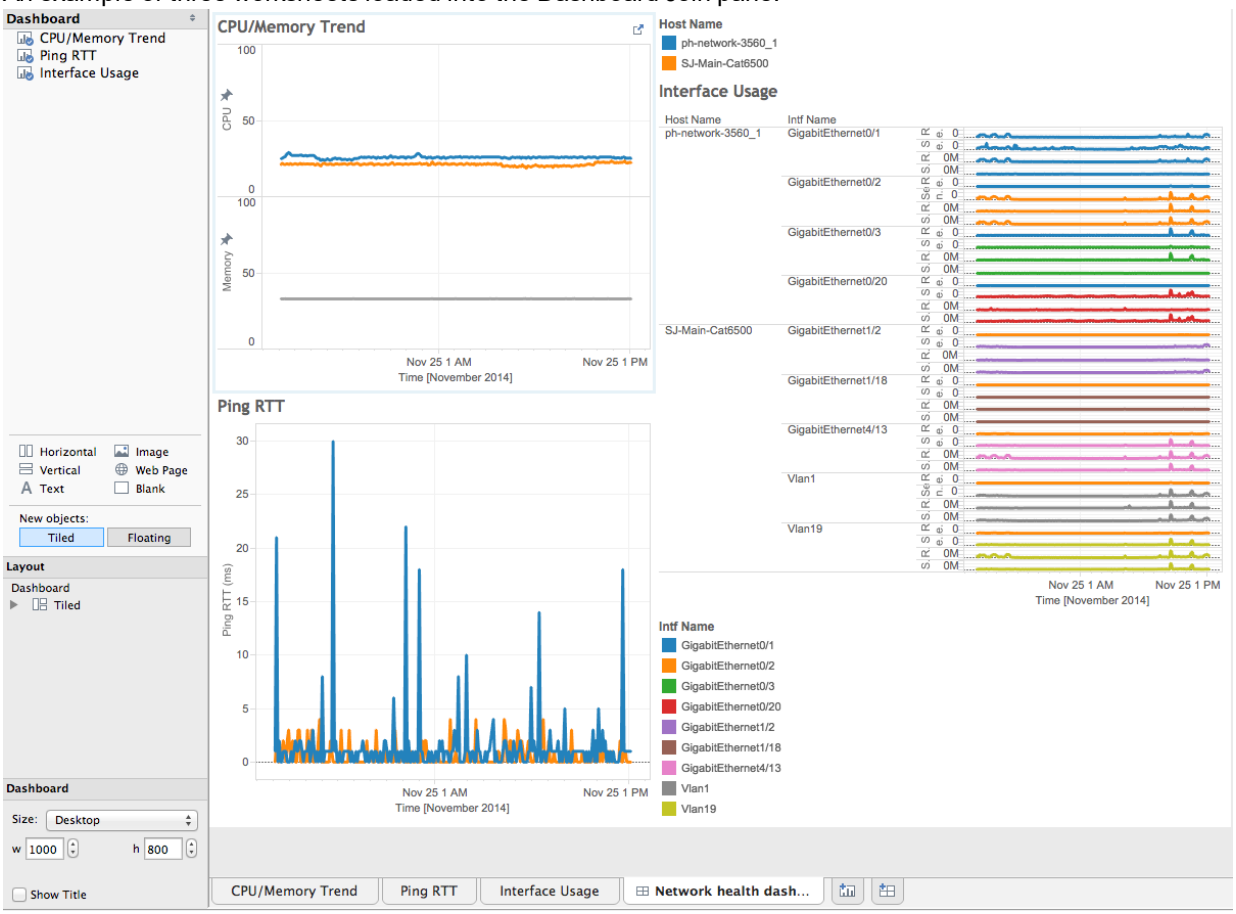

3. Under **Dashboard**, select an appropriate **Size** and screen resolution.
4. Open the Dashboard **Options** menu and select **Rename**.
5. In the **File** menu, select **Save** .

Publish the Workbook

1. In the **Server** menu, select **Sign In...**
2. Enter the IP address and port number for the Visual Analytics Server.
3. Enter the **Username** and **Password** for the Visual Analytics Server admin user, and then click **Sign In.**
4. In the **Server** menu, select **Publish Workbook**.
5. Enter attributes for the workbook, such the associated **Project**, **Name**, **View Permissions**, and **Views to Share**.
   See Adding Users to Workbooks for more information about user permissions for workbooks.

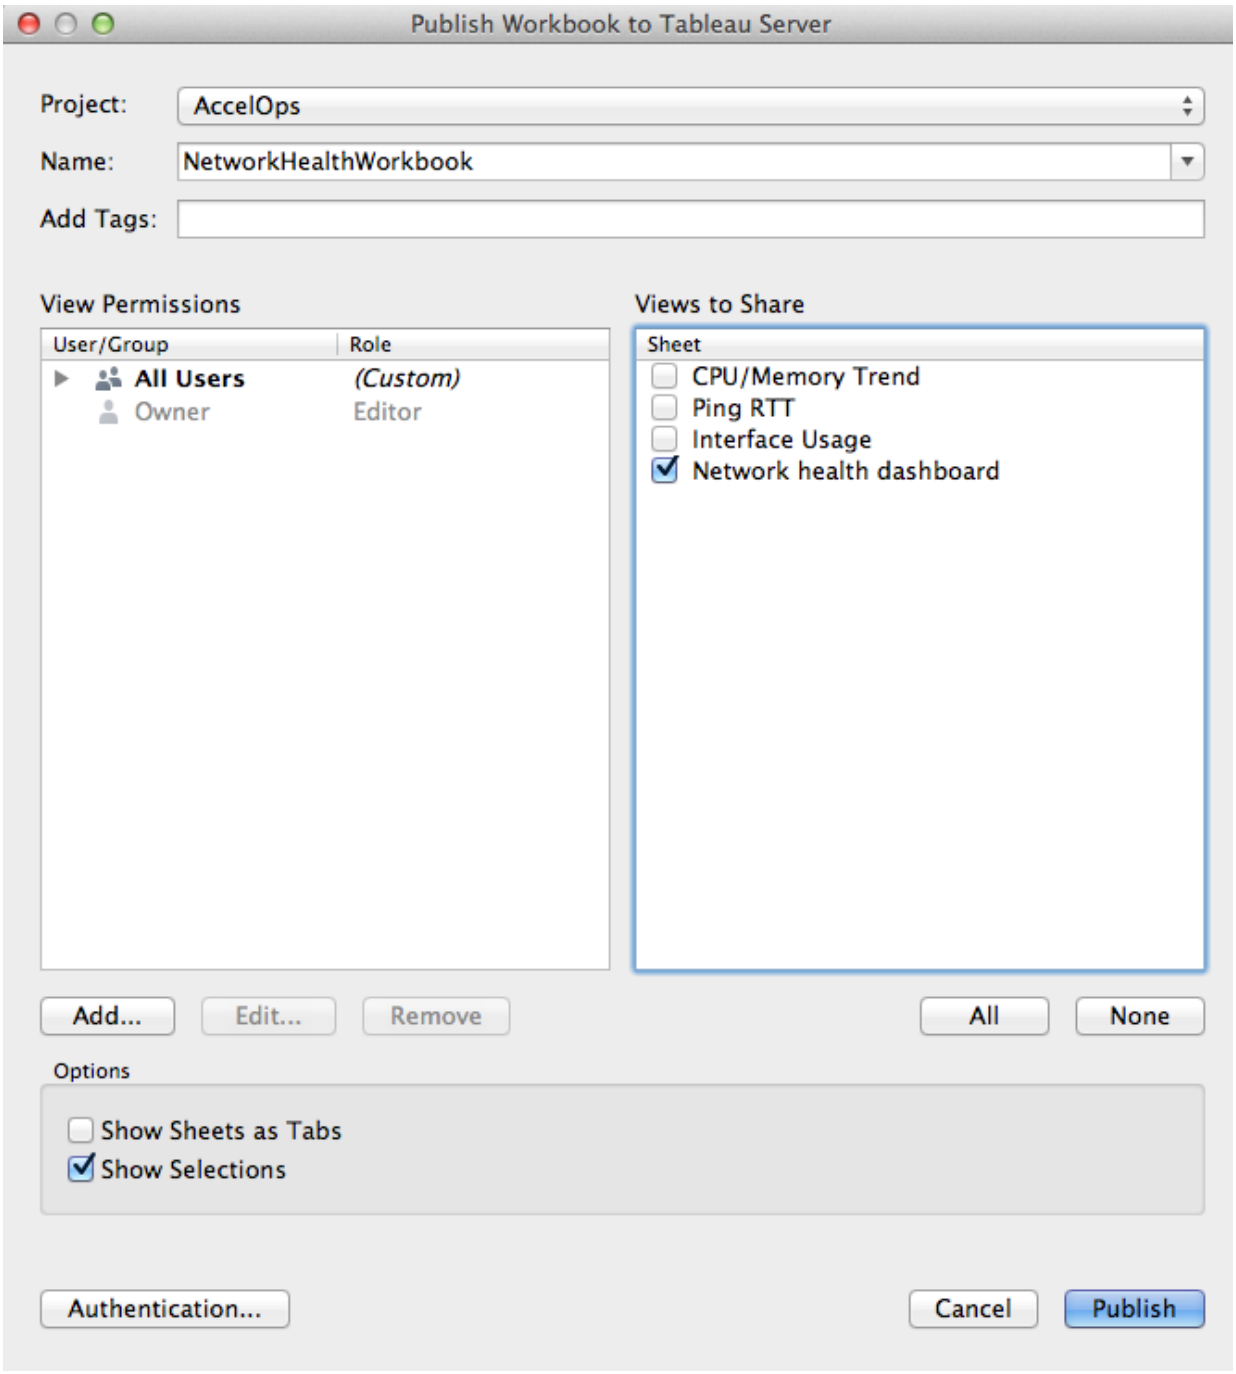

6. Click **Publish**.

Using FortiSIEM Workbooks with Tableau Visual Analytics Desktop and Server

You can use any of the workbooks provided by FortiSIEM, which are attached to this page, to create visualizations of FortiSIEM data.

1. Download a workbook attached to this page to your local device where Tableau Visual Analytics Desktop is installed.
2. In Visual Analytics Desktop, go to **File > Open...**.

3. Browse to the file you downloaded and open it.
4. You can make any changes you want to the workbook, but you can upload it to the server and start using it as is. Follow the instructions in the **Publish the Workbook** section of Creating a Single Sheet Workbook to publish to the Tableau Visual Analytics Sever, and add user permissions as described in Adding Users to Workbooks.

Adding Users to Workbooks

Only the workbook publisher can give access to specific users during report creation time. As the FortiSIEM Visual Analytics Server Administrator, you can add users to the system and view which workbooks users can access.

- Adding Users to Visual Analytics Server
- Viewing User Access to Workbooks

Adding Users to Visual Analytics Server

1. Log in to FortiSIEM Visual Analytics Server.
2. In the **Admin** tab click **Users**.
3. Click **Add**.
4. Enter the user name as it appears in Active Directory.
5. Select the **License Level** for the user and assign **User Rights** as necessary.
6. Click **OK**.

Viewing User Access to Workbooks

1. Log in to Visual Analytics Server.
2. In the **Admin** tab click **Users**.
3. Select a user name to see the workbooks that the user can access.

Real Time Performance Probe

This section describes how to probe monitored devices for real time performance metrics.

Available metrics

- CPU utilization
- Memory utilization
- Network interface statistics
- Uptime
- Disk utilization
- SNMP Ping Statistics
- Process Utilization

GUI launch locations

Real time Performance Metrics option is available from the following GUI locations

- CMDB > Device > IP Address > Right click
- CMDB > Device > Interfaces > Name > Right click
- Incident > Incident Source and Incident Target > Right click

Running a real time probe

1. From any of the above locations, select **Real Time Performance Metrics**
2. Select the parameters
    a. Select **Job Name** as the metric of interest
    b. Select polling **Frequency** in seconds
    c. Select the number of **Runs** as the number of times the device will be polled
    d. Select the **Collector** which should communicate to the device
    e. Depending on the job name, you may also need to select a Filter. For example, select Interface Name for Network Interface Statistics.
3. Click **Start**
4. The data will start to be displayed in the chart below
5. You can select two fields to be displayed side-by-side by
    a. selecting one attribute in the **Left Chart** drop down
    b. selecting another attribute in the **Right Chart** drop down and
    c. selecting Right Chart

The probe will stop after the device has been probed for the specified number of **Runs.**

6. To stop a probe, click **Probe.**

Implementation Note

FortiSIEM uses the same event framework to collect data from the devices and display them in the GUI. However these events are neither stored in the database, nor do they trigger incidents.

Example - Real time Interface Statistics Display

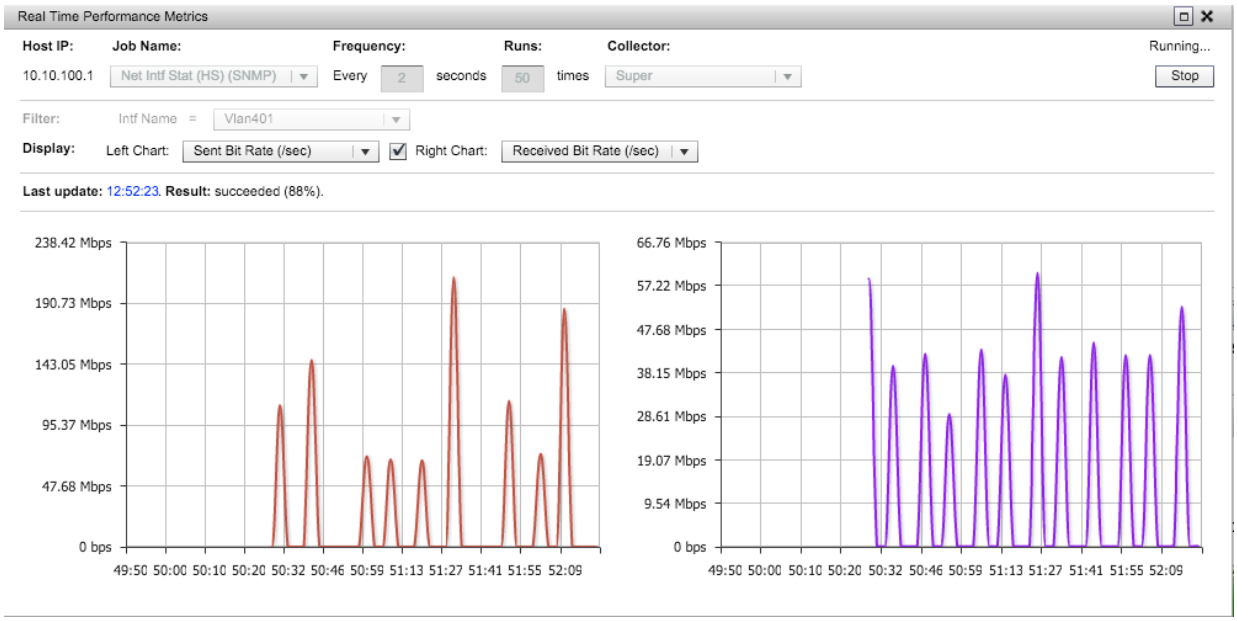

Incidents - Flash version

Incidents are a category of events that are are triggered when a set of rule conditions have been met. When an incident occurs, it appears in both the **Dashboard > Incident Dashboard** and in the **Incidents** tab, and, in some cases, a notification is sent based on the notification policies you have set. You can also create tickets from FortiSIEM incidents. Topics in this section cover the incident information that is available in the dashboard, how to create incident notification policies, how to create tickets for FortiSIEM and other ticket-handling systems, and how to manage the IPS Vulnerability Map.

- Incident Information
- The IPS Vulnerability Map
- Incident Notifications
- Creating Tickets
- Using Incidents in Searches and Rules

Viewing and Searching Incidents

The Incident Dashboard displays incident information for your IT infrastructure based on the filter conditions you set. You can also view incidents grouped by incident attributes, use values in incident attributes to refine your searches, view information about rules that triggered incidents, and use incident information to create rule exceptions and event dropping rules.

- List View of Incidents
- Device Risk View of Incidents
- Calendar View of Incidents
- Fishbone View of Incidents

List View of Incidents

There are two ways you can view the incidents that are occurring in your IT infrastructure.

- The **Incidents** tab, shown in the screenshot for this topic, where you can view incidents and incident details
- **Dashboard > Incident Dashboard**, which includes the same incident summary and user interface controls found in the Incidents tab, but which also provides other views of incidents, including a fishbone view of incidents in your infrastructure, a topology view with the number and severity of incidents overlaid on devices, a calendar view, and a location view that includes both a summary view of incident source and target IP locations and a map view, along with the number and severity of incidents for that location overlaid on the map.

In both locations you can filter the incidents in the dashboard, find out more information about sources and targets of incidents, customize the dashboard layout, and manage the rules associated with incidents.

- Incident Attributes
- Incident Dashboard User Interface Controls
- Incident Details

Incident Attributes

An Incident has the following attributes.

Attribute Name	Description
Event Severity Category	The severity of the incident, **High**, **Medium**, or **Low**
Last Seen Time	The last time that the incident was triggered
First Seen Time	The first time that the incident was triggered
Incident Name	The name of the rule that triggered the incident
Incident ID	The unique ID assigned to the incident
Incident Source	The source IP or host name that triggered the incident
Incident Target	The IP or host name where the incident occurred
Incident Detail	Event attributes that triggered the incident
Status	The status of the incident, **Active**, **Cleared**, **Cleared Manually**, **System Cleared**
Cleared Reason	For manually cleared incidents, this displays the reason the incident was cleared
Cleared Time	The time an incident was cleared
Cleared User	The person who cleared the incident
Comments	Any comments that users have entered for the incident
Ticket Status	Status of any tickets associated with the incident
Ticket ID	The ID number of any tickets generated by the incident
Ticket User	The person assigned to any tickets generated by the event
External User	If the ticket was cleared in an external ticket-handling system, this lists the name of the person the ticket was assigned to
External Cleared Time	If the ticket was cleared in an external ticket-handling system, this lists the time it was cleared.
External Resolved Time	If the ticket was resolved in an external ticket-handling system, this lists the time it was resolved.
External Ticket ID	The ID of the incident in an external ticket-handling system
External Ticket State	The state of the incident ticket in an external ticket-handling system

Attribute Name	Description
External Ticket Type	The type assigned to the incident ticket in an external ticket-handling system.
Organization	The organization reporting the event
Impacts	Organizations impacted by the event
Business Service	Business services impacted by the incident
Incident Notification Status	Status of any notifications that were sent because of the incident
Notification Recipients	Who received notification of the incident
Incident Count	How many times the incident has occurred during the selected time interval

Incident Dashboard User Interface Controls

This screenshot shows the Incidents tab with the major user interface controls outlined in red.

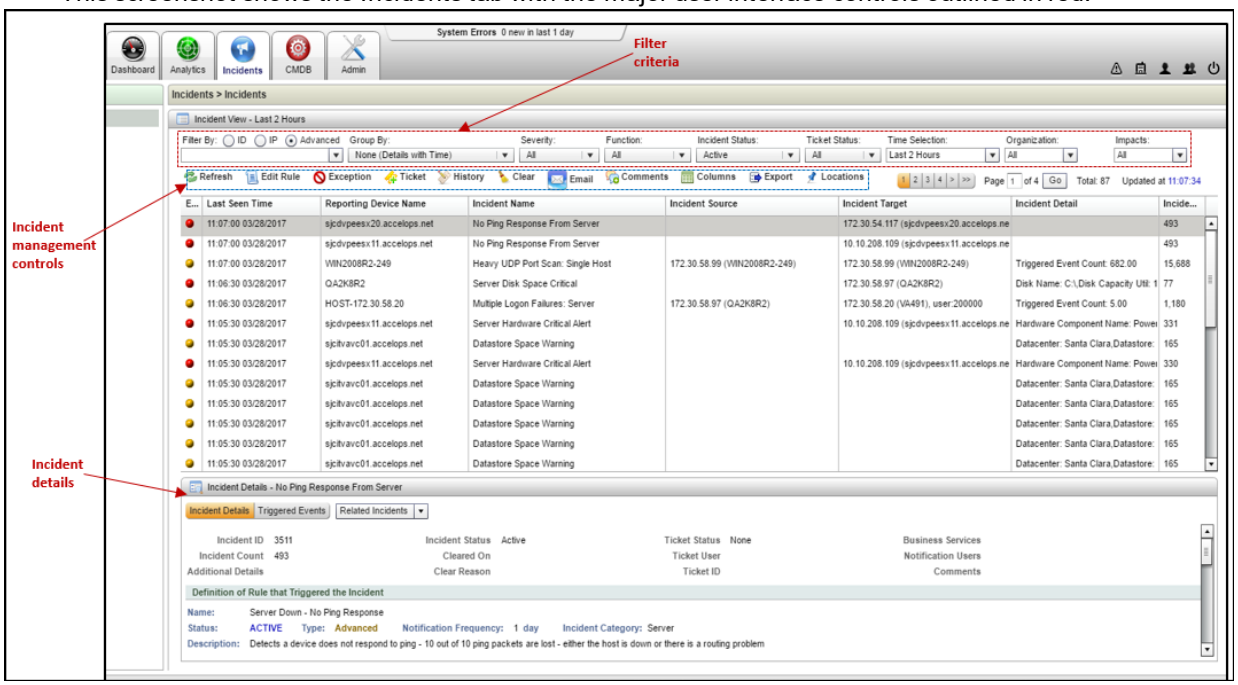

Incident Dashboard Filter Controls

The filter controls let you control which incidents are shown in the dashboard.

Filter Control	Description
Filter Criteria	ou have three options for the filter conditions: • **ID** Search for an incident by ID • **IP** Search for an incident based on an IP address • **Advanced** Use this option to set filter conditions based on event attributes as described in Creating a Structured Real Time Search. See Selecting Attributes for Structured Searches, Display Fields, and Rules for more information about using attributes in search filters.
Group By	Use these options to group incidents in the dashboards based on incident attributes. See Using Group By Attributes to View Incidents for more information.
Severity	Use these options to only see incidents with the selected severity level
Function	Use these options to view incidents related to a specific infrastructure functional area, such as **Performance** or **Security**.
Incident Status	Filter incidents to view according to their status
Ticket Status	Filter incidents based on the status of their associated tickets. See Creating Tickets In FortiSIEM In-built Ticketing System for more information.
Time Selection	Select the time interval during which incidents should have occurred. The default is **Last 2 Hours**.
Organization	For Service Provider deployments, select the organization you want to view incidents for.
Impacts	For multi-tenant deployments, select an organization to view the incidents that are impacting it

Incident Management Controls

Filter Control	Description
Refresh	Refresh the dashboard view
Edit Rule	Edit the rule associated with the incident. See the topics under Rules for more information.
Exception	Create an exception to the rule associated with the incident. See Defining Rule Exceptions for more information.
Ticket	Create a ticket from the incident. See Creating Tickets In FortiSIEM In-built Ticketing System for more information.
History	View the ticket history associated with an incident.
Clear	Clear the incident. See Defining Clear Conditions for more information on how to set rule conditions that will automatically clear incidents. All non-security related incidents are cleared from the system every night at midnight local time, and will show a status of **System Cleared**. A status of **Manual Clear** means that a user cleared the incident from the Incident Dashboard, while **Clear** means it was cleared by a rule condition.
Email	Select one or more incidents from the incident view and send email using any specific email template.
Comments	Add comments to the incident.
Columns	Change the columns displayed in the summary table. Incident Columns describes all the columns that can be added to the Incident Dashboard.
Export	Export the incident information to a PDF or CSV file
Locations	View geolocation information about the incidents. Pin colors on the map indicate incident severity: • Red: HIGH Severity • Yellow: MEDIUM Severity • Green: LOW Severity • Black: Incidents with multiple severity levels at the same location

Contextual Menus

Clicking on an item within a column of the incident summary will open a contextual menu, with options depending on whether the incident attribute you selected includes an IP address (**Source IP** or **Target IP**, for example), or some other kind of incident attribute. Shared between both menus are an **Add to Filter** option, which enables you to select a result attribute and add it to the **Filter By** conditions. Both menus also include most of the same options available in the Incident Management controls to edit and add exceptions to rules. The IP address

contextual menu provides options to view more information about the associated device, with many of the same options you would find in the Analysis menu used in search summary dashboards.

This screenshot shows the IP contextual menu open after selecting an IP address in the **Incident Source** column of the **Incidents** tab.

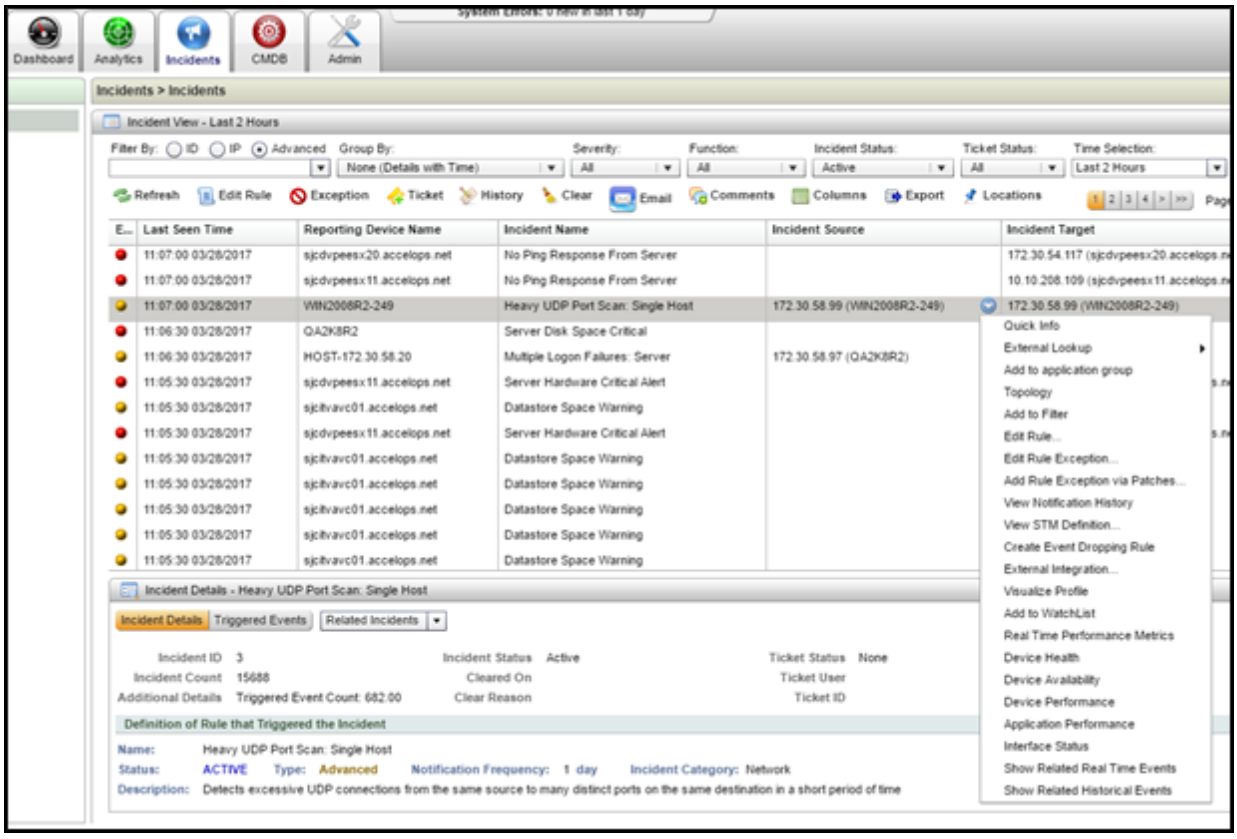

Incident Summary in the Dashboard > Incident Dashboard Contextual Menu

The **Dashboard > Incident Dashboard** contextual menu includes an option not found in the **Incident** tab view of incidents. Click in any column for an incident in the Incident Dashboard to open the contextual menu, and you will see the option **Show incident details**. If you select this option the **Incidents** tab view of incidents will load, and you will see detailed information about the incident you selected in the **Incident Details** pane.

Incident Details

The Incident Details pane at the bottom of the Incidents Dashboard provides you with information about a selected incident in three areas: **Incident Details**, **Triggered Events**, and **Related Incidents**.

Incident Details

The Incident Details include the ID of the incident, specific details about the event that triggered the incident, and the definition of the rule associated with the incident.

Triggered Events

The list of events that triggered the incident. For columns containing an event type, or host or IP information, click on an item to open a contextual menu and view more information about it.

Selecting Sub Patterns

If the rule that triggered the incident has multiple sub patterns, you can select a sub pattern to see which events met its conditions.

Related Incidents

Use this menu to view related incidents based on the **Source**, **Target**, **Rule Name**, or **Reporting IP** associated with the selected incident.

Searching for Incidents by Incident Attributes

As your review incidents in your dashboard, you may want to build searches based on attributes from selected incidents. For example, you may want to use the value for the **Incident Target** attribute in an incident as a filter condition to find similar or related incidents, and then add more conditions based on the results of that search.

1. Log in to your Supervisor node.
2. Go to **Incidents**.
3. In the Incident Dashboard, select an incident.
4. Click on the attribute value for the selected incident that you want to add to the **Filter By** condition to open the **Options** menu, and then select **Add to Filter**.
   The type of search will change to **Advanced**, and the attribute value you selected will be added to the Filter By conditions.
5. Click in the Filter By Conditions field to open the Conditions Builder and add other incident attributes.
6. Click **Refresh** when you're done creating filter conditions to see the results.

Using Group By Attributes to View Incidents

The Incident Dashboard presents a view of all incidents based on the filter conditions you select. However, there may be situations in which you want to view incidents grouped on incident attributes like **Incident Source**, **Incident Target**, **Severity**, or **Incident Name**. Once incidents are grouped by their attributes, you can view **Incident Details** for the entire group.

1. Log in to your Supervisor node.
2. Go to **Incidents**.
3. In the **Group By** menu, select the attributes you want to use to group the incidents, and then click **Refresh**.
   The Incident Dashboard will refresh and display incidents grouped according to the attributes you selected, with a **COUNT(Matched Events)** column that indicates how many incidents are in each group.
4. Select a group and then click on it to open the **Options** menu.
5. In the Options menu, select **Show Incident Details for This Group**.
   The Incident Dashboard will refresh to show all incidents in the selected incident group, and you can use the Contextual Menus to find out more information about them.

Device Risk View of Incidents

Viewing Devices Sorted By Risk:

1. Go to Incident tab
2. Set Group By to Host Risk Score.

3. Left pane shows Devices Sorted By Risk
4. Right pane shows incidents for the device selected in left panel

Calendar View of Incidents

The calendar view of incidents provides a summary view of the number of incidents that have occurred on a calendar day, grouped by severity. Clicking a group loads a summary of those incidents.

This screenshot shows the calendar view of incidents for the month of February 2015.

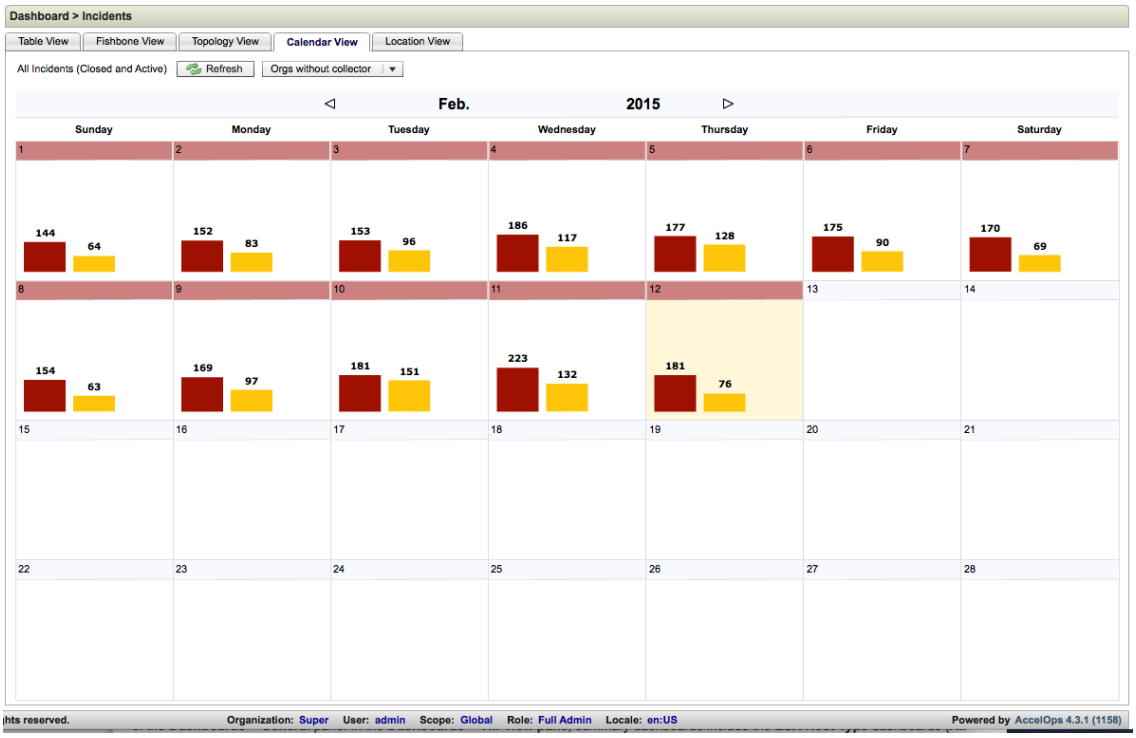

Fishbone View of Incidents

The fishbone view of incidents presents a view of networks and devices in those networks, along with the incidents triggered for those devices over the last 24. This view is derived from the Network Segments in the CMDB, with the devices associated with those segments overlaid. The numbers and colors for each device indicate the number and severity of incidents associated with that device.

- Clicking on an incident number will show you a summary of those incidents. Clicking on **Last Seen**, **First Seen**, **Incident Name**, or **Incident Details** in that summary will let you select **Incident Details** to view more information. Clicking on any IP addresses associated with the device will open a contextual menu that will let you find out more information about that device.
- Clicking on an IP number or hostname in the fishbone view will let you view the Quick Info for that device, or you can select Topology to view it within the context of your network topology.
- Hovering your mouse cursor over a device or incident number will show you the IP address and host name for that device, as well as the type of device.

This screenshot shows an example fishbone view of network segments, devices, and associated incidents.

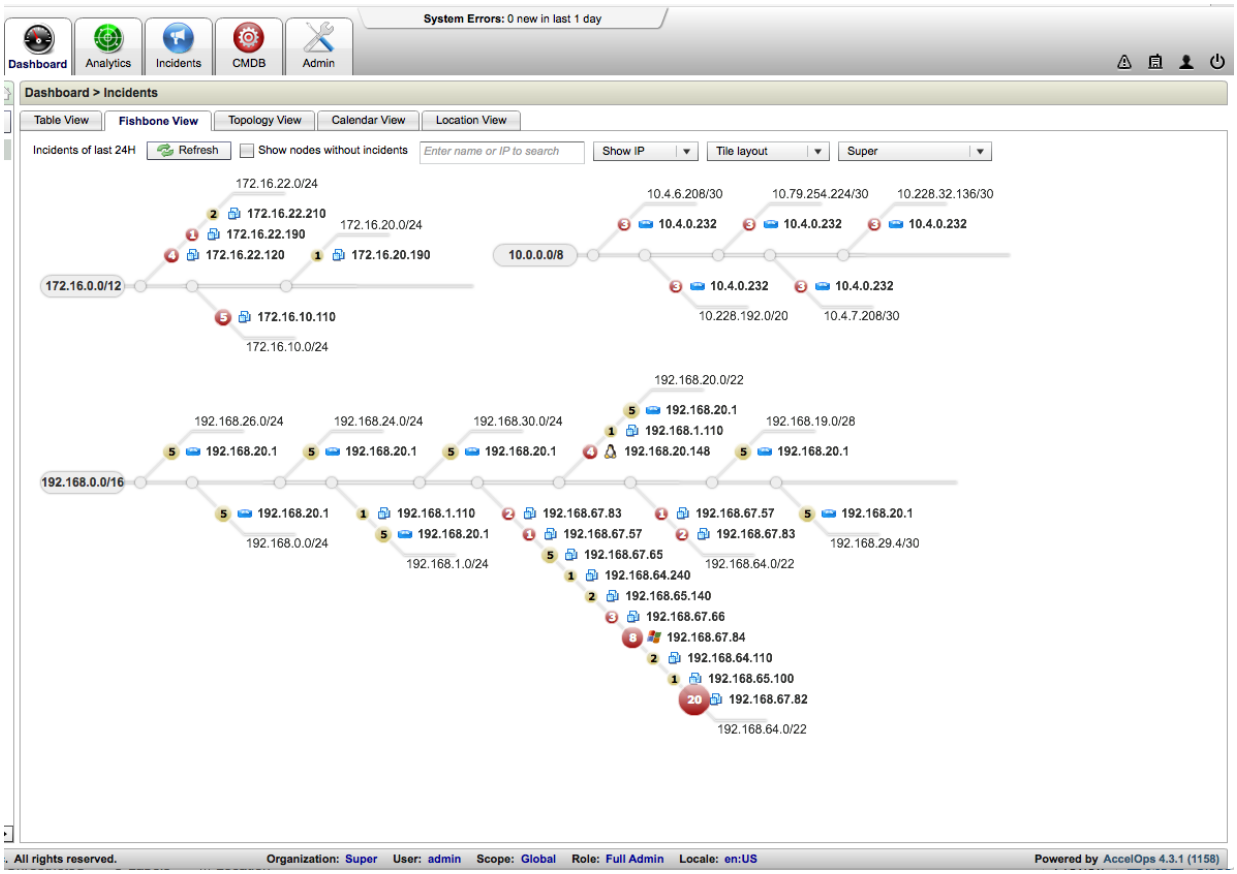
System Errors: 0 new in last 1 day
Dashboard
Analytics
Incidents
CMDB
Admin
Dashboard > Incidents
Table View
Fishbone View
Topology View
Calendar View
Location View
Incidents of last 24H
Refresh
Show nodes without incidents
Enter name or IP to search
Show IP
Tile layout
Super
172.16.22.0/24
172.16.22.210
172.16.22.190
172.16.22.120
172.16.20.0/24
172.16.20.190
172.16.0.0/12
172.16.10.110
172.16.10.0/24
10.4.6.208/30
10.79.254.224/30
10.228.32.136/30
10.4.0.232
10.0.0.0/8
10.228.192.0/20
10.4.7.208/30
192.168.20.0/22
192.168.20.1
192.168.1.110
192.168.20.148
192.168.19.0/28
192.168.26.0/24
192.168.24.0/24
192.168.30.0/24
192.168.0.0/16
192.168.0.0/24
192.168.1.0/24
192.168.67.83
192.168.67.57
192.168.67.65
192.168.64.240
192.168.65.140
192.168.67.66
192.168.67.84
192.168.64.110
192.168.65.100
192.168.67.82
192.168.64.0/22
192.168.29.4/30
All rights reserved.
Organization: Super
User: admin
Scope: Global
Role: Full Admin
Locale: en:US
Powered by AccelOps 4.3.1 (1158)

Incident Notifications

The sending of notifications when an incident occurs is handled by Notification Policies, which you can see listed in the **Analytics > Incident Notification Policies** page. Instead of having notifications set for each rule, you can create a policy and have it apply to multiple rules.

When viewing the notification policies, think of the columns on the page as representing a series of "If … and … then" statements that lead to the notification action. For example, you could read the table columns as a sentence:

"IF **Incident Severity** is X1 AND **Rule** is X2 AND **Time Range** is X3 AND **Affected Items** includes X4 AND **Affected Organizations** is X5, THEN take the actions specified in the **ACTION** column."

When FortiSIEM evaluates whether a notification action should be triggered based on the notification conditions, it evaluates all notification policies, and will trigger the actions of all policies that meet the condition, instead of just the first policy that meets the conditions. This means that the order of policies in the list doesn't matter, and that you can write policies with overlapping conditions that could also, for example, include different actions.

See also the topics under Incident Notification for more information about the methods that are available for sending notifications from FortiSIEM, including the FortiSIEM API.

- Creating an Incident Notification Policy
- Sending Email and SMS Notifications for Incidents
- Setting Scripts as Notification Actions
- Viewing Incident Notification History

Creating an Incident Notification Policy

Prerequisites

- Make sure you have enabled the settings for sending email or other notification actions as described in Setting Up Routing Information for Reports and Incident Notifications.
- You should read the introductory topic on incident notifications to understand how policy conditions are processed..

Procedure

1. Log in to your Supervisor node.
2. Go to **Analytics > Incident Notification Policy**.
3. Click **New**.
4. Select the **Incident Severity**.
   Only incidents matching the severity level you select will trigger a notification.
5. For **Rules**, click **...** and select the rule or rules you want to trigger this notification.
6. Set a **Time Range** during which this notification will be in effect.
   Notifications will be sent only if an incident occurs during the time range you set here.
7. For **Affected Items**, click **...** and use the CMDB Browser to select the devices or applications for which this policy should apply.
   Instead of individual devices or groups, you can apply the notification policy to an IP address or range by clicking **Add** under **IP/Range**. You can also select a group, and then select the **Not** option to explicitly exclude that group of applications or devices from the notification policy.
8. For Service Provider deployments, select the **Organizations** to which the notification policy should apply.
   Notifications will be sent only if the triggering incidents affect the selected organization.
9. Select the **Actions** to take when the notification is triggered.
   See the topics under Sending Email and SMS Notifications for Incidents, Creating Tickets, Creating Inbound Policies for Updating Ticket Status from External Ticketing Systems, and Setting Scripts as Notification Actions for more information about notification actions.
10. Enter any **Comments** about the policy.
11. When you are finished creating the notification policy, select **Enabled** to make it active in your deployment.
12. Click **Save**.

Sending Email and SMS Notifications for Incidents

When you set actions for an incident notification, one option is to send an email or SMS message to groups or individuals, and you also have an option to specify a template that should be used in the email.

- Prerequisites
- Procedure
- Related Links

Prerequisites

- Make sure the email gateway has been configured for your deployment.
- You should also have set up any email templates that you want to use for notifications.

Procedure

1. Log in to your Supervisor node.
2. Go to **Analytics > Incident Notification Policy**.
3. Select the policy that you want to set up the email or SMS notification for.
4. Under **Actions**, next to the email/sms notification table, click **...** .
5. For multi-tenant deployments, select the **Organization** that contains the individuals or groups you want notified. Under **Folders**, you will see the user groups for that organization listed.
6. In the **Folders** pane, select a group**.**
   In the **Items** pane, you will see a list of users for that group.
7. Select a group and click **Folder >>** to add a group to the **Notification Actions** list, or select individual users and click **Items >>**.
   Adding Individual Email Addresses

   If you want to set up notifications for individual email addresses without selecting from a user group, click **Add** under **Email Addr** and enter the addresses separated by commas.
8. Under **Notification** Actions, select the **Method**, Email or SMS, that you want to use sending the notification.
9. Select an **Email Template** if you are sending an email notification.
   If you leave this blank, the default email template will be used.
10. Click **Save**.
    Run On

    The **Run On** column only applies if your notification action is to execute a script. For email and SMS notification actions, it will be auto-populated with **Super**.

You can send incident notification emails for multiple incidents based on customer requirements by selecting multiple incidents and clicking **Email** button.

Related Links

- Setting Up the Email Gateway
- Setting Scripts as Notification Actions

Customizing Email Templates for Notifications

Email templates for incident notifications are based on incident variables that you put into the subject and body of the template, which are then populated with the actual attribute values in the incident.

Incident Attribute Variables

These are the incident attribute variables you can use for your email template.

- $organization
- $status
- $hostName
- $incidentId
- $incidentTime
- $firstSeenTime
- $lastSeenTime
- $incident_severityCat
- $incident_severity
- $incident_incidentCount
- $ruleName
- $ruleDescription
- $incident_source
- $incident_target
- $incident_detail
- $affectedBizService

Example Email Template

This example first shows a template with the incident attribute variables, and then an email based on this template with the variables populated from an incident.

Template

Email Subject:
$ruleName was triggered at $incidentTime

Email Body:
The host, $incident_target, was being scanned by $incident_source starting at $firstSeenTime and ending at $lastSeenTime. There were $incident_incidentCount hits.

Please investigate and report as necessary.

Generated Email

Subject: Server Memory Warning was triggered at Jan 10 22:43 UTC

Body: The host, Host IP: 192.168.1.23 Host Name: QA-V-WIN03-ORCL, was being scanned by 10.1.1.1 starting at Jan 10 22:05 UTC and ending at Jan 10 22:11 UTC. There were 2 hits.

Please investigate and report as necessary.

Creating an Email Template

1. Log in to your Supervisor node.
2. Go to **Admin > General Settings > Incident Email Templates**.
3. Click **Add**.
4. For Service Provider deployments, select the organization for which you are creating the email template.
5. Enter a **Name** for the template.
6. Enter the **Email Subject** and **Email Body**.
   You can select attribute variables from the **Insert Content** menu to enter into your template, rather than having to type them out by hand.
7. Click **OK**.

Setting a Default Template

When you are creating a notification policy and need to select an email template, if you leave the option blank, the default template will be used. To set an email template as default, select the template in the list on the **Incident Email Templates page,** and then click **Set as Default**. For Service Provider deployments, to select a template as default for an organization, first select the organization, then set the default email template for that organization.

Setting Scripts as Notification Actions

One of the actions you can specify for an incident notification is to execute a script. For example, suppose you are monitoring Windows services that are in Auto mode, and you have rules that will trigger an incident if one of those services is stopped. The notification action for that incident can include the running of a script by FortiSIEM that will re-start the service, as shown in the example scripts in this topic.

How Script Notification Actions are Processed

1. When a notification policy is triggered, the policy actions are handled in sequential order. That means, if there are multiple script actions, the first one will be processed before the second.
2. When you specify the notification action as a script, you must provide the full path to the script in the notification policy settings, for example `/tmp/Myscript.py`
3. When the script action is processed,
   a. FortiSIEM notification module will first generate an incident XML file and put it in the same directory as the script.
   b. FortiSIEM will then call the script with the XML file name as an argument. The full incident XML file path will be passed as the first command line parameter to the script, e.g. $1 for shell and sys.argv[1] for Python.
4. You must write the script so it expects the incident XML file to be located in the same directory as the script, for example /tmp if the script location is /tmp/Myscript.py. Use absolute path to refer to the incident XML file.
5. When the script returns, the incident XML file that was created by FortiSIEM is deleted, so there is no confusion with the next script action which involves a new incident XML file and is processed only after the previous script action is complete.

See here for an example of a notification script.

Setting a Script Notification Action

1. Log in to your Supervisor node.
2. Go to **Analytics > Incident Notification Policy**.
3. Select the notification policy where you want to add the script action.
4. Under **Actions**, next to the **Methods** table, click **...** .
5. Under **Run Script**, click **Add**.
6. For **Script Name**, enter the name of the script and the absolute directory path to it.
7. Click **OK**.
   The script will be added to the **Notification Actions**.

Selecting the Collector Where the Script Will Run

If your deployment includes Collectors, you can specify the Collector where the script will run.

1. Prepare the script on the Collector(s) and make sure they run properly.
2. When incident triggers, Collector will download the incident XML from the Supervisor and run the script with incident XML as argument.
3. The Collector will then return the results to the Supervisor.

In the **Notification Actions** table, select the Collector from the list in the **Run On** menu after you have added the script to the notification actions.

Example of a Windows Restart Script as a Notification Action

This topic provides an example of a script that could be used as a notification action, following the example of re-starting a Windows service that has stopped an triggered an incident as described in Setting Scripts as Notification Actions.

This example requires two scripts: one located on the Windows server that hosts the service, and a script on the FortiSIEM Supervisor host machine that will be triggered by the incident notification and will execute the Windows server script.

Windows Script

1. Create a script named `installWinexeSvc.bat` for starting the remote `winexe` provider service.

```
sc create AoWinexeSvc binPath= C:\WINDOWS\WINEXESVC.EXE start= auto
DisplayName= AoWinexeSvc
sc description AoWinexeSvc "Remote command provider for FortiSIEM
monitoring"sc start AoWinexeSvc
```

2. Run `installWinexeSvc.bat` on the monitored Windows server and make sure that the `AoWinexeSvc` service starts.

```
C:\>installWinexeSvc.bat
```

You should see an output similar to this as Windows installs the service and verifies that it is running.

```
C:\>sc create AoWinexeSvc binPath= C:\WINDOWS\WINEXESVC.EXE start= auto
DisplayName= AoWinexeSvc
[SC] CreateService SUCCESS
C:\>sc description AoWinexeSvc "Remote command provider for FortiSIEM
monitoring"[SC] ChangeServiceConfig2 SUCCESS
C:\>sc start AoWinexeSvc

SERVICE_NAME: AoWinexeSvc
        TYPE               : 10  WIN32_OWN_PROCESS
        STATE              : 2  START_PENDING
                                (NOT_STOPPABLE, NOT_PAUSABLE, IGNORES_
SHUTDOWN)
        WIN32_EXIT_CODE    : 0  (0x0)
        SERVICE_EXIT_CODE  : 0  (0x0)
        CHECKPOINT         : 0x0
        WAIT_HINT          : 0x7d0
        PID                : 1580
        FLAGS              :
```

FortiSIEM Script

This script, `restartWinService.py`, reads the incident XML file, parses out the target IP and stopped service, and issues a `winexe` command to restart the service.

```
 #!/usr/bin/python
importos, re, sys, time
importxml.dom.minidom
```

```
iflen(sys.argv) != 2:
   print "Usage: parseTargetIP.py incident.xml"   exit()
else:
   fileName = sys.argv[1]
print "parsing incident xml file : ", fileName
#os.system("cp "+ fileName + " "+ fileName + ".txt")
# /incident/incidentTarget/entry[@attribute='hostIpAddr']
doc = xml.dom.minidom.parse(fileName)
nodes = doc.getElementsByTagName('incidentTarget')
ifnodes.length < 1:
   print "no incident Target found!"else:
   targeNode = nodes[0]
targetIP = ""fornode in targeNode.childNodes :
   ifnode.nodeType == node.ELEMENT_NODE:
      ifnode.getAttribute("attribute") == "hostIpAddr":
         targetIP = node.firstChild.data
iftargetIP == "":
   print "no incident target found!"# trim IP, e.g. 10.1.20.189(SH-Quidway-SW1)
targetIP = re.sub(r'\(.+\)', "", targetIP)
print "restart service for target IP: ", targetIP
# parse process name
nodes = doc.getElementsByTagName('incidentDetails')
ifnodes.length < 1:
   print "no incidentDetails found!"else:
   targeNode = nodes[0]
fornode in targeNode.childNodes :
   ifnode.nodeType == node.ELEMENT_NODE:
      ifnode.getAttribute("attribute") == "serviceName":
         targetService = node.firstChild.data

################################################################################-
###############
# NOTE:  You need to replace the user and password with an account on your Windows
server that #
#        has permissions to run thiswindows command.
          #

################################################################################-
###############
# stop the service
stopCmd = "winexe --user Administrator --password ProspectHills! //"+ targetIP + "
'sc stop "+ targetService + "'"ret = os.system(stopCmd)
print "stop service with return code ,", ret
print "waiting service stop"time.sleep(10)

################################################################################-
###############
# NOTE:  You need to replace the user and password with an account on your Windows
server that #
#        has permissions to run thiswindows command.
          #
```

```
#############################################################################-
###############
## start the service
startCmd = "winexe --user Administrator --password ProspectHills! //"+ targetIP +
" 'sc start "+ targetService + "'"print "start command: ", startCmd
ret = os.system(startCmd)
print "start service with return code ,", ret
```

Incident XML File Format

This topic includes an example of the XML file that is generated for incidents, and descriptions of its contents.

Example Incident XML File

```
 <?xml version="1.0" encoding="UTF-8" ?><incident incidentId="5672" ruleType="PH_
RULE_AUTO_SRVC_DOWN" severity="10" repeatCount="1" organization="Super" status-
s="Cleared">  <name>Auto Service Stopped</name>  <description>Detects that an auto-
matically running service stopped. Currently this works for windows servers and is
detected via WMI.</description>  <displayTime>Fri Jun 29 15:51:10 PDT 2012</dis-
playTime>  <incidentSource>  </incidentSource>  <incidentTarget>    <entry attrib-
ute="hostIpAddr" name="Host IP">172.16.10.15</entry>    <entry
attribute="hostName" name="Host Name">QA-V-WIN03-ADS</entry>  </incidentTarget>
<incidentDetails>    <entry attribute="serviceName" name="OS Service Name">S-
pooler</entry>    <entry attribute="servicePath" name="OS Service
Path">C:\WINDOWS\system32\spoolsv.exe</entry>  </incidentDetails>  <affectedB-
izSrvc>Auth Service</affectedBizSrvc>  <identityLocation>  </identityLocation>
<rawEvents>    [SrvcDown]
        [PH_DEV_MON_AUTO_SVC_START_TO_STOP]:[eventSeverity]=PHL_INFO,[fileName]-
]=phPerfJob.cpp,[lineNumber]=6005,[hostName]=QA-V-WIN03-ADS,[hostIpAd-
dr]=172.16.10.15,[serviceName]=Spooler,
[servicePath]=C:\WINDOWS\system32\spoolsv.exe,[serviceDesc]=Manages all local and
network print queues and controls all printing jobs. If this service is stopped,
printing on the local machine will be unavailable. If this service is disabled,
any services that explicitly depend on it will fail to start.,[phLogDetail]=
 </rawEvents>
</incident>
```

XML Tag and Attribute Definitions

XML Tag	Attributes	Description
<incident>		
	incidentID	Unique id of the incident in FortiSIEM. You can search for the incident by using this ID.
	ruleType	Unique id of the rule in FortiSIEM
	severity	The severity of the incident, HIGH MEDIUM LOW
	repeatCount	How many times this incident has occurred
	organization	In Service Provider deployments, the organization affected by the incident
	status	The status of the incident
<name>		The name of the rule that triggered the incident
<description>		The description of the rule that triggered the incident
<displayTime>		The time when the incident occurred
<incidentSource>		The source of the incident. It includes the event attributes associated with the source presented as name:value pairs. Common attributes for source and target tributes here are srcIpAddr , destIpAddr , hostIpAddr .
<incidentTarget>		Where the incident occurred, or the target of an IPS alert. It includes the event attributes associated with the target presented as name:value pairs. Common attributes for source and target tributes here are srcIpAddr , destIpAddr , hostIpAddr .
<incidentDetails>		The event attributes associated with the rule definition that triggered the incident
<affectedBizSrvc>		Any business services impacted by the event
<identityLocation>		Information associated with the Identity and Location Report
<rawevents>		The contents of the raw event log for the incident.

Viewing Incident Notification History

There are two ways you can view the notification history for an incident.

1. In the **Incident Notification Status** column of the **Incident Dashboard**.
2. Click on an incident in the **Incident Name** column of the Incident Dashboard, and then select **View Notification History** from the **Options** menu.

Creating Tickets In FortiSIEM In-built Ticketing System

FortiSIEM includes a feature that will let you create and assign tickets for IT infrastructure tasks, and create tickets directly from incidents. You can see all tickets that have been created by going to Incidents > Tickets, and then use the filter controls to view tickets by assignee, organization, priority, and other attributes. You can also configure FortiSIEM and you Remedy system so that Remedy will take tickets created by incident notification actions.

- Configuring Remedy to Accept Tickets from FortiSIEM Incident Notifications
- Ticket Related Operations

Configuring Remedy to Accept Tickets from FortiSIEM Incident Notifications

This topic describes how to configure Remedy to accept tickets as notification actions from FortiSIEM.

Prerequisites

- Make sure you have configured the Remedy server settings in FortiSIEM.

Procedure

1. In Remedy, create a new form, **FortiSIEM_Incident_Interface**, with the incident attributes listed in the table at the end of this topic as the form fields.
2. When you have defined the fields in the form, right-click on the field and select the **Data Type** that corresponds to the incident attribute.
3. After setting the form field data type, click in the form field again to set the **Label** for the field.
4. When you are done creating the form, go to **Servers > localhost > Web Service** in Remedy, and select **New Web Service**.
5. For **Base Form**, enter **FortiSIEM_Incident_Interface**.
6. Click the **WSDL** tab.
7. For the **WSDL Handler URL**, enter `http://<midtier_server>/arsys/WSDL/public/<servername>/FortiSIEM_Incident_Interface`.
8. Click the **Permissions** tab and select **Public.**
9. Click **Save.**

You can test the configuration by opening a browser window and entering the WSDL handler URL from step 7, substituting the Remedy Server IP address for `<midtier_server>` and `localhost` for `<servername>`. If you see an XML page, your configuration was successful.

Incident Attributes for Defining Remedy Forms

Incident Attribute	Data Type	Description
biz_service	text	Name of the business services affected by this incident
cleared_events	text	
cleared_reason	text	The reason for clearing the incident if it was cleared,
cleared_time	bigint	The time at which the incident was cleared
cleared_user	character varying (255)	The user who cleared the incident
comments	text	Comments
cust_org_id	bigint	The organization id to which the incident belongs
first_seen_time	bigint	Time when the incident occurred for the first time
last_seen_time	bigint	Time when the incident occurred for the last time
incident_count	integer	Number of times the incident triggered between the first and last seen times
incident_detail	text	Incident Detail attributes that are not included in incident_ src and incident_target
incident_et	text	Incident Event type
incident_id	bigint	Incident Id
incident_src	text	Incident Source
incident_status	integer	Incident Status
incident_target	text	Incident Target
notif_recipients	text	Incident Notification recipients
notification_action_ status	text	
orig_device_ip	text	
ph_incident_category	character varying (255)	FortiSIEM defined category to which the incident belongs: Network, Application, Server, Storage, Environmental, Virtualization, Internal, Other
rule_id	bigint	Rule id

Incident Attribute	Data Type	Description
severity	integer	Incident Severity 0 (lowest) - 10 (highest)
severity_cat	character varying (255)	LOW (0-4), MEDIUM (5-8), HIGH (9-10)
ticket_id	character varying (2048)	Id of the ticket created in FortiSIEM
ticket_status	integer	Status of ticket created in FortiSIEM
ticket_user	character varying (1024)	Name of the user to which the ticket is assigned to in FortiSIEM
view_status	integer	
view_users	text	

Ticket Related Operations

Creating a ticket without an Incident

1. Go to **Incidents > Tickets**.
2. Click **New**.
3. Enter a **Summary** and **Description** for the ticket.
   Both of these fields are required.
4. For **Assigned To**, select a user from the menu.
5. Set any **Due Date** for the ticket.
6. Select a **Priority** for the ticket.
7. Click **Save**.

Creating a ticket from an Incident

1. In the Incident Dashboard, select the incident you want to create a ticket for.
2. Click **Ticket**.
   The **Incident ID**, **Summary** and **Description** for the ticket will be populated from the incident information.
3. Select the person you want to assign the ticket to.
4. Enter a **Due Date** for the ticket.
5. Set a **Priority** for the ticket.
6. Click **Save**.

Closing a ticket

1. Go to **Incidents > Tickets**.
2. Select a ticket
3. Click **Edit**

4. For **State** drop down, select **Closed**
5. **Click** Save.

Changing the assignee in a ticket

1. Go to **Incidents > Tickets**.
2. Select a ticket
3. Click **Edit**
4. For **Assigned** drop down, select the new Assignee
5. **Click** Save.

Changing the due date in a ticket

1. Go to **Incidents > Tickets**.
2. Select a ticket
3. Click **Edit**
4. For **Due Date** edit box, select the **date** and then the **time**
5. **Click** Save.

Adding notes to a ticket

1. Go to **Incidents > Tickets**.
2. Select a ticket
3. Click **Edit**
4. Add to **Description**
5. **Click** Save

Adding attachments to a ticket

1. Go to **Incidents > Tickets**.
2. Select a ticket
3. Click **Edit**
4. Click **PDF or PNG** under **Attach file**
5. Include the file and Click **Upload**.
6. **Click** Save

Exporting a ticket

1. Go to **Incidents > Tickets**.
2. Select a ticket
3. Click **Export**

Viewing Ticket History

1. Go to **Incidents > Tickets**.
2. Select a ticket

3. Click **Edit**
4. See **Action History** on bottom right pane

Searching tickets

This can be done in two ways:

- Type in key words in Search box - 

Summary column will be searched.
- Type in key words in Search box - 

Summary column will be searched.

Creating Tickets in External Ticketing System

See External Helpdesk System Integration.

Using Incidents in Searches and Rules

Creating an Historical Search from an Incident

When you are viewing an incident, you may want to about other events related to the source or target of the incident. This topic describes how to create an historical search from an incident.

1. In the Incident Dashboard, select the incident you want to use.
2. Select the Incident Source or Incident Target you want to use, and then select **Show Related Historical Events**.
   The Historical Search interface will load, with the IP address of the selected incident attribute loaded in the **Filter By** conditions, and the **Display Fields** set to the incident attributes.
3. Click **Run**.
4. You will see a list of events for the Incident Source or Target, which you can further analyze as described in Refining the Results from Historical Search.

Creating a Real Time Search from an Incident

When you are viewing an incident, you may want to about other events related to the source or target of the incident. This topic describes how to create a real time search from an incident.

1. In the Incident Dashboard, select the incident you want to use.
2. Select the **Incident Source** or **Incident Target** you want to use, and then select **Show Related Real Time Events**.
   The real time search interface will load, with the IP address of the selected incident attribute loaded in the **Filter By** conditions, and the **Display Fields** set to the incident attributes.
3. Click **Run**.
4. You will see a list of events for the Incident Source or Target, which you can further analyze as described in Viewing and Refining Real Time Search Results.

Editing Rules from Incidents

If you need to edit the rule associated with an incident, you can do so directly from the Incident Dashboard.

1. In the Incident Dashboard, select an incident based on the rule you want to edit.
2. Click in any column of the selected incident to open the **Options** menu, and then select **Edit Rule**.
3. Edit the rule as necessary, and then click **Save**.

Incidents - HTML5 version

Incident tab allows users to view and manage incidents.

Incident Attributes

This topic describes all the columns that can be used to create views in the Incident Dashboard. You can add or remove columns from the dashboard by clicking the **Columns** icon.

Column Name	Description
Severity	The severity of the incident, **High**, **Medium**, or **Low**
Last Occurred	The last time that the incident was triggered
First Occurred	The first time that the incident was triggered
Incident	The name of the rule that triggered the incident
Incident ID	The unique ID assigned to the incident
Source	The source IP or host name that triggered the incident
Target	The IP or host name where the incident occurred
Detail	Event attributes that triggered the incident
Status	The status of the incident, **Active**, **Cleared**, **Cleared Manually**, **System Cleared**
Cleared Reason	For manually cleared incidents, this displays the reason the incident was cleared
Cleared Time	The time an incident was cleared
Cleared User	The person who cleared the incident
Comments	Any comments that users have entered for the incident
Ticket Status	Status of any tickets associated with the incident
Ticket ID	The ID number of any tickets generated by the incident
Ticket User	The person assigned to any tickets generated by the event
External User	If the ticket was cleared in an external ticket-handling system, this lists the name of the person the ticket was assigned to
External Cleared Time	If the ticket was cleared in an external ticket-handling system, this lists the time it was cleared

Column Name	Description
External Resolved Time	If the ticket was resolved in an external ticket-handling system, this lists the time it was resolved
External Ticket ID	The ID of the incident in an external ticket-handling system
External Ticket State	The state of the incident ticket in an external ticket-handling system
External Ticket Type	The type assigned to the incident ticket in an external ticket-handling system
Organization	The organization reporting the event
Impacts	Organizations impacted by the event
Business Service	Business services impacted by the incident
Incident Notification Status	Status of any notifications that were sent because of the incident
Notification Recipients	Who received notification of the incident
Incident Count	How many times the incident has occurred during the selected time interval

Viewing Incidents

- Device Risk View of all incidents
- Viewing incident details
- Grouped View of all incidents

Device Risk View of all incidents

This is the default view when user clicks the Incident tab. It shows a list of devices that triggered incidents. Devices are ranked by a risk score that is computed by combining asset criticality, triggered incidents and found security vulnerabilities (details - here).

To see the incidents for a device, click that device. The incidents show up in a time line view.

- Active Incidents over the last 2 hours are displayed
- The following incident attributes are shown
  - **Severity** - High, Medium, Low - shown by colored icons
  - **Last Occurred** - the last time the Incident happened
  - **Reporting Device Name** - names of devices that reported the events that led to the incident
  - **Incident** - rule name
  - **Source** - incident source
  - **Target** - incident target

- **Detail** - incident parameters other than source and target
- **Count** - number of times the same incident has triggered

To show incidents over a different time interval:

1. Click **Time Range** Button. A search window appears
2. To choose a relative time window
   - Choose **Time Range** Operator as **LAST.**
   - Specify the number of Minutes/Hours/Days/Weeks.
   - Click **Check** button.
   - The Incident page will automatically refresh to show all the incidents over the time window.
3. To choose an absolute time window
   - Choose **Time Range** Operator as **FROM.**
   - Specify the starting and end times.
   - Click **Check** button.
   - The Incident page will automatically refresh to show all the incidents over the time window

An incident can be in any of the following states:

- Active
- Cleared
- Cleared Manually
- System Cleared

By default only Active Incidents are shown. To show Incidents in other states

1. Click **Incident Status** Button. A search window appears.
2. To add a new value, click on the white space next to the selected value. A menu appears. Select the needed values one by one.
3. Click **Check** button. The Incident page will automatically refresh to show all the incidents in selected state(s).

To select a different set of Incident attributes

1. Click **Choose columns** icon
2. In the popup, select the columns you want to display by moving them to the right. You can re-order the position of the columns.
3. Click**OK.**To force a refresh of the incident view, click the **Refresh** icon

Incidents may be displayed over multiple pages. To see incidents on a different page,

1. Select the **Page Selector** icon
2. Either enter the page number or click on the Next or Previous icon to go to the right page

To view incidents for a different organization (Service Provider version),

1. Click the **User icon** on top right
2. Choose the right organization
3. Click **Change View**

Viewing incident details

In the default view, an incident is shown in a single line. To see the details of the incident,

- Click anywhere on the incident line
- Basic incident attributes are shown immediately below the incident

- More advanced incident attributes are shown in a bottom pane

To revert to the single line incident view, click anywhere on the incident line. Detailed views will disappear.

To view the rule that triggered the incident,

- Click anywhere on the incident line in the single line incident view
- In the bottom pane, Click **Rule** tab**.** Rule details are displayed.

To view the events that triggered the incident

- Click anywhere on the incident line in the single line incident view
- In the bottom pane, Click **Events** tab**.** Basic Event attributes are displayed in a single line.
- To see the raw events, click on the Basic Event line. Raw events are displayed.

Grouped View of all incidents

Sometimes user may need a grouped view of incidents to get an overview of what incidents have triggered and involves which devices. The following grouped views are provided

- **Severity** - Ranks Incident Severities By Count
- **Name** - Ranks the Incidents By Count
- **Name, Target** - Ranks Incident Name and Incident Target By Count
- **Name, Source** - Ranks Incident Name and Incident Source By Count
- **Name, Source, Target** - Ranks Incident Name, Incident Source and Incident Target By Count
- **Name, Source, Target, Business Service** - Ranks Incident Name, Incident Source, Incident Target and Business Services By Count
- **Name, Source, Target, Business Service, Organizations** - Ranks Incident Name, Incident Source, Incident Target, Business Services and Organizations By Count

To get a grouped view

- Choose the desired view from **Group By** drop down

Group view works with Search

Grouped view works with Search filters. In other words, Grouped view includes the incidents where the search conditions are satisfied.

To get to a un-grouped view from grouped view,

- Choose **None**in the **Group By** selector.

Searching Incidents

Searchable Incident Attributes

Incident Attribute	Description
Time Range	In
ID	Incident ID
IP	Incident Source IP or Incident Target IP
Host	Host name associated with Incident Source IP or Incident Target IP
User	User field specified in Incident Target or Incident Details
Severity	Incident Severity category - High, Medium or Low
Function	Security, Availability, Performance or Change. This is a property of an Incident.
Incident Status	Possible values are Active, Cleared, Cleared Manually, System Cleared
Ticket Status	Possible values are New, Open, Closed, External, reopened, None. External means opened in an external system.
Incident	Rule name
Biz Service	Business Service name
Organization	Organization name

Constructing Search Condition

To construct a Search condition from a displayed Incident,

- Mouse over the cell containing the specific Incident attribute
- Right click and choose **Add to filter**
- The condition will be added to existing search string
- Matching incidents will be displayed

To construct a Search condition from scratch:

- Click on the **Add filter** edit area. Three fields are displayed
  - Incident Attribute
  - Operator
  - Value

- Select one of the **Incident Attributes** from the drop down
- Select an **Operator** from =, != IN, NOT IN, CONTAINS, NOT CONTAINS
- Select one or more **Values** from the displayed choices
- Click the Check button.
- Matching incidents will be displayed

Managing Incidents

Adding Comments

1. Click on an Incident in the un-grouped view
2. From **Actions** drop down, select **Add Comments**
3. Write the comment and click **OK.**

Clearing Incidents

1. Click on an Incident in the un-grouped view
2. If you have more incidents to clear, then press Shift and click on the second incident. This will will select all incidents between the first one and this one. To get this approach to work effectively,
    - Create a filter to get all the incidents to be cleared in view
    - Select the first incident
    - Press Shift and click on the last incident - all incidents are now selected
3. From **Actions** drop down, select **Clear**
4. Click **OK**

Exporting Incidents to a PDF document

1. Click on an Incident in the un-grouped view
2. If you have more incidents to export, then press Shift and click on the second incident. This will will select all incidents between the first one and this one. To get this approach to work effectively,
    - Create a filter to get all the incidents to be exported in view
    - Select the first incident
    - Press Shift and click on the last incident - all incidents are now selected
3. From **Actions** drop down, select **Export**
4. Click **OK**

Device Risk Score Computation

Risk computation algorithms are proprietary and this section presents only the knobs that user is able to tweak to change the score.

Risk score components

The following factors affect risk score of a device

1. Device Importance (also called Asset Weight)
2. Count and CVS Score for non-remediated vulnerabilities found for that device
3. Severity and Frequency of Security incidents triggering with that device as source or destination
4. Severity and Frequency of Other (performance, availability and change) incidents triggering on that device

Overall Score (0-100) is a weighted average of 3 components - Vulnerability Score, Security Incident Score and Other Incident Score, computed as follows.

```
Overall Risk = vul_weight * Vulnerability Score + security_inci_weight *
Security Incident Score + other_inci_weight * Other Incident Score.
```

User controllable constants

1. Device Importance - this can be set in CMDB > Device > Summary. You can select multiple devices and set the Importance in one shot. Values are
    a. Mission Critical - 10
    b. Critical - 7
    c. Important - 4
    d. Normal - 1
2. Relative weights of Vulnerabilities, Security and Other incidents to the risk score. The default values of the constants are defined in phoenix_config.txt:
    a. vul_weight = 0.6
    b. security_inci_weight = 0.3
    c. security_inci_weight = 0.1
3. Maximum number of high-severity events that a mission-critical host can tolerate for each of the 3 score components. These default thresholds are defined in ‘phoenix_config.txt:
    a. vul_threshold = 1
    b. security_inci_threshold = 3
    c. other_inci_threshold = 6

Time varying Risk score

Risk scores are computed for each day. Current risk score is a exponentially weighted average of today's risk and yesterday's risk.

The algorithm also reduces the score for earlier vulnerabilities that are now patched. Such vulnerabilities have a weight of 0.7 while new and old but existing vulnerabilities have weight 1

Miscellaneous Operations

Exporting Events to Files

You can run the phExportEvent tool from a Supervisor or Worker node to export events to CSV files. The file will contain these fields:

- Customer Id (applicable to SP license)
- Reporting Device IP
- Reporting Device Name
- Event Received Time
- Raw Message

This code block shows the commands that you can use with phExportEvent, followed by a table that describes them in more detail.

```
phExportEvent {--dest DESTINATION_DIR} {--starttime START_TIME | --relstarttime RELATIVE_
START_TIME} {--endtime END_TIME | --relendtime RELATIVE_END_TIME} [--dev DEVICE_NAME] [--
org
ORGANIZATION_NAME] [-t TIME_ZONE]
```

phExportEvent Command	Description
DESTINATION_DIR	Destination directory where the exported event files are saved
START_TIME	Starting time of events to be exported. The format is YYYY-MM-DD HH:MM:SS {+|-} TZ. If TZ is not given, local time zone of the machine where the script is running will be used. Example: `2010-03-10 23:00:00 -8` means Pacific Standard Time, 23:00:00 03/10/2010. `2010-07-29 10:20:00 +5:30` means India Standard Time 10:20:00 07/29/2010.
RELATIVE_START_TIME	Must be used together with `END_TIME`. Starting time of events to be exported relative backward to the end time as specified using `--endtime END_TIME`. The format is `{NUM}{d|h|m}` where `NUM` is the number of days or hours or minutes. For example, `--relstarttime 5d` means the starting time is 5 days prior to the ending time.
END_TIME	Ending time of events to be exported. The format is the same as `START_TIME`.
RELATIVE_END_TIME	Must be used together with `START_TIME`. Ending time of events to be exported relative forward to the start time as specified using START_TIME. The format is same as RELATIVE_START_TIME.

phExportEvent Command	Description
DEVICE_NAME	Host name or IP address of the device with the events to be exported. Use a comma-separated list to specify multiple IPs or host names, for example, `--dev 10.1.1.1,10.10.10.1,router1,router2`. Host name is case insensitive
ORGANIZATION_NAME	Used only for Service Provider deployments. The name of the organization with the events to be exported. To specify multiple organizations, enter a commandeach for one organization, for example, `--org "Public Bank" --org "Private Bank"`. The organization name is case insensitive.
TIME_ZONE	Specifies the time zone used to format the event received time in the exported event files. The format is `{+|-}TZ`, for example, `-8` means Pacific Standard Time, `+5:30` means India Standard Time.

Dynamic Population of Location, User, and Geolocation Information for Events

In most cases, network logs only contain IP address information, but to investigate incidents involving that IP, you need additional context for that IP address such as host name, user, and geolocation information. Because FortiSIEM collects detailed IT infrastructure information in the CMDB, it is able to correlate that information to the IP address to create a context for the event, and insert that context information into events in real time as parsed attributes. This topic describes the way in which this context information is populated into events.

Correlating Event Information

Event information is derived from several different sources.

1. During the discovery process, FortiSIEM discovers the host name and network interface address information during discovery and stores them in the CMDB. If any IP address other than the **Access IP** changes, then running a rediscovery will update the CMDB with the right information.
2. FortiSIEM collates information from various authentication logs and forms a time-based Identity and Location Report containing the IP address, MAC address, Host Name, Domain, User, Network Access Point, and Network Access Point Port for the event.
3. The geolocation database maps IP addresses to Country, State, City, Organization, Longitude, and Latitude information.

Assigning Attributes to Events

When FortiSIEM parses an event, attributes are assigned to it following this process:

Host Name Attribute

For each IP address (**Host IP**, **Source IP**, **Destination IP**, **Reporting IP**):

1. FortiSIEM checks the CMDB for an associated host name, and if one is found, then the host name is added to the event.
2. If the host name is not found in then CMDB, then FortiSIEM checks the Identity and Location Report for the host name, and if one is found, then it is added to the event.
3. If the host name is not found in either the CMDB or Identity and Location Report, then FortiSIEM runs DNS lookup for the host name, and if one is found, then it is added to the event. For performance reasons the DNS result is cached, and because excessive DNS lookups can cause event processing delays, FortiSIEM has an algorithm to dynamically bypass DNS lookup if it begins falling behind in event processing.

User Name Attribute

For **Source IP**, FortiSIEM checks for user information in the Identity and Location Report, and if anything is found, it is added to the event.

Geolocation Attribute

For each IP address (**Host IP**, **Source IP**, **Destination IP**, **Reporting IP**), FortiSIEM checks the geolocation database. If geolocation information is found for that IP, then Country, State, City, Organization, Longitude, and Latitude information is added to it.

Dynamic Updating of Attribute Information

For any of these attributes, when there is a change in the infrastructure (for example, a network device has a new IP or a new user logs on to the system), the change is populated into the CMDB and/or Identity and Location Report, and the event parsing module learns of the change and starts populating events with the new metadata.

Because the FortiSIEM approach to populating event attributes is dynamic and change driven, it is always able to map the right IP address to host names and users in the face of dynamic changes in the IT infrastructure.

Attributes Added to Events

IP Type	Attributes
Source IP	• Source Host Name • User (corresponding to Source IP) • Source Country • Source State • Source City • Source Organization • Source Longitude • Source Latitude
Destination IP	• Destination Host Name • Destination Country • Destination State • Destination City • Destination Organization • Destination Longitude • Destination Latitude
Host IP	• Host Name • Host Country • Host State • Host City • Host Organization • Host Longitude • Host Latitude
Reporting IP	• Reporting Host Name • Reporting Country • Reporting State • Reporting City • Reporting Organization • Reporting Longitude • Reporting Latitude
PostNAT (Network Address Translation) IP	• PostNAT Country • PostNAT State • PostNAT City • PostNAT Organization • PostNAT Longitude • PostNAT Latitude

Monitoring Custom Applications

While FortiSIEM provides support for many applications, there may also be situations in which you have a custom application running in your infrastructure that you want to monitor. This topic explains how to set up FortiSIEM to monitor that application, and add it to a business service.

1. Log in to your Supervisor.
2. Go to **CMDB > Applications**, and either select a group where you want to add the application, or create a new one.
3. Click **New**, and enter an **Application Name** and a **Process Name**.
4. Click **Save**.
5. Initiate discovery of the server where the application is running.
6. Go to **CMDB > Devices** and select the server.
7. Click the **Software** tab and make sure the application has been discovered.
8. Go to **General Settings > Monitoring > Important Processes**.
9. Click **Add** and enter the name of the process that the application is running on.
10. Click **Apply All**.
11. Run a structured historical search using these attributes to make sure the process utilization metrics are being received by FortiSIEM.

Attribute	Value
Reporting IP	The IP address of the server where the application is running
Event Type	`PH_DEV_MON_PROC_RESOURCE_UTIL`
Application Name	The name of the application

12. Add your application to a business service.
You should now be able to go **Dashboard > Summary Dashboards > Biz Service Summary** and see your process running under **Top Monitored Processes** when you select the associated business service.

IPS Vulnerability Map

The IPS Vulnerability Map lists devices that have a known vulnerability. You can view the IPS Vulnerability Map by going to **Incidents > IPS Vunerability Map**, and you can also add new devices to the map.

The IPS Vulnerability Map includes these columns.

Column	Description
IPS Event Types	The event types associated with the vulnerability
Vendor Vulnerability ID	The vulnerability ID provided by the device vendor
CVE IDs	The vulnerability ID provided by Common Vulnerabilities and Exposures
Vulnerability Description	A brief description of the device's vulnerability
Found in Device Type	Specific devices or applications that have the vulnerability
Found in Version	The version of the device or application that has the vulnerability
Fixed in Version	The version in which the vulnerability was fixed
Fixed via Patches	The patch version in which the vulnerability was fixed

Adding Entries to the IPS Vulnerabilities Map

Updating IPS Vulnerability Map Entries: You can update existing entries, for example when a patch is released to fix a vulnerability, by selecting an entry in the IPS Vulnerabilities Map and clicking **Edit**.

1. Log in to your Supervisor node.
2. Go to **Incidents > IPS Vulnerability Map**.
3. Click **Add**.
4. Select the **IPS Event Type** associated with the vulnerability.
5. Enter any **Vendor Vulnerability ID**s.
6. Enter any **CVE ID**s.
   See the Common Vulnerability and Exposures website for CVE IDs. Separate multiple IDs with commas.
7. Enter a **Vulnerability Description**.
8. For **Affected Software**, click **Add**, and then select the affected devices or applications from the **Found in Device Type** menu.
9. Enter any **Found in Version** information for the affected software.
10. Enter any fix information for the vulnerability.
11. Click **OK**.
12. Click **Save**.