# Sentinel™ 5

# User's Guide v5.1.1

- Linux
- Solaris
- Windows

## Volume II of V

**March 2006**
**www.esecurity.net**

# Preface

The e-Security Technical documentation is general-purpose operation and reference guide. This documentation is intended for Information Security Professionals. The text in this documentation is designed to serve as a source of reference about e-Security's Enterprise Security Management System. There is additional documentation available on the e-Security web portal.

e-Security Technical documentation is broken down into five different volumes. They are:

- Volume I – Sentinel™ 5 Install Guide
- Volume II – Sentinel™ 5 User's Guide
- Volume III – Sentinel™ 5 Wizard User's Guide
- Volume IV – Sentinel™ 5 User's Reference Guide
- Volume V – Sentinel™ 3rd Party Integration

## Volume I – Sentinel Install Guide

This guide explains how to install:

- Sentinel Server
- Sentinel Console
- Sentinel Correlation Engine
- Sentinel Crystal Reports
- Wizard Agent Builder
- Wizard Agent Manager
- Advisor

## Volume II – Sentinel User's Guide

This guide discusses:

- Sentinel Console Operation
- Sentinel Features
- Sentinel Architecture
- Sentinel Communication
- Shutdown/Startup of Sentinel
- Vulnerability assessment
- Event monitoring
- Event filtering
- Event correlation
- Sentinel Data Manager
- Event Configuration for Business Relevance
- Mapping Service
- Historical reporting
- Wizard Host Management
- Incidents
- Cases
- User management
- Workflow

## Volume III – Wizard User's Guide

This guide discusses:

- Wizard Agent Builder Operation
- Wizard Agent Manager
- Agents
- Wizard Host Management
- Building and maintaining agents

## Volume IV - Sentinel User's Reference Guide

This guide discusses:

- Wizard scripting language
- Wizard parsing commands
- Wizard administrator functions
- Wizard and Sentinel meta-tags
- Sentinel correlation engine
- User Permissions
- Correlation command line options
- e-Security database schema

### Volume V - Sentinel 3rd Party Integration Guide

- Remedy
- HP OpenView Operations
- HP Service Desk

# Chapter 1 – Sentinel Introduction

Sentinel™ 5 is the leading security information management and compliance monitoring solution that receives information collected from many sources throughout an enterprise, standardizes it, prioritizes it and performs correlation all in real-time. Sentinel collects data from many security products on the market and provides the flexibility to collect data from new technologies and products as installations and business requirements evolve.

Many of the capabilities in Sentinel 5 are the result of architectural re-design of Sentinel 4.0 and driven by the needs of e-Security's customers. As security threats and regulatory pressure increase, organizations are looking for one solution that will enable them to:

- Gain the visibility and insight required to manage a security environment more cost-effectively.
- Continuously monitor compliance with internal policies and government regulations (e.g. Sarbanes-Oxley, HIPAA, GLBA, FISMA, NISPOM, DCID 6/3 and DITSCAP).
- Identify and resolve incidents faster and more cost-effectively through centralized, automated collection and resolution of threat and policy data.
- Provide operational and executive metrics to continually assess security and compliance posture and address both tactical and strategic goals.
- Reduce operational costs associated with security and compliance monitoring, incident identification and remediation.

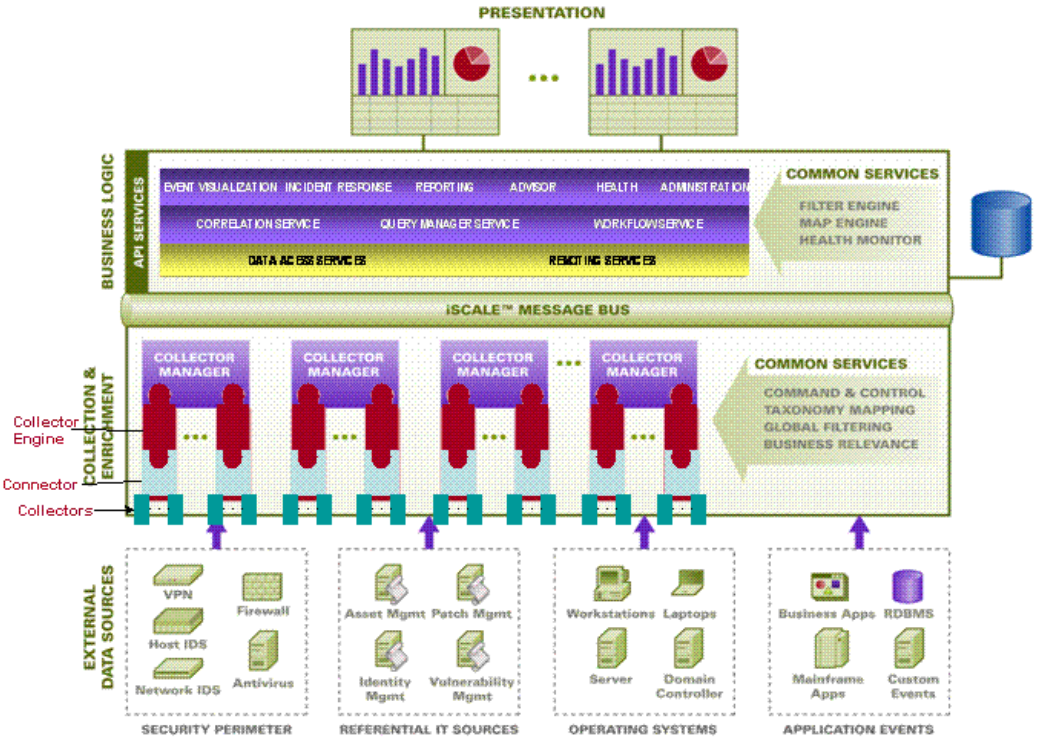


An event is an action or occurrence reported to Sentinel. An event received from a security device is called an external event and an event generated by Sentinel is called an internal event. Events can be security-related, performance-related

or information related. For example, an external event could be an attack detected by an Intrusion Detection System (IDS), a successful login reported by an operating system or a customer-defined situation such as a user accessing a file. Internal events are generated by Sentinel to indicate a noteworthy change to the state of the system such as an agent being stopped or a correlation rule being disabled.

Correlation is the process of analyzing security events to identify patterns within an event or a stream of events. For example, a correlation rule can be created to detect when thirty or more ICMP events occur within one minute time period. High volume traffic (flood) of ICMP could result in a denial of service attack. Correlation can detect patterns in a stream of events from a single device, a set of similar devices or an arbitrary collection of devices. This allows the user to make a better determination of the risk and of the severity of the incident.

Sentinel also incorporates additional information into the feed, such as information about the machines in the network and their known services and vulnerabilities. This information is made available in real-time further refining the significance of the events being monitored.

Sentinel Control Center uses background processes to display real-time events and event summaries (Active Views™), Incidents, historical reports (Analysis) and Advisor reports.

Events that are deemed of significant importance can be grouped together in an object called an *Incident*. An incident can be created manually by the user or automatically by the correlation engine. The incident can hold additional information such as information about assets that are being attacked, the vulnerabilities of these assets, information about the attack that is retrieved from the e-Security Advisor component. In addition, other information can be appended as attachments.

This guide assumes that you are familiar with the basics of network security, database administration, Windows and UNIX operating system environments.

This chapter describes Sentinel 5's functional and logical architecture, followed by its key product modules.

## Functional Architecture

Sentinel 5 is composed of three component subsystems, which form the core of the functional architecture:

- iSCALE Platform – an event-driven scalable framework
- Data Source Integration – an extensible collector framework
- Application Integration – an extensible application framework

Sentinel treats both "services" and "applications" as abstract service endpoints that can readily respond to asynchronous events. Services are "objects" that do not need to understand protocols or how messages get routed to the peer services.

## Sentinel Features

Sentinel is a feature-rich end-user application that allows one to monitor and manage a variety of functions. Some of the main functions include:

- Provides real time views of large streams of events
- Provides reporting capabilities based off real time and historical events
- Regulates users and what they are able to see and do by permission assignment
- Allows you to restrict which events users have access to
- Allows you to organize events into incidents for efficient response management and tracking
- Allows you to detect patterns in events and streams of events

## Architecture Overview

The Sentinel system is responsible for receiving events from the Wizard Agent Manager. The events are then displayed in real-time and logged into a database for historical analysis.

At a high level, the Sentinel system uses a relational database and is comprised of Sentinel processes and a reporting engine. The system accepts events from the agent manager as its input. The agent manager interfaces with third-party products and normalizes the data from these products. The normalized data is then sent to the Sentinel processes and database.

Historical analysis and reporting can be done using e-Security's integrated reporting engine. The reporting engine extracts data from the database and integrates the report displays into the Sentinel Control Center using HTML documents over an HTTP connection.

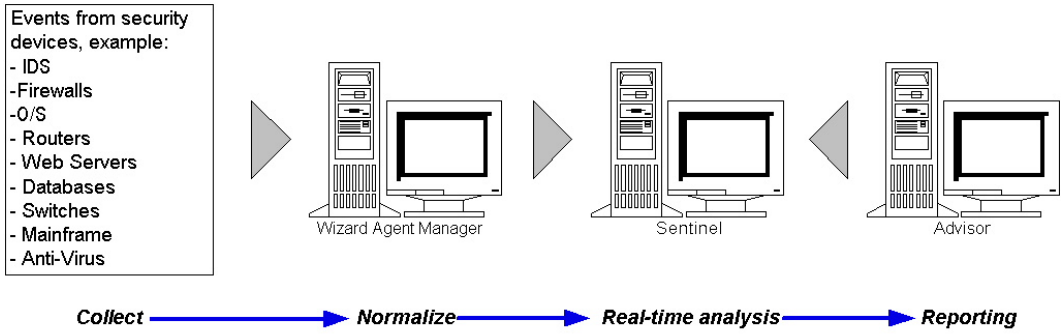


Sentinel Features are:

- Real-time processing of events that are received from the Wizard Agent Manager
- An intuitive and flexible rule-based language for correlation
- Rules compiled for high performance
- Scalable, multi-threaded, distributable and extensible architecture

Sentinel processes communicate with each other through a message-oriented middleware (MOM).

### iSCALE Platform

e-Security's iSCALE™ architecture is built using a standards-based, Service-Oriented architecture (SOA) that combines the advantages of in-memory processing and distributed computing. At the heart of iSCALE is a specialized message bus capable of handling high data volumes. Built from the ground up

using a best of breed, standards-based approach, iSCALE can scale cost-effectively.

**Message Bus**

The iSCALE Message Bus allows for independent scaling of individual components while also allowing for standards-based integration with external applications. The key to scalability is that unlike other distributed software, no two peer components communicate with each other directly. All components communicate via the message bus, which is capable of moving thousands of message packets per second.

Leveraging the message bus' unique features, the high-throughput communication channel can maximize and sustain a high data throughput rate across the independent components of the system. Events are compressed and encrypted on the wire for secure and efficient delivery from the edge of the network or collection points to the hub of the system, where real-time analytics are performed.

The iSCALE message bus employs a variety of queuing services that improve the reliability of the communication beyond the security and performance aspects of the platform. Using a variety of transient and durable queues, the system offers unparalleled reliability and fault tolerance. For instance, important messages in transit are saved (by being queued) in case of a failure in the communication path. The queued message is delivered to the destination after the system recovers from failure state.

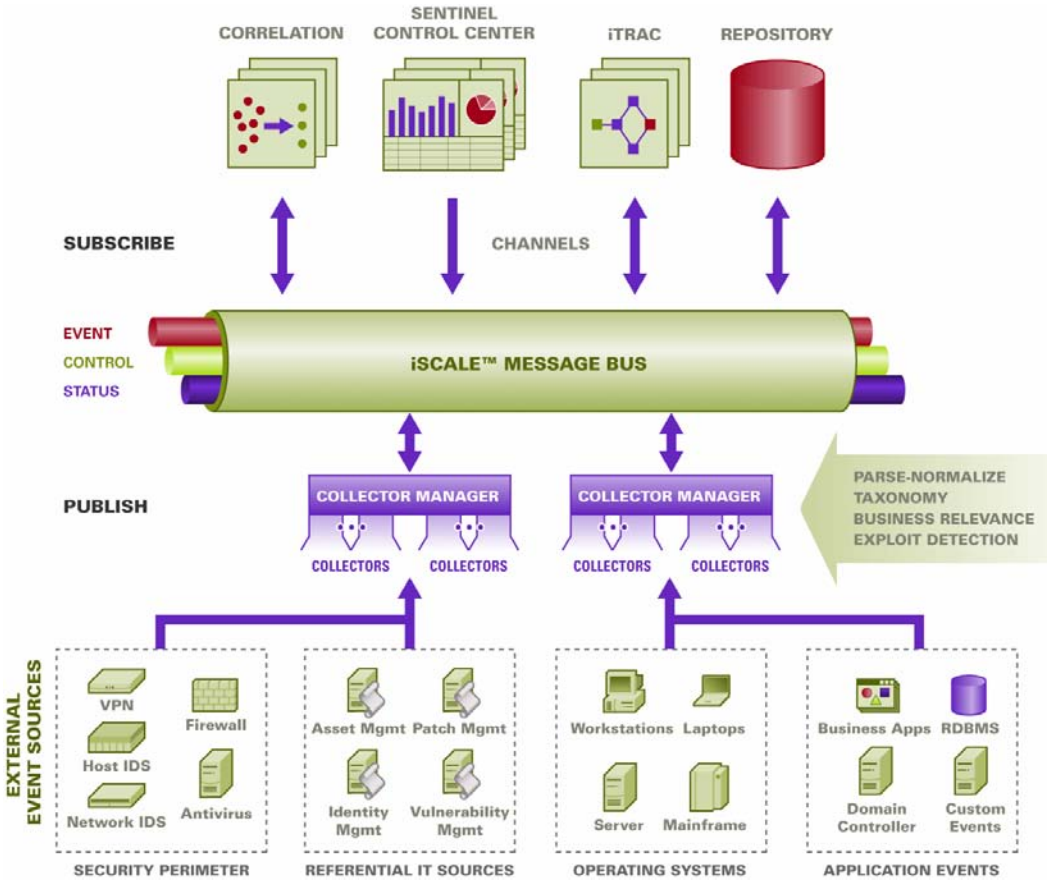

### Channels

The iSCALE platform employs a data-driven or event-driven model that allows independent scaling of components for the entire system based on the workload. This provides a flexible deployment model since each customer's environment varies: one site may have a large number of devices with low event volumes; another site may have fewer devices with very high event volumes. The event densities (i.e., the event aggregation and event multiplexing pattern on the wire from the collection points) are different in these cases and the message bus allows for consistent scaling of disparate workloads.

iSCALE takes advantage of an independent, multi-channel environment, which virtually eliminates contention and promotes parallel processing of events. These channels and sub-channels work not only for event data transport but also offer fine-grain process control for scaling and load balancing the system under varying load conditions. Using independent service channels such as control channels and status channels, in addition to the main event channel, allows sophisticated and cost-effective scaling of event-driven architecture.

## Sentinel Event

Sentinel receives information from devices, normalizes this information into a structure called a *Sentinel Event,* or *Event* for short and sends the event for processing. Events are processed by the real time display, correlation engine and the backend server.

An event comprises of more than 200 tags. Tags are of different types and of different purposes. There are some predefined tags such as severity, criticality, destination IP and destination port. There are two sets of configurable tags: Reserved Tags are for eSecurity internal use to allow future expansion and Customer Tags, which are for customer extensions.

Tags can be repurposed by renaming them. The source for a tag can either be *external*, which means that it is set explicitly by the device or the corresponding agent or *referential.* The value of a referential tag is computed as a function of one or more other tags using the mapping service. For example, a tag can be defined to be the building code for the building containing the asset mentioned as the destination IP of an event. For example, a tag can be computed by the mapping service using a customer defined map using the destination IP from the event.

### Mapping Service

Map Service allows a sophisticated mechanism to propagate business relevance data throughout the system. This facility aids scalability and provides an extensibility advantage by enabling intelligent data transfer between different nodes of the distributed system.

Map Service is a data propagation facility that gives the ability to cross-reference Vulnerability Scanner data with Intrusion Detection System signatures and more (e.g. asset data, business-relevant data). This allows immediate notification when an attack is attempting to exploit a vulnerable system. Three separate components provide this functionality:

- collection of real time events from an intrusion detection source;

- comparing those signatures to the latest vulnerability scans; and
- cross referencing an attack feed via Sentinel Advisor (an optional product module, which cross-references between real-time IDS attack signatures and the user's vulnerability scanner data).

Map Service dynamically propagates information throughput the system without impacting system load on the system. When important data sets (i.e., "maps" such as asset information or patch update information) are updated in the system, the Map Service propagates the updates across the system, which can often get to be hundreds of megabytes in size.

iSCALE's Map Service algorithms handle large referential data sets across a production system processing large real-time data volumes. These algorithms are "update-aware" and selectively push only the changes or "delta data sets" from the repository to the edge or system perimeter.

**Streaming Maps**

Map Service employs a dynamic update model and streams the maps from one point to another, avoiding the build up of large static maps in dynamic memory. The value of this streaming capability is particularly relevant in a mission-critical real-time system such as Sentinel where there needs to be a steady, predictive and agile movement of data independent of any transient load on the system.

**Exploit Detection (Mapping Service)**

Sentinel provides the ability to cross-reference event data signatures with Vulnerability Scanner data. Users are notified automatically and immediately when an attack is attempting to exploit a vulnerable system. This is accomplished through:

- Advisor Feed
- Intrusion detection
- Vulnerability scanning
- Firewalls

Advisor provides a cross-reference between event data signatures and vulnerability scanner data. Advisor feed has an alert and attack feed. The alert feed contains information about vulnerabilities and threats. The attack feed is a normalization of event signatures and vulnerability plug-ins. For information about Advisor installation, see the Sentinel Installation Guide.

The supported systems are:

**Intrusion Detections Systems**
- Cisco Secure IDS
- Enterasys Dragon Host Sensor
- Enterasys Dragon Network Sensor
- Intrusion.com (SecureNet_Provider)
- ISS BlackICE
- ISS RealSecure Desktop
- ISS RealSecure Network
- ISS RealSecure Server
- ISS RealSecure Guard
- Snort
- Symantec Network Security 4.0 (ManHunt)

**Vulnerability Scanners**
- eEYE Retina
- Foundstone Foundscan
- ISS Database Scanner
- ISS Internet Scanner
- ISS System Scanner
- ISS Wireless Scanner
- Nessus
- nCircle IP360
- Qualys QualysGuard

**Firewalls**
- Cisco IOS Firewall

- Symantec Intruder Alert
- McAfee IntruShield

You will require at least one vulnerability scanner and either an IDS or firewall from each category above. Your IDS and Firewall DeviceName (rv31) has to appear in the event as hi-lighted in gray above. Also, your IDS and Firewall must properly populate the DeviceAttackName (rt1) field (such as, WEB-PHP Mambo uploadimage.php access).

The Advisor feed is sent to the database and then to the Exploit Detection Service. The Exploit Detection Service will generate one or two files depending upon what kind of data has been updated.

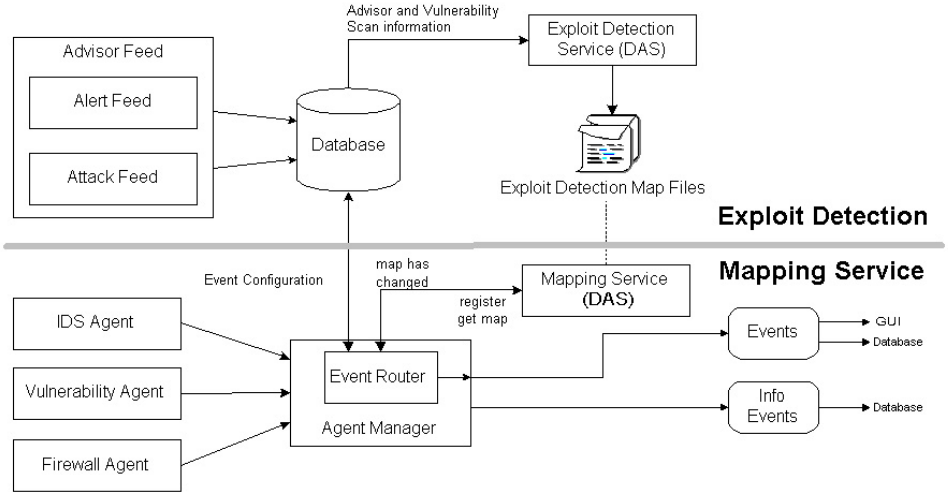


The Exploit Detection Map Files are used by the Mapping Service to map attacks to exploits of vulnerabilities.

Vulnerability Scanners scan for system (asset) vulnerable areas. IDS' detect attacks (if any) against these vulnerable areas. Firewalls detect if any traffic is against any of these vulnerable area. If an attack is associated with any vulnerability, the asset has been exploited.

The Exploit Detection Service generates two files located in:

```
$ESEC_HOME/sentinel/bin/map_data
```

The two files are attackNormalization.csv and exploitDetection.csv.

The attackNormalization.csv is generated after

- Advisor feed
- DAS Startup (if enabled in das_query.xml, disabled by default)

The exploitDetection.csv is generated after one of the following:

- Advisor feed
- Vulnerability scan
- Sentinel Server Startup (if enabled in das_query.xml, disabled by default)

By default, there are two configured event columns used for exploit detection and they are referenced from a map (all mapped tags will have the scroll icon).

- Vulnerability

- AttackId

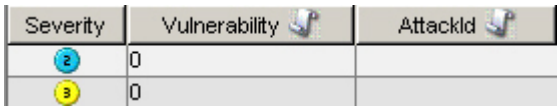


When the vulnerability field (*vul*) equals 1, the asset or destination device is exploited. If the vulnerability field equals 0, the asset or destination device is not exploited.

Sentinel comes pre-configured with the following map names associated with attackNormalization.csv and exploitDetection.csv.

| Map Name | csv File Name |
|---|---|
| ▪ AttackSignatureNormalization | ▪ attackNormalization.csv |
| ▪ IsExploitWatchlist | ▪ exploitDetection.csv |

There are two types of data sources:

- External - retrieves information from the agent
- Referenced from Map - retrieves information from a map file to populate the tag.

The AttackId tag has the Device (type of the security device, e.g. - Snort) and AttackSignature columns set as Keys and uses the NormalizedAttackID column in the attackNormalization.csv file. In a row where the DeviceName event tag (an IDS device such as Snort, information filled in by Advisor and Vulnerability information from the Sentinel Database) is the same as Device and where the DeviceAttackName event tag (attack information filled in by Advisor information in the Sentinel Database via the Exploit Detection Service) is the same as AttackSignature, the value for AttackId is where that row intersects with the NormalizedAttackID column.

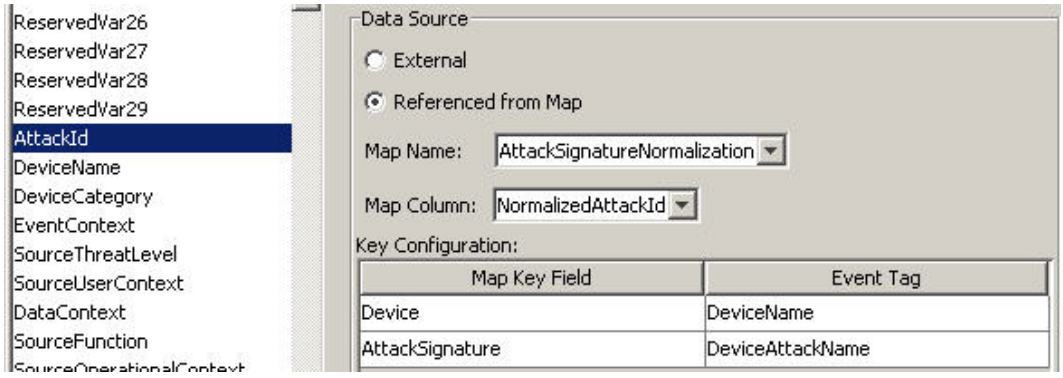


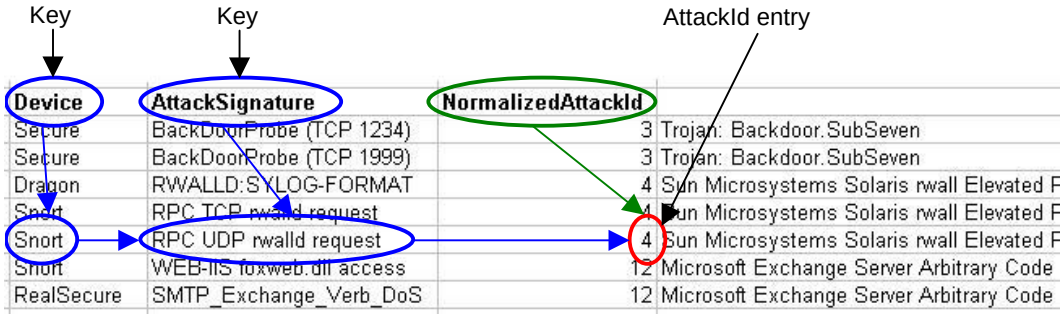

The Vulnerability tag has a column entry "_EXIST_", which means that map result value will be 1 if the key is in IsExploitWatchlist (exploitDetection.csv file) or 0 if it is not. The key columns for the vulnerability tag are IP and NormalizedAttackId. When an incoming event with a DestinationIP event tag that matches the IP column entry and an AttackId event tag that matches the NormalizedAttackId column entry in the same row, the result is a one (1). If no match is found in a common row, the result is zero (0).

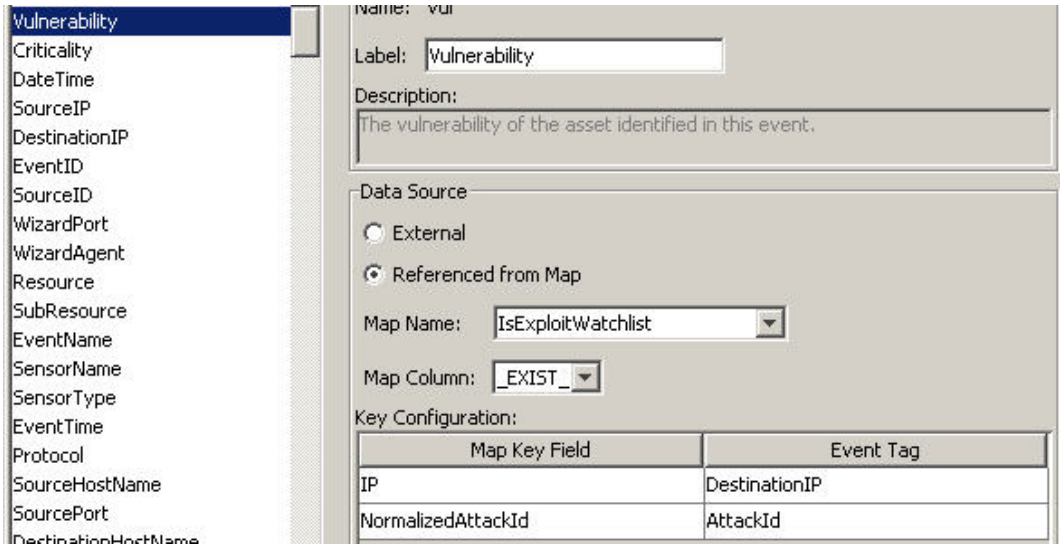


### Data Source Integration

Using adaptable and flexible technology is central to Sentinel's data source integration strategy, which is achieved via interpretive Agents (also referred to as Collectors) that parse and normalize the events in the data stream.

These Agents can be modified as needed and are not tied to a specific environment. The creation, modification, deployment and maintenance of Agents are simple and can be done by users directly. An integrated development environment allows for interactive creation of Agents using a "drag and drop" paradigm from a graphical user interface. Non-programmers can create Agents, ensuring both current and future requirements are met in an ever-changing IT environment. The command and control operation of Agents (e.g. start, stop) is performed centrally from the Sentinel Control Center.

### Application Integration

External application integration via standard APIs is central to Sentinel. For example, a bi-directional API for trouble ticketing systems including Remedy® and HP OpenView's ServiceDesk® allows straightforward integration with external systems.

The API is Web Services-based and therefore allows any external systems that are SOAP-aware to take advantage of pervasive integration with the Sentinel system.

## Time

The time of an event is very critical to its processing. It is important for reporting and auditing purposes as well as for real time processing. The correlation engine

processes time ordered streams of events and detects patterns within events as well as temporal patterns in the stream. However, the device generating the event may not know the real time when the event is generated. In order to accommodate this Sentinel allows two options in processing alerts from security devices: trust the time the device reports and use that as the time of the event, or, do not trust the device time and instead stamp the event at the time it is first processed by Sentinel (by the agent).

Sentinel is a distributed system and comprises several processes that can be in different parts of the network. In addition, there can be some delay introduced by the device. In order to accommodate this, the Sentinel processes reorder the events into a time ordered stream before processing.

The following illustration explains the concept of Sentinel Time.

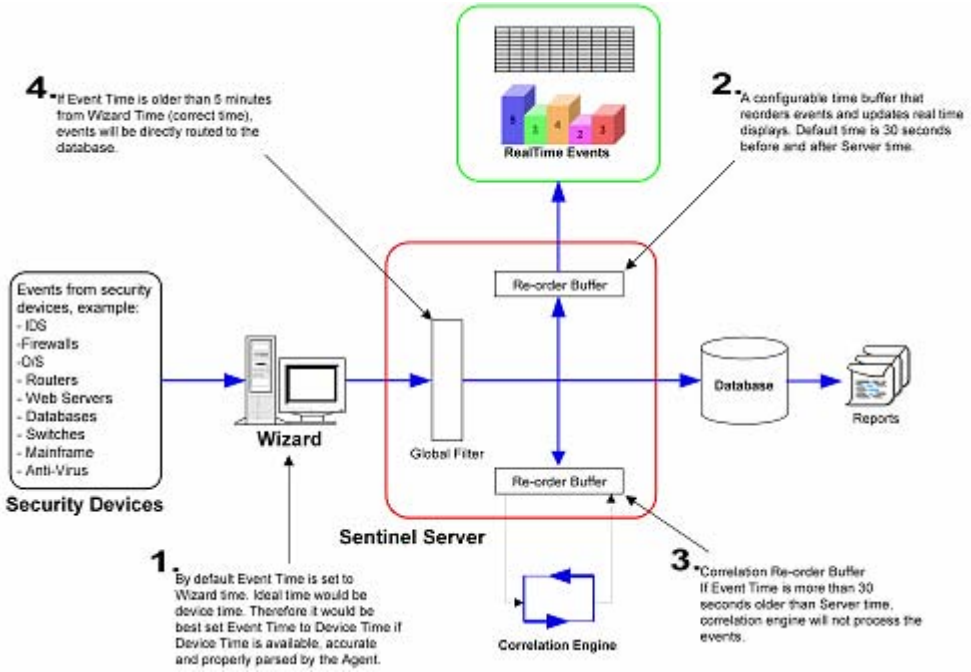


1. By default, Event Time is set to Wizard time. Ideal time would be device time. Therefore it would be best to set Event Time to Device Time if Device Time is available, accurate and properly parsed by the Agent.
2. A configurable time buffer that reorders events and updates real time displays. Default time is 30 seconds before and after server time.
3. Correlation Re-order buffer, if event time is more than 30 seconds older than Server time, correlation engine will not process the events.
4. If event time is older than 5 minutes from Wizard Time (correct time), events will be directly routed to the database.

## Internal or System Events

Internal or System Events is a means to report on the status and status change of the system. There are two types of events generated by the internal system, they are:

- Internal events
- Performance events

Internal events are informational and describe a single state or change of state in the system. They report when a user logs in or fails to authenticate, when a process is started or a correlation rule is activated. Performance events are generated on a periodic basis and describe average resources used by different parts of the system.

All system events populate the following attributes

- ST (Sensor Type) field: for internal events it is set to 'I' and for performance events it is set to 'P'
- Event ID: a unique UUID for the event
- Event Time: the time the event was generated
- Source: the UUID of the process that generated the event
- Sensor Name: the name of the process that generated the event (for example, DAS_Binary)
- RV32 (Device Category): set to 'ESEC'
- Agent: 'Performance' for performance events and 'Internal' for internal events

In addition to the common attributes, every system event also sets the resource, subresource, the severity, the event name and the message tags. For internal events, the event name specific enough to identify the exact meaning of the event (for example, UserAuthenticationFailed). The message tags adds some specific detail; in the above example the message tag will contain the name of the user, the OS name if available and the machine name). For performance events the event name is generic describing the type of statistical data and the data itself is in the message tag.

Performance events are sent directly to the database. To view them, do a quick query.

See Appendix A for list of System Events.

## Processes

The following processes and Windows service communicate with each other through iSCALE--the message-oriented middleware (MOM).

- Watchdog
- Event Statistics
- Data Synchronizer (Data Controller)
- Correlation Engine
- RuleLg Checker (Correlation Rule Checker)
- Data Access Service (DAS) – binary, query and Active Views™
- Query Manager
- eSecurity Service (MSSQL only) – see Watchdog

The following is the architecture for e-Security Server.

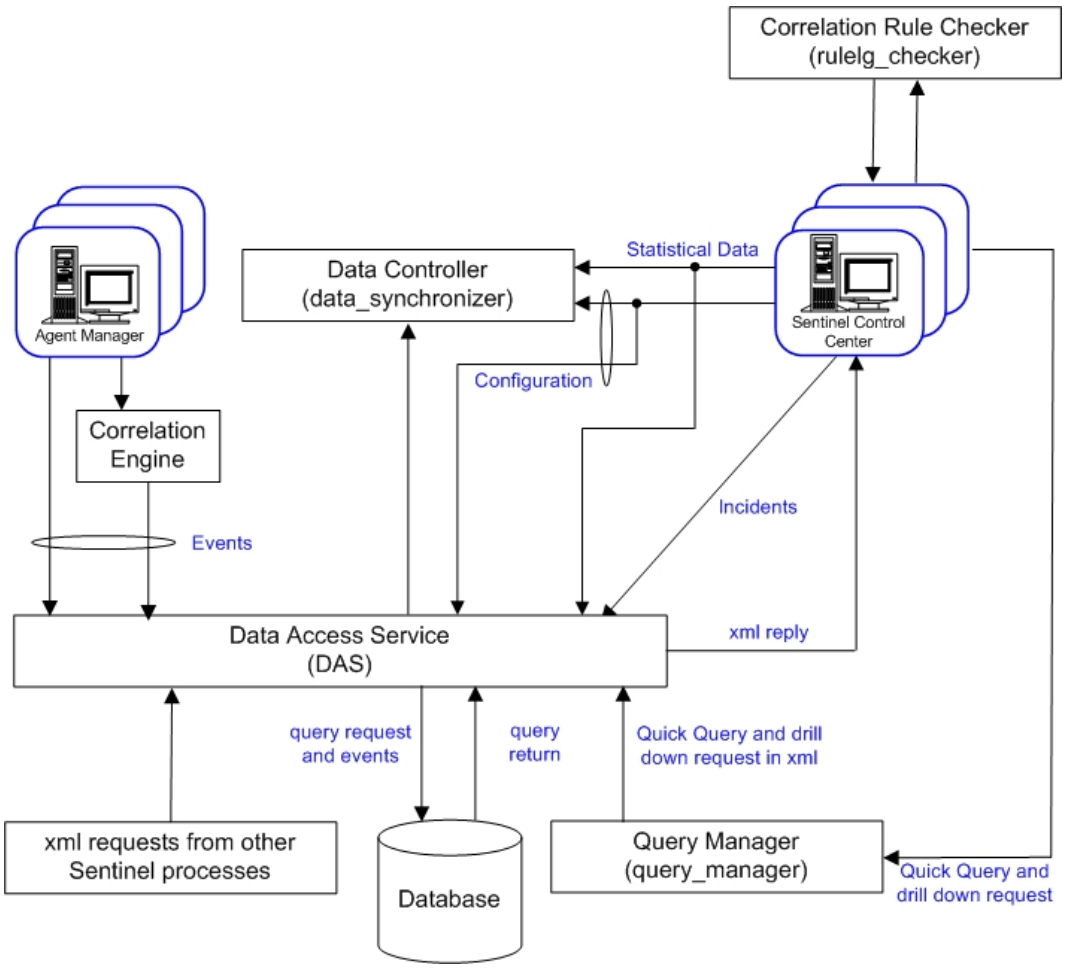


### Watchdog Process

Watchdog is a Sentinel Process that manages other Sentinel Processes. If a process other than Watchdog stops, Watchdog will report this and will then restart that process.

For Windows, watchdog is a service and is called eSecurity. If this service is stopped, it will stop all Sentinel processes on that machine.

### Event Statistics

The Event Statistics engine is a component of the das_binary process. It manages the data used by the Active Views charts and event tables in the Sentinel Control Center.

The engine maintains a set of events and statistical data for each filter and event attribute combination specified in the Active Views wizard. The first time a user creates an Active View with a given filter and event attribute, a new data set is created. This data set contains the counts of that attribute across fixed intervals, as well as the most recent events for each of those intervals. Each data set is configured to hold the most recent 24 hours of data.

Intervals are sent to the Sentinel Control Center after a brief delay, to stabilize the data that might have arrived late due to network delays and time skew.

Active Views are automatically shared by multiple users if the desired event attribute and filter are the same. When an Active View is no longer in use by any user, it will be discarded after an hour period. However, if an Active View is saved in user preferences, it will continue to collect data for up to 100 hours.

**Data Synchronizer Process (Data Controller)**

The Data Synchronizer (data_synchronizer) process manages the modification of configuration data by multiple users. When a user requests to modify data through the Sentinel Control Center, the data record is locked by the data_synchronizer. The details of who locked the data are published to the other active Sentinel Control Centers and no other users may modify that data. If a Sentinel Control Center is closed before it unlocks any data that it has locked, the locks will timeout.

**Correlation Engine Process (correlation_engine)**

The Correlation Engine (correlation_engine) process receives events from the Wizard Agent Manager and publishes correlated events based on user-defined correlation rules.

**RuleLg Checker Process (rulelg_checker)**

The RuleLg Checker (rulelg_checker) process validates the syntax of filter and correlation rule expressions. The Sentinel Control Center uses these results to determine if a filter or a correlation rule can be saved.

**Data Access Service Process (DAS)**

The Data Access Service (DAS) process is Sentinel Server's persistence service and provides an interface to the database. It provides data driven access to the database backend.

DAS is a container, composed of five different processes. Each process is responsible for different types of database operations. These processes are controlled by the following configuration files:

- das_binary.xml: used for event and correlated event insertion operations
- das_query.xml: all other database operations
- activity_container.xml: used for executing and configuring activity service
- workflow_container.xml: used for configuring the workflow (iTRAC) service
- das_rt.xml: used for configuring the Active Views function within the Sentinel Control Console

DAS receives requests from the different Sentinel processes, converts them to a query against the database, processes the result from the database and converts it that back to a reply. It supports requests to retrieve events for Quick Query and Event Drill Down, to retrieve vulnerability information and advisor information and to manipulate configuration information. DAS also handles logging of all events being received from the Wizard Agent Manager and requests to retrieve and store configuration information.

**Query Manager Process (query_manager)**

The query manager (query_manager) process receives quick query and drill down requests from Sentinel Control Center and forwards them to the database through DAS. The requests from Sentinel Control Center define the events

needed from a filter. If a filter is used, the Query Manager retrieves the filter definition and converts the filter to an xml criterion. Query Manager then sends the request to DAS. Not all filters can be completely converted to queries that can be processed by the database. If the filter is fully converted, the Query Manager instructs DAS to send the reply directly to the Sentinel Control Center. If the filter contains regular expressions that cannot be converted to SQL the query manager converts what it can and generates a conservative criterion that returns a superset of the required events. In that case, Query Manager instructs DAS to return the result to the Query Manager. When the reply comes back to the Query Manager it filters it in memory and sends those events that pass the filter to the Sentinel Control Center.

## Logical Architecture

Sentinel 5 is composed of three logical layers:

- collection and enrichment layer
- business logic layer
- presentation layer.

The collection/enrichment layer aggregates the events from external data sources, transforms the device-specific formats into e-Security format, enriches the native events source with business-relevant data and dispatches the event packets to the message bus. The key component orchestrating this function is the Agent, aided by a taxonomy mapping and global filter service.

The business logic layer contains a set of distributable components. The base components are a Remoting service that adds messaging capabilities to the data objects and services to enable transparent data access across the entire network and Data Access service that is an object management service to allow users to define objects using metadata. Additional services include Correlation, Query Manager, Workflow, Event Visualization, Incident Response, Health, Advisor, Reporting and Administration.

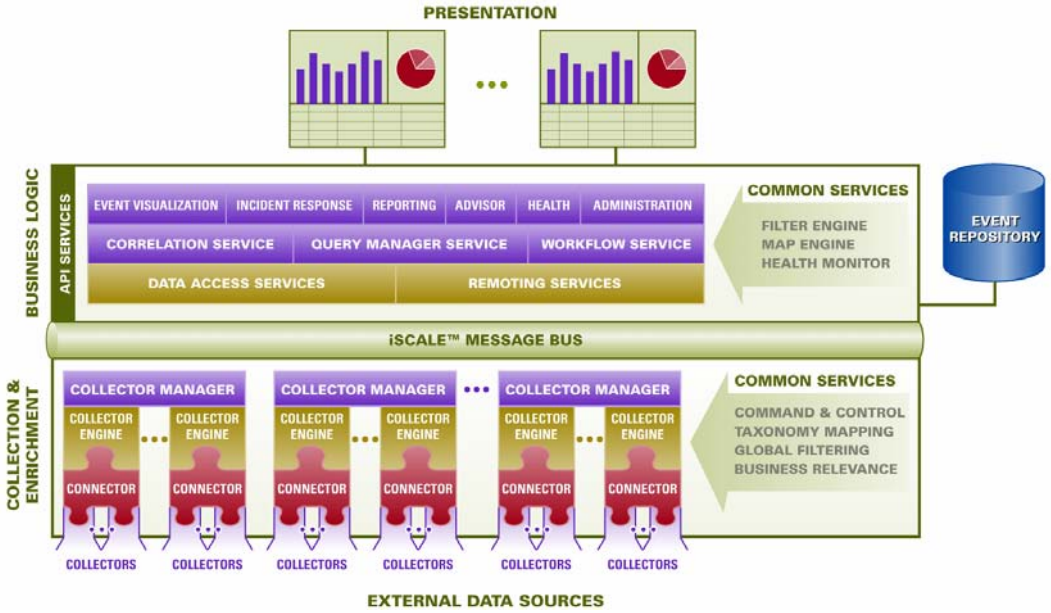

The presentation layer renders the application interface to the end user. A comprehensive dashboard called the Sentinel Control Center offers an integrated user workbench consisting of an array of seven different applications accessible via a single common framework. This cross-platform framework is built on Java™ 1.4 standards and provides a unified view into independent business logic components – real-time interactive graphs, actionable incident response, automated enforceable incident workflow, reporting, incident remediation against known exploits and more.

Each of the layers are illustrated in the figure above and subsequently discussed in detail in the following sections.

## Collection and Enrichment Layer

Events are aggregated using a set of flexible and configurable Agents, which collect data from a myriad of sensors and other devices and sources. User can use pre-built Agents, modify existing Agents or build their own Agents to ensure the system meets all requirements.

Data aggregated by the Agents in the form of events is subsequently normalized and transformed into XML format, enriched with a series of metadata (i.e., data about data) using a set of business relevance services and propagated to the server-side for further computational analysis using message bus platform. The Collection and Enrichment layer consists of the following components:

- Connectors and Agent
- Agent Manager and Engine
- Agent Builder

### Connectors and Agents

A Connector is a concentrator or multiplexed adapter that connects the Agent Engine to the actual monitored devices.

Agents are the component-level aggregator of event data from a specific source. Sentinel 5 primarily supports remote “agent-less” connections to sources; however, Agents can be deployed on specific devices where a remote approach is less efficient.

Agents are controlled from the Sentinel Control Center, which orchestrates the communication between the Collectors and the Sentinel platform for real time analysis, correlation computation and incident response.

### Agent Manager and Engine

Collector Manager manages the Collectors, monitors system status messages and performs event filtering as needed. Main functions of the Collector Manager include transforming events, adding business relevance to events via taxonomy, performing global filtering on events, routing events and sending health messages to the Sentinel server.

A Collector Engine is the interpreter component that parses the Collector code.

### Agent Builder

Agent Builder is a standalone application that is used to build, configure and debug Agents. This application serves as an integrated development environment (or IDE) that allows the user to create new Agents to parse data

from source devices using a special-purpose interpretive language designed to handle the nature of network and security events.

### Common Services

All of the above-described components in this Collection and Enrichment layer are driven by a set of common services. These utility services form the fabric of the data collection and data enrichment and assist in filtering the noise from the information (via global filters), applying user-defined tags to enrich the events information (via business relevance and taxonomy mapping services) and governing the data collectors' functions (via command and control services).

**Taxonomy** - Nearly all security products produce events in different formats and with varying content. For example, Windows and Solaris report a failed login differently.

Sentinel's taxonomy automatically translates heterogeneous product data into meaningful terms, which allows for a real-time homogeneous view of the entire network security. Sentinel Taxonomy formats and filters raw security events before adding event context to the data stream. This process formats all the security data in the most optimal structure for processing by the Sentinel Correlation engine, as you can see in the following diagram.

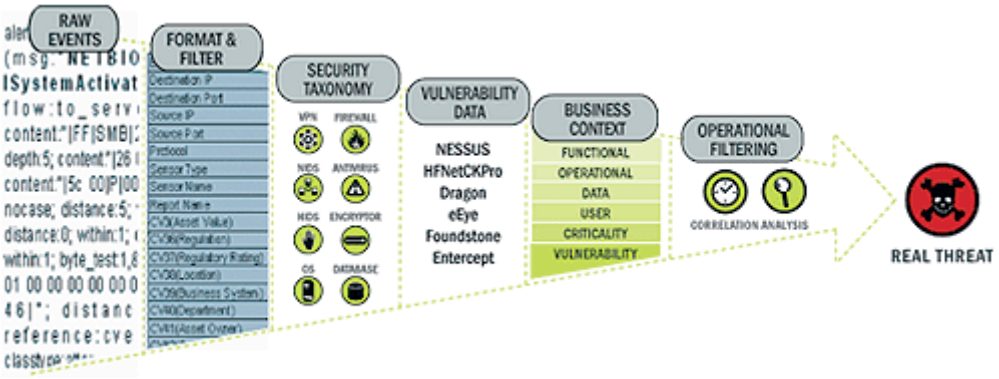


**Business Relevance** - Sentinel 5 injects business-relevant contextual data directly into the event stream. It includes up to 135 customizable fields where users can add in asset specific information such as business unit, owner, asset value, geography. Once this information is added into the system, all other components can take advantage of the additional context.

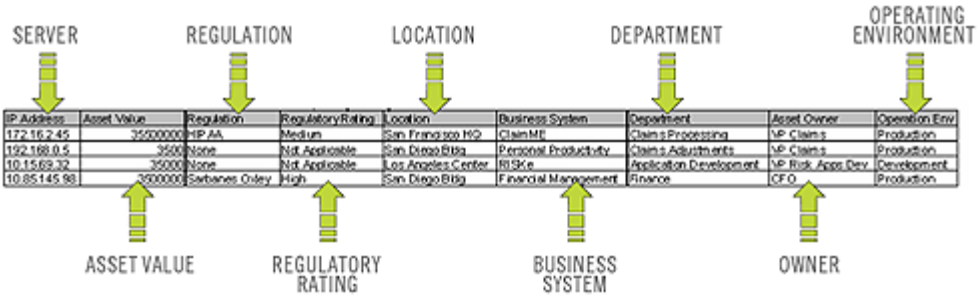


**Exploit Detection** - Exploit Detection enables immediate, actionable notification of attacks on vulnerable systems. It provides a real-time link between IDS signatures and vulnerability scan results, notifying users automatically and

immediately when an attack attempt to exploit a vulnerable system. This dramatically improves the efficiency and effectiveness of incident response.

Exploit Detection provides users with updates of mappings between IDS and vulnerability scanner product signatures. The mappings include a comprehensive list of IDS and vulnerability scanners, Users simply upload vulnerability scan results into Sentinel. Exploit Detection automatically parses them and updates the appropriate IDS Collectors. It uses the embedded knowledge of vulnerability status to efficiently and effectively prioritizes responses to security threats in real time.

When an attack is launched against a vulnerable asset, Exploit Detection alerts users with the corresponding severity level of the exploited vulnerability. Users can then take immediate action on high-priority events. This takes the guesswork out of alert monitoring and increases incident response efficiency by focusing reaction on known attacks against vulnerable assets.

Exploit Detection also enables users to map or “un-map” signatures and vulnerabilities to tune out false positives and negatives and to leverage custom signatures or vulnerability scans.

## Business Logic Layer

The kernel of the Sentinel 5 platform consists of a set of loosely-coupled services that can run in a standalone configuration or in a distributed topology. This service-oriented architecture (SOA) is called iSCALE. Specifically, Sentinel’s SOA comprises a set of engines, services and APIs working together for linear scaling of the solution against increasing data load and/or processing workload.

Sentinel services run in specialized containers and allow unparalleled processing and scaling because they are optimized for message-based transport and computation. The key services that make up the Sentinel Server include:

- Remoting Service
- Data Access Service
- Query Manager Service
- Correlation Service
- Workflow Service
- Event Visualization
- Incident Response
- Reporting
- Advisor
- Health
- Administration

### Remoting Service

Sentinel 5’s Remoting Service provides the mechanism by which the server and client programs communicate. This mechanism is typically referred to as distributed object application.

Specifically, the Remoting Service provides:

- Locate remote objects: This is achieved through metadata that describes the object name or registration token, although the actual location is not required, since the iSCALE message bus allows for location transparency.
- Communicate with remote objects: Details of communication between remote objects are handled by the iSCALE message bus.

- Object streaming and chunking: When large amounts of data need to pass back and forth from the client to the server, these objects are optimized to load the data on demand.
- Callbacks: Another pattern and layer of abstraction built into the Remoting Service that allows for PTP remote object communication.
- Service monitoring and statistics: This provides performance and load statistics for usage of these remote services.

**Data Access Service**

Data Access Service (DAS) is an object management service, which allows users to define objects using metadata. DAS manages the object and access to objects and automates transmission and persistence. DAS also serves as a facade for accessing data from any persistent data store such as databases, directory services or files. The operations of DAS include uniform data access via JDBC and optionally high-performance event insert strategies using native connectors (i.e., OCI for Oracle 9i and ADO for Microsoft SQL Server).

**Query Manager Service**

The Query Manager Service orchestrates drill-down and event history requests from the Sentinel Control Center. This service is an integral component for implementing the paging algorithm used in the Event History browsing capability. It converts user-defined filters into valid criteria and appends security criteria to it before events are retrieved. This service also ensures that the criteria do not change during a paged event history transaction.

**Correlation Service**

Sentinel 5's correlation algorithm computes correlated events by analyzing the data stream in real time. It publishes the correlated events based on user-defined rules before the events reach the database. Rules in the correlation engine can detect a pattern in a single event of a running window of events. When a match is detected, the correlation engine generates a correlated event describing the found pattern and may create an incident or trigger a remediation workflow via iTRAC. The correlation engine works with a rules checker component which computes the correlation rule expressions and validates syntax of filters. In addition to providing a comprehensive set of correlation rules, e-Security's correlation engine provides specific advantages over database-centric correlation engines.

- By relying on in-memory processing rather than database inserts and reads, the correlation engine performs during high steady-state volumes as well as during event spikes when under attack, the time when correlation performance is most critical.
- Correlation volume does not slow down other system components, so the user interface remains responsive, especially with high event volumes.
- Distributed correlation – Organizations can deploy multiple correlation engines, each on its own server, without the need to replicate configurations or add databases. Independent scaling of components provides cost-effective scalability and performance.
- The correlation engine can add events to incidents after an incident has been determined.

Users are encouraged to measure a metric called Event Rules per Second (ERPS). ERPS is the measure of the number of events that can be examined by a correlation rule per second. This measure is a good performance indicator as it estimates the impact on performance when two factors intersect: events per second and number of rules in use.

### Workflow Service (iTRAC)

The Workflow Service receives triggers on incident creation and initiates workflow processes based on pre-defined workflow templates. It manages the lifecycle of these processes by generating work items or executing activities. This service also maintains a history of completed processes that may be used for auditing incident responses.

### Event Visualization

Active Views™, the interactive graphical user interface for event visualization, provides one integrated, security management dashboard with a comprehensive set of real-time visualization and analytical tools to facilitate threat detection and analysis. Users can monitor events in real time and perform instant drill-downs from seconds to hours in the past. A wide array of visualization charts and aids allow monitoring of information via 3D bar, 2D stacked, line and ribbon chart representation and others. Additional valuable information can be viewed from the Active Views dashboard, including notification of asset exploits (exploit detection), viewing asset information and graphical associations between pertinent source IPs and destination IPs.

Because Active Views uses the iSCALE architecture, analysts can quickly drill down for further analysis because Active Views provides direct access to the real-time memory-resident event data, which easily handles thousands of events per second without any performance degradation. Data is kept in memory and written to the database as needed (Active Views can store up to 8 hours of data in memory with typical event loads). This uninterrupted, performance-oriented real-time view is essential when under attack or in steady-state.

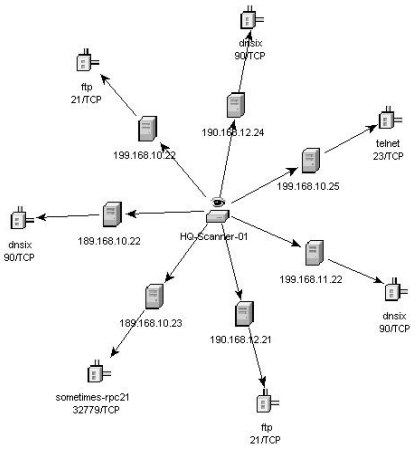

### Incident Response Through iTRAC

iTRAC transforms traditional security information management from a passive "alerting and viewing" role to an "actionable incident response" role by enabling organizations to define and document incident resolution processes and then guide, enforce and track resolution processes once an incident or violation has been detected.

Sentinel 5 comes with "out-of-the-box" process templates that use the SANS Institute's guidelines for incident handling. Users can start with these pre-defined processes and configure specific activities to reflect their organization's best practices. iTRAC processes can be automatically triggered from incident creation or correlation rules or manually engaged by an authorized security or audit professional. iTRAC keeps an audit trail of all actions to support compliance reporting and historical analysis.

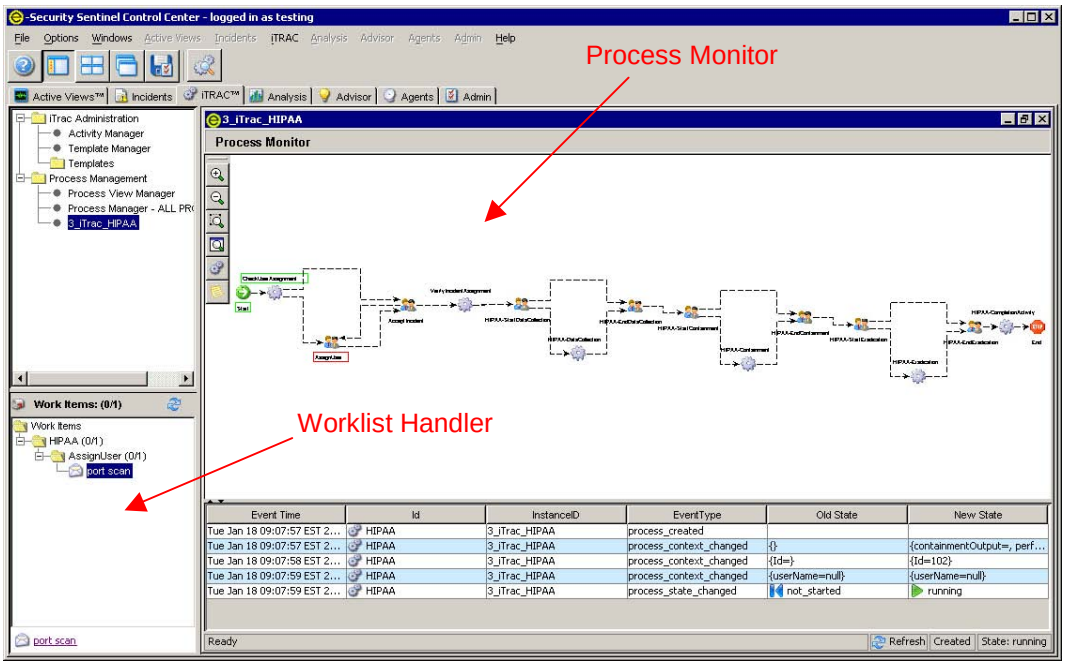


A worklist provides the user with all tasks that have been assigned to the user and a process monitor provides real-time visibility into process status during a resolution process lifecycle.

iTRAC's activity framework enables users to customize automated or manual tasks for specific incident-resolution processes. The iTRAC process templates can be configured using the activity framework to match the template with an organization's best practices. Activities are executed directly from the Sentinel Control Center.

iTRAC's automation framework works using two key components – the activity container and the workflow container. The former automates the activities execution for the specified set of steps based on input rules and the latter automates the workflow execution based on activities via a work-list. The input rules are based on the XPDL (XML Processing Description Language) standard and provide a formal model for expressing executable processes in a business enterprise. This standards-based approach to the implementation of business-

specific rules and rule sets ensures future-proofing of process definitions for customers.

**Reporting Service**

The Reporting service allows for reporting, including historical and vulnerability reports. Sentinel 5 comes with out-of-the-box reports and enables users to configure their own reports using Crystal Reports. Some examples of reports included with Sentinel 5 are:

- Trend analysis
- Security status of lines of business or critical assets
- Attack types
- Targeted assets
- Response times and resolution
- Policy compliance violations

**Advisor**

Sentinel Advisor, an optional module, cross-references Sentinel's real-time alert data with known vulnerabilities and remediation information, bridging the gap between incident detection and response. With Advisor, organizations can determine if events exploit specific vulnerabilities and how these attacks impact their assets. Advisor also contains detailed information on the vulnerabilities that attacks intend to exploit, the potential effects of the attacks if successful and necessary steps for remediation. Recommended remediation steps are enforced and tracked using iTRAC incident response processes.

**Health**

The Health service enables users to get a comprehensive view of the distributed Sentinel 5 platform. It aggregates health information from various processes that are typically distributed on various servers. The health information is periodically displayed on the Sentinel Control Center for the end user.

**Administration**

The Administration facility allows for user management and settings setup facilities typically needed by application administrators of Sentinel 5.

**Common Services**

All of the above described components in this business logic layer of the architecture are driven by a set of common services. These utility services assist in fine-grain filtering (via Filter Engine) of events to users, continuous monitoring of system health statistics (via Health Monitor) and dynamic updates of system wide data (via Map Service). Together, these utility services form the fabric of the loosely-coupled services that allow for unparalleled processing and scaling over the message bus-based transport for real-time analytics and computation.

## Presentation Layer

The presentation layer renders the application interface to the end user. The Sentinel Command Center is a comprehensive dashboard that presents information to the user.

# Product Modules

## Sentinel Control Center

The Sentinel Control Center provides an integrated, powerful security management dashboard. Intuitive displays enable analysts to quickly identify new trends or attacks, manipulate and interact with real-time graphical information and respond to incidents. Key features include:

- Active Views – Real-time analytics and visualization
- Incidents – Incident creation and management
- Analysis – Correlation rules definition and management
- iTRAC – Process management for documenting, enforcing and tracking incident resolution processes.
- Reporting – Historical reports and metrics

## Sentinel Wizard

Sentinel Wizard collects data from source devices and delivers a richer event stream by injecting taxonomy, exploit detection and business relevance into the data stream before events are correlated and analyzed and sent to the database. A richer event stream means that data is correlated with the required business context to identify and remediate internal or external threats and policy violations. In any configuration, there may be one or more Wizards deployed, providing customers with the ability to deploy product components into their infrastructure based on their network topology.

## Sentinel Advisor

Sentinel Advisor, an optional module, cross-references Sentinel's real-time alert data with known vulnerabilities and remediation information.

# Contents

This guide contains the following:

## Conventions Used

### Notes and Cautions

**NOTE:** Notes provide additional information that may be useful.

**CAUTION:** Cautions provide additional information that may keep you from performing damage or loss of data to your system.

### Commands

Commands appear in courier font. For example:

```
useradd –g dba –d /export/home/oracle –m –s /bin/csh
   oracle
```

## Other e-Security References

The following manuals are available with the e-Security install CDs.

- Sentinel™ 5 Install Guide
- Sentinel™ User's Guide
- Sentinel™ 5 Wizard User's Guide
- Sentinel™ 5 User Reference Guide
- Sentinel™ 5 3rd Party Integration Guide
- Release Notes

## Contacting e-Security

- For Technical Support, email at mailto:support@esecurity.net
- For information, email at mailto:info@esecurity.net
- Website: http://www.esecurity.net
- For 24x7 support, call Technical Support directly at 800-474-3131

# Chapter 2 – Navigating Sentinel Control Center

The Sentinel Control Center is made up of:

- Menu bar
- Toolbar
- Tabs

In addition, the following is discussed in this chapter:

- Starting the Sentinel Control Center
- Changing the Sentinel Control Center's Look
- Saving User Preferences
- Changing Sentinel Password

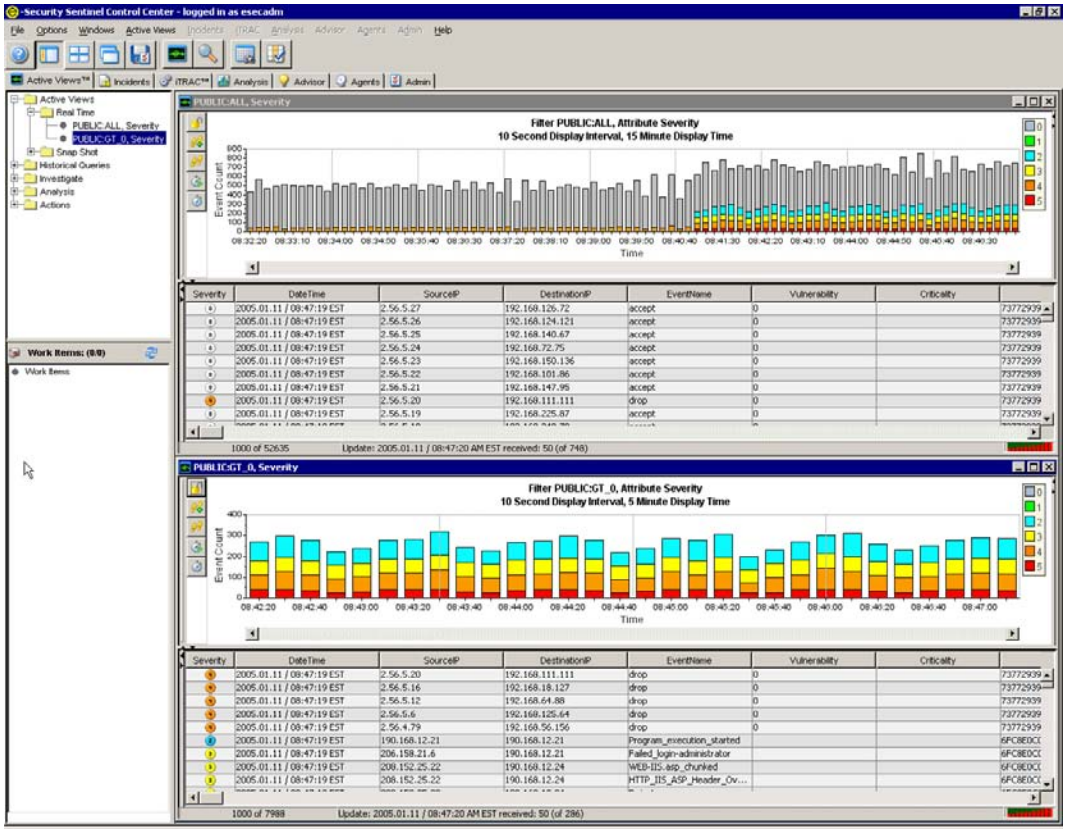


## Starting the Sentinel Control Center

### Starting the Sentinel Control Center in Windows

Starting the Sentinel Control Center in Windows

1. Click Start > e-Security > Sentinel Control Center or click the Sentinel Control Center on the desktop.
2. Enter your username, password and click OK.

### Starting the Sentinel Control Center in UNIX

Starting the Sentinel Control Center in UNIX

1. As user esecadm, cd to:

   ```
   $ESEC_HOME/sentinel/console
   ```

2. Run the following command:

   ```
   ./run.sh
   ```

3. Enter your username, password and click OK.

## Menu Bar

Beneath the title bar are ten menus. From the left across the top of the window they are File, Options, Windows, Active Views, Incidents, iTRAC, Advisor, Agents, Admin and Help.

File, Options, Windows and Help options are always available. Other options are available, depending on which tab is active and what permissions you have.

### File Menu

- Save Preferences
- Exit

### Options Menu

- Change Password
- Tab Placement
  - Top
  - Bottom
- Dock Navigator
- Show Navigator

### Windows Menu

- Cascade All
- Title All
  - Tile Best Fit
  - Tile Horizontal
  - Tile Vertical
- Minimize All
- Restore All
- Close All

### Active Views™

- Properties
- Create Active View
- Event Query

### Incidents

- Display Incident View Manager
- Create Incident
- Attachment Viewer Configuration

### iTRAC™

- Display Process Manager

### Analysis

- Create Report

### Advisor

- Create Report

### Agents

- Display Agent view Manager

### Admin

- Reporting Configuration
- Correlation Rules
- Correlation Engine Manager
- Global Filter Configuration
- Menu Configuration
- Filter Configuration
- User Configuration

### Help

- Help
- About Sentinel

## Tool Bar

Five system-wide toolbar buttons are always displayed. Other buttons display depending on which tab or window you have active and on user permissions.

### System-Wide Toolbar

The five system-wide toolbar buttons are:

-  View Sentinel Help
-  Show/Hide Navigation Window
-  Tile All Display Windows
-  Cascade All Display Windows
-  Save User Preferences

### Active Views™ Tab

When the ActiveViews™ Tab is active, the following is available.

-  Active Views
-  Launch Event Query

### Event Count Over Time Window

When an Event Counts Over Time window is active, the following is available.

-  Snapshot of an Event Count Over Time Table
-  Manage Columns of Event Count Over Time Table

### Events Counts Over Time Chart

When the Events Counts Over Time Chart is active, the following is available within the Events Counts Over Time Chart.

- Lock/Unlock the Chart
- Increase Display Interval
- Decrease Display Interval
- Increase Display Time
- Decrease Display Time

When you click the Lock button, the available buttons are:

- Lock/Unlock the Chart
- Increase Display Interval
- Decrease Display Interval
- Increase Display Time
- Decrease Display Time
- Zoom In
- Zoom Out
- Drill Down to Events
- Save as a html File

### Snapshot Window

When the Snapshot Window is active, the following is available.

-  Manage Columns

## Incidents Tab

When the Incidents tab is active, the following is available.

-  Display Incident View Manager
-  Create a New Incident
-  Configure Attachment Viewers

### Incident

When an Incident is open, the following is available.

-  Manage Columns of Associated Events

### iTRAC

When the iTRAC tab is active, the following is available.

-  Displays Process View Manager

### Analysis and Advisor Tab

When either the Analysis or Advisor tab is active, the following is available.

-  Create Report

### Agents Tab

When the Agents tab is active, the following is available.

-  Displays Agent Manager View Manager
-  Displays Agent View Manager

### Admin Tab

When the Admin tab is active, the following is available.

-  Display Reporting Configuration
-  Display Correlation Engine Manager
-  Display Menu Configuration
-  Display User Manager
-  Display Correlation Rule
-  Display Global Filter Configuration
-  Display Filter Manager
-  Server View Manager (Linux only)

#### Filter Manager Window

When the Filter Manager window is active, the following is available.

-  Create a New Filter
-  Delete the Selected Filter (active when a filter is selected)

#### Menu Configuration Menu

When the Menu Configuration window is active and in modify mode, the following is available.

-  Create New Menu Item
-  Delete Menu Item
-  Activate Menu Item
-  Deactivate Menu Item

## Tabs

Depending on your user permissions, your Sentinel Control Center will display the following tabs. You must have the specific permission to view each tab.

- Active Views™
- Incidents
- iTRAC™
- Analysis
- Advisor
- Agents
- Admin

For more information about Tabs, see the individual chapters for each tab.

## Changing the Sentinel Control Center's Look

You can change the Sentinel Control Center's look by:



### Setting the Tab Position

To set the tab position

1. Click Options > Tab Placement.
2. Select either Top or Bottom.

### Showing or Hiding the Navigator window

To show or hide the Navigator window

1. Click Options > Show Navigator on or off.

### Docking or Floating the Navigator window

To dock or float the Navigator window

1. Click Options > Dock Navigator on or off.

### Cascading Windows

To cascade windows

1. Click Windows > Cascade All. All open windows in the right panel will cascade.

### Tiling Windows

To Tile Windows

1. Click Windows > click Tile All.
2. Point to either:
   - Tile Best Fit
   - Tile Vertical
   - Tile Horizontal

### Minimizing and Restoring All Windows

To minimize all windows

1. Click Windows > Minimize All. All open windows in the right panel will minimize.

### To restore all windows to original size

To restore all windows to original size

1. Click Windows > Restore All. All open windows in the right panel will restore to their original size.

### To restore an individual window

To restore an individual window

1. Click the minimized window. The window will restore to its original size.

### Closing All Open Windows at Once

To close all windows

1. Click Windows > Close All.

## Saving User Preferences

You must have the user permission Save Workspace.

Preferences that can be saved are:

- Permanent windows, those that can be recreated because they are not dependent on data that was available at the time of their original creation. For example, Summary displays and Active Views can be saved. However, temporary windows, such as snapshots and quick queries cannot be saved. All the windows listed in the Admin Navigator are saved, but none of the secondary windows you open by double-clicking a selection in one of those windows is saved.
- Window positions

- Window sizes, including the application window
- Tab positions
- Navigator docked or floating and showing or hidden

To save your preferences

1. Click File > click Save Preferences or click the Save Preferences button.

## Changing Sentinel Control Center Password

**NOTE**: In order to meet stringent security configurations required by Common Criteria Certification, e-Security requires a strong password with the following characteristics:

1. Choose passwords of at least 8 with characters in length that includes at least one UPPER CASE, one lower case, one special symbol (!@#$%^&*()_+), and one numeric (0-9).

2. Your password may not contain your e-mail name or any part of your full name.

3. Your password should not be a "common" word (for example, it should not be a word in the dictionary or slang in common use).

4. Your password should not contain words from any language, because numerous password-cracking programs exist that can run through millions of possible word combinations in seconds.

5. You should choose a password you can remember and yet is complex. For example, Msi5!YOld (My son is 5 years Old) OR IhliCf5#yN (I have lived in California for 5 years now).

To change your Sentinel Control Center password

1. Click Options > click Change Password.
2. Enter the old password.
3. Enter the new password and re-enter the new password to verify it.

**NOTE**: e-Security recommends as a best practice a minimum password length of 8 characters that includes alphanumerics.

4. Click OK.

# Chapter 3 – Active Views™ Tab

You must have the proper permission to use Active Views™ tab. If this permission is not assigned, none of the permissions related to actions using this tab will be available.

In the Active Views tab, you have the ability to monitor, near real time, events as they are happening and perform queries on these events. You can monitor them in a table form or though 3D bar, 2D stacked, Line or ribbon chart representation.

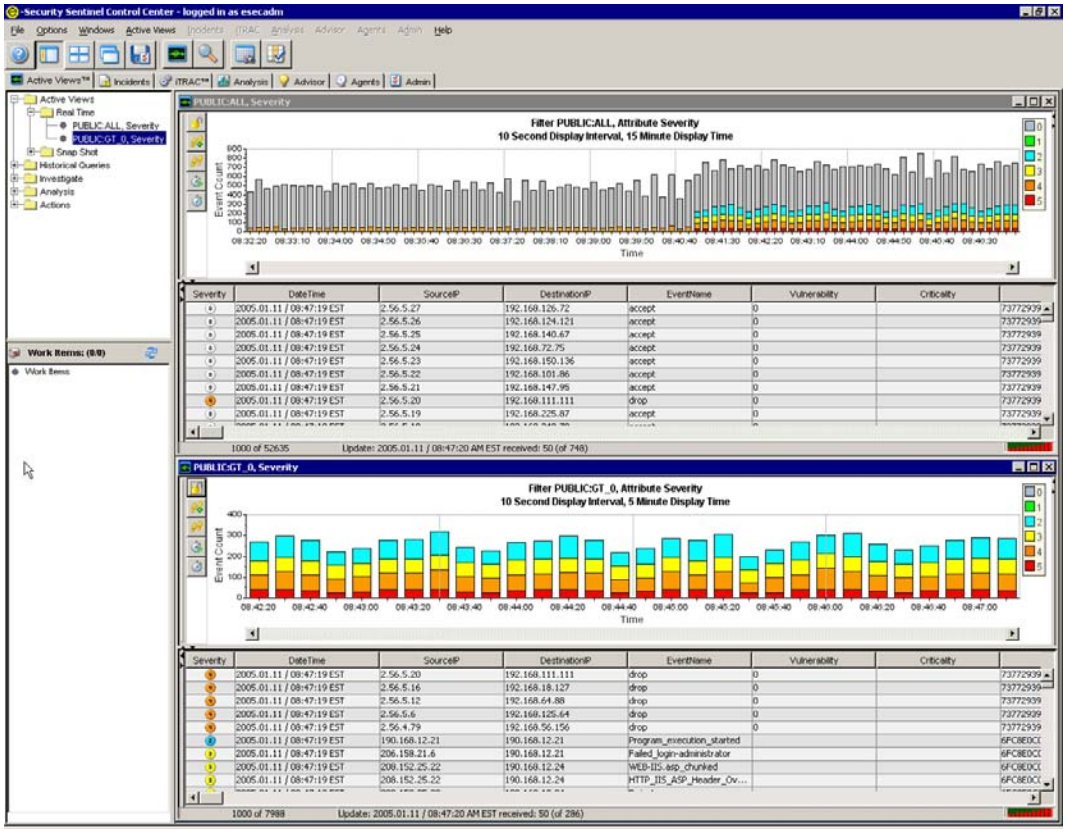


## Active Views Tab - Description

Event views are formatted as tables. Active view configuration is determined by the das_rt.xml file. A near Real Time Event Table with graphical presentation and Snapshot are the two types of Active Views.

- Near Real Time Event Table
  - Holds up to 750 events per 30-second period.
  - By default, the client maintains a 24-hour period of cached events. This is configurable through Active View Properties.
  - By default, the event table will display a maximum of 30,000 events. This is configurable through Active View Properties.
  - By default, the event table refreshes every 30 seconds (send time delay). This represented by a gray line in the event table.

| | | | | |
|---|---|---|---|---|
| 3 | 2005.06.21 / 06:34:38 EDT | | | Threshold_ex |
| 2 | 2005.06.21 / 06:34:38 EDT | 206.158.21.6 | 192.168.10.1 | Password_ex |
| 2 | 2005.06.21 / 06:34:28 EDT | 190.168.12.21 | 190.168.12.21 | Program_exe |

In the event when there are more than 750 per 30-second time period, a red separation line will appear indicating that there are more events than what is displayed.

| | | | | |
|---|---|---|---|---|
| 3 | 2005.06.21 / 07:07:00 EDT | 172.16.112.50 | 172.16.0.65 | unsuccessfu |
| 3 | 2005.06.21 / 07:07:00 EDT | 172.16.112.50 | 172.16.0.65 | suspicious-fil |
| 3 | 2005.06.21 / 07:06:58 EDT | 172.16.112.50 | 172.16.0.65 | successful-a |

  - When saving user preferences, it will continue to collect data for 4 days. For instance, if you save your preferences, log out and log back in the following day your Active View will display any data as if you never logged off.
  - If an Active View is created and not saved, it will continue to collect data for an hour. Within that hour time frame if an identical Active View is created, the Active View will display data for the last hour.
- Snapshot - time-stamped views of a Real Time Event View table.

The following is what makes an Active View unique.

- Filter assigned to an Active View
- The z-axis attribute
- The security filter assigned to a user

The Active Views Tab allows you to:

- Reconfigure Active Views
- Add Events to an incident
- Close a Snapshot or a Visual Navigator Window
- Create an incident
- Custom Menu Options with Events
- Delete a Snapshot or a Visual Navigator Window
- Event Query
- Graph Map
- View Advisor Data
- Manage Columns
- Send messages about Events by e-mail
- Show or Hide Event Details
- Snapshot of a Visual Navigator Window
- View Events that triggered a correlated event
- View Vulnerability Visualization
- View Asset Data
- Perform HP – OpenView Operations and Service Desk
- Perform Remedy Operations

As a user, you can change values (column names) to display logical names and have it populate throughout the system. You can apply attributes to the event stream that are relevant to your business. For more information, see Chapter 10 - Sentinel Data Manager, the Wizard User's Guide and e-Security User Reference Guide.

## Reconfigure Active Views Maximum Event and Cache Value

Active View Properties allows you to configure the maximum number events that can be display in an Active Views and the cached time in each client. The default maximum total number of events in an Active View is 30,000 events. The default cache time value in an Active View is 24 hours.

To reconfigure Active Views Maximum Events and Cache Value

1. Click the Active Views tab.
2. Click Active Views > Properties.

3. Make your changes.

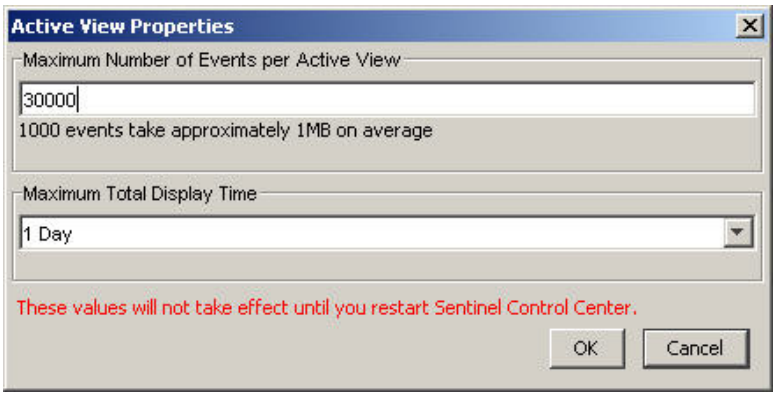


The new values will not take effect until you restart the Sentinel Control Center.

## To View Real Time Events

### To View Real Time Events

1. Click the Active Views tab.
2. Click Active Views > Create Active View or click the Create Active View Button.

3. In the Event Visualization Wizard window, click the down arrows to select your Z-axis, Filter and to Display Events (Yes or No).

**NOTE**: In the filter selection window you can build your own filter or select one of the already built filters. Selecting the 'All' filter will allow all events to appear in your window. When creating an Active View, if the filter assigned to the Active View is changed or deleted after creation of the Active View, the Active View is unaffected.

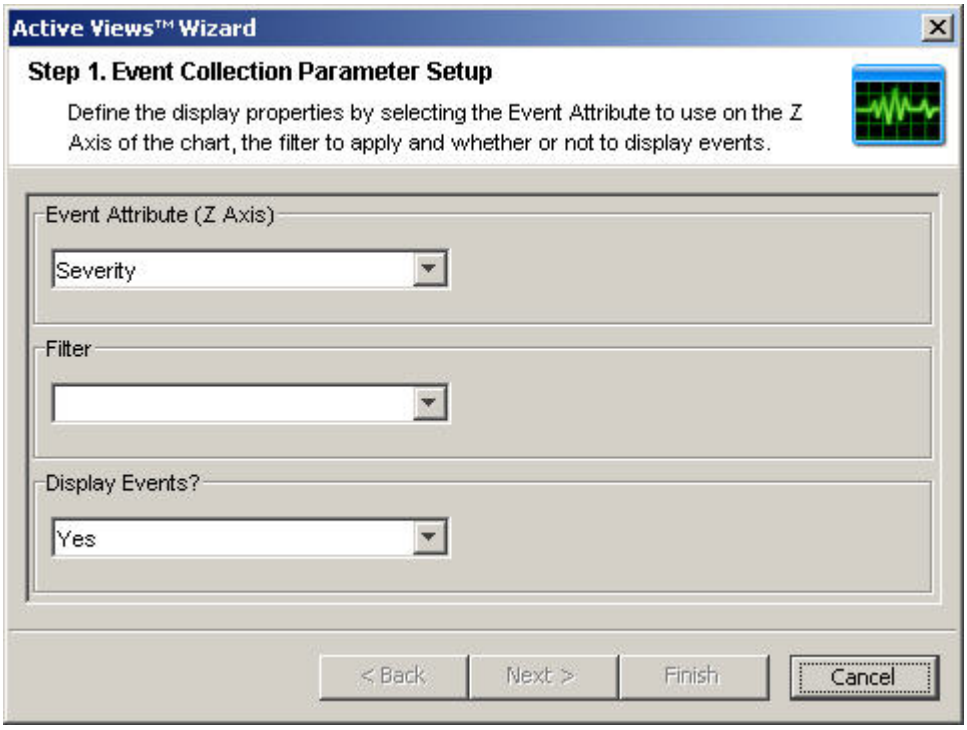


After making your selection, you can click Next or Finish. If you select Finish, the following default values will be chosen:

- Display and Refresh rate of 30 seconds
- Display time of 15 minutes
- Y-axis as Event Count
- Chart type – Stacked Bar 2D

4. If you click Next, Click the down arrows to select your:
   - Display and Refresh rate – number of seconds for event rate to be updated
   - Display time – amount of time to display the chart
   - Y-axis – either total Event Count or Event Count per Second

   Click Next.
5. Select your chart type. Click Next.
   - Chart type - 3D bar, 2D stacked, Line or ribbon charts
6. In addition to your filter selection, you may further refine your event table. You have the option conditions of:
   - None
   - is exactly
   - is not
   - is < (is less than)
   - is <= (is less than or equal to)
   - is > (is greater than)
   - is >= (is greater than or equal to)
   - contains
   - doesn't contain
   - is empty
   - is not empty

   After you create your criteria, click the 'Add to list' button. Click Finish.

**NOTE**: After creating your view, you can edit or remove this refinement to the event table by right-clicking in the graph area and selecting properties. For more information, see To Reset Parameters, Chart Type or Event Table of an Active View.

You graph will look similar to:

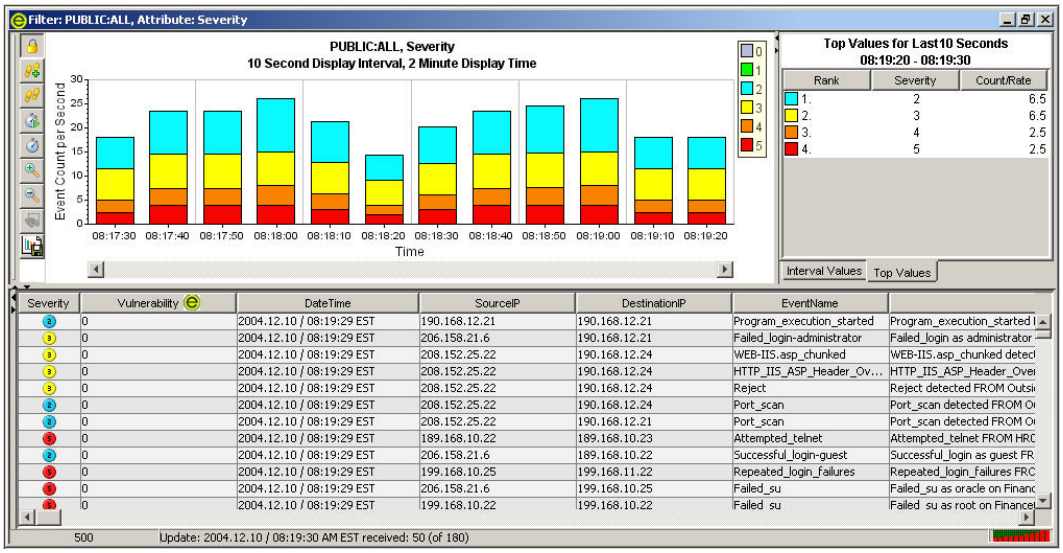


**NOTE**: Active View Properties – Refine Event Table will not affect the graphical representation.

The five buttons to the left of the chart perform the following functions:

- Lock/Unlock the Chart – used when performing a drill-down, zoom in, zoom out, zoom to selection and saving a chart as an html file.

- Increase Display Interval – increases the display time interval for incoming events

- Decrease Display Interval – decreases the display time interval for incoming events

- Increase Display Time – increase the time interval along the x-axis

- Decrease Display Time – decreases the time interval along the x-axis

When you click the Lock button, additional available buttons are:

- Lock/Unlock the Chart – used when performing a drill-down, zoom in, zoom out, zoom to selection and saving a chart as an html file.

- Zoom In – zooms in without changing any of the time settings of the chart

- Zoom Out – zooms out without changing any of the time settings of the chart

- Zoom to Selection – zooms in on a selection of time intervals of events.

- Save navigator details as an html file with chart as images and events in a tabular format.

## To Reset Parameters, Chart Type or Event Table of an Active View

When viewing an Active View, you can reset your chart parameters, change your chart type and if there are events of interests you can filter out other events versus creating a new Active View and filter.

### To Reset Parameters, Chart Type or Event Table of an Active View

1. Within an Active View displaying a chart, right-click and select Properties.

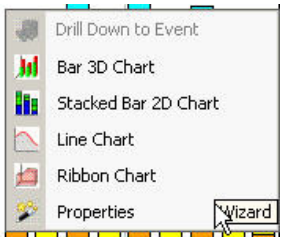


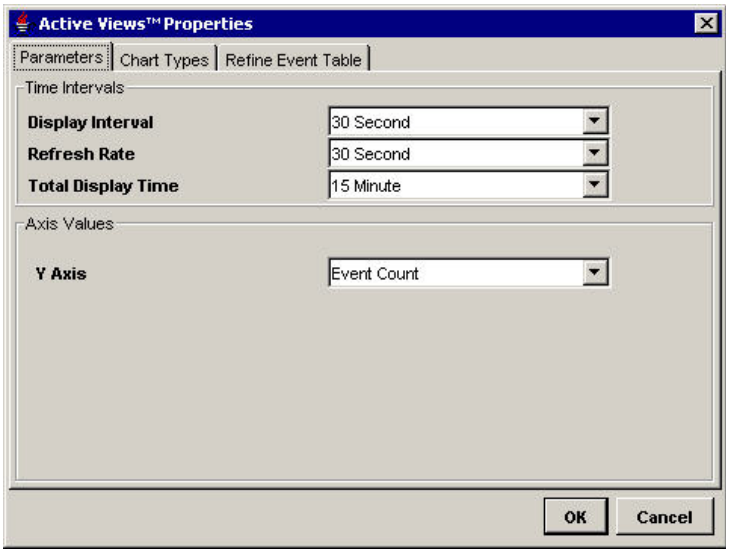


Under the Parameters tab, you can set:

- Display Interval – time between each interval
- Refresh Rate – number of seconds for event rate to be updated
- Total Display Time – amount of time to display the chart
- Y-axis – either total Event Count or Event Count per Second

Under the Chart Types tab, you can set your chart to 3D bar, 2D stacked, Line or ribbon charts.

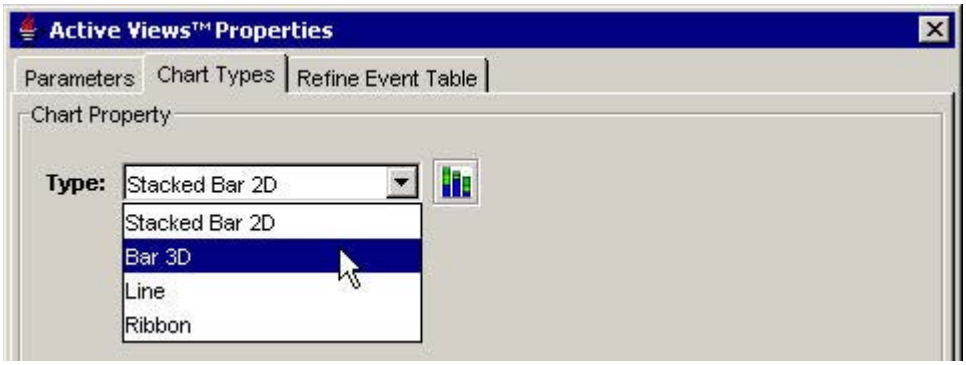


Under the Refine Event Table you can filter Event Field within your Active View.

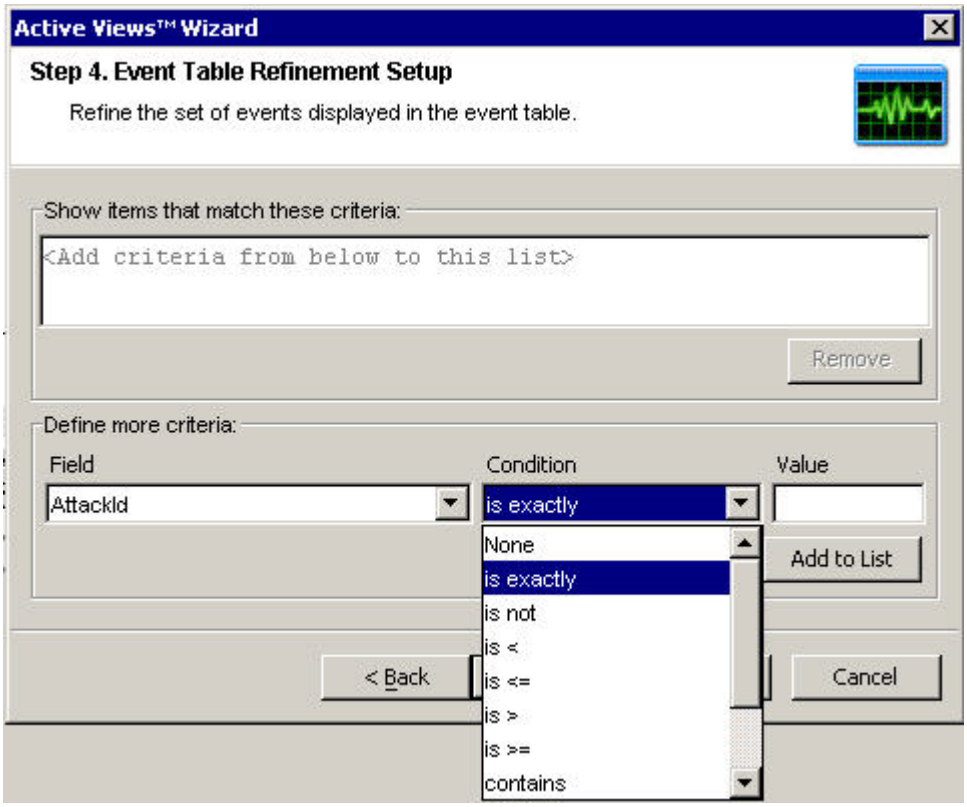


For example you may filter in events with a specific entry in field, such as DeviceAttackName is exactly Back_Door_Probe (TCP 3128). This will result in an Event table with events that only contain DeviceAttackName equivalent to Back_Door_Probe (TCP 3128).

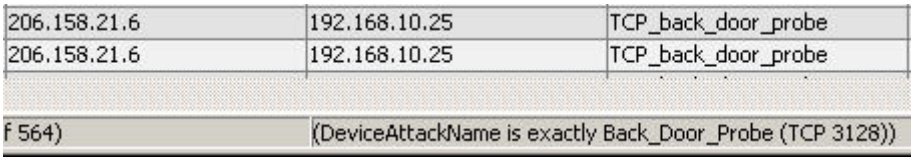


When refining an event table, you see your filter criteria in the bottom right of the events table.

### Rotating a 3D Bar or Ribbon Chart

To rotate a 3D bar or ribbon chart

1. Click anywhere on the chart and hold down the mouse button.
2. Reposition the chart as desired by moving the mouse while holding the button down.

## Showing and Hiding Event Details

To show event details

1. In an Event Real Time table of the Visual Navigator or Snapshot, double-click or right-click an event and click Show Details. An event details will display in the left panel of the Event Real Time table.

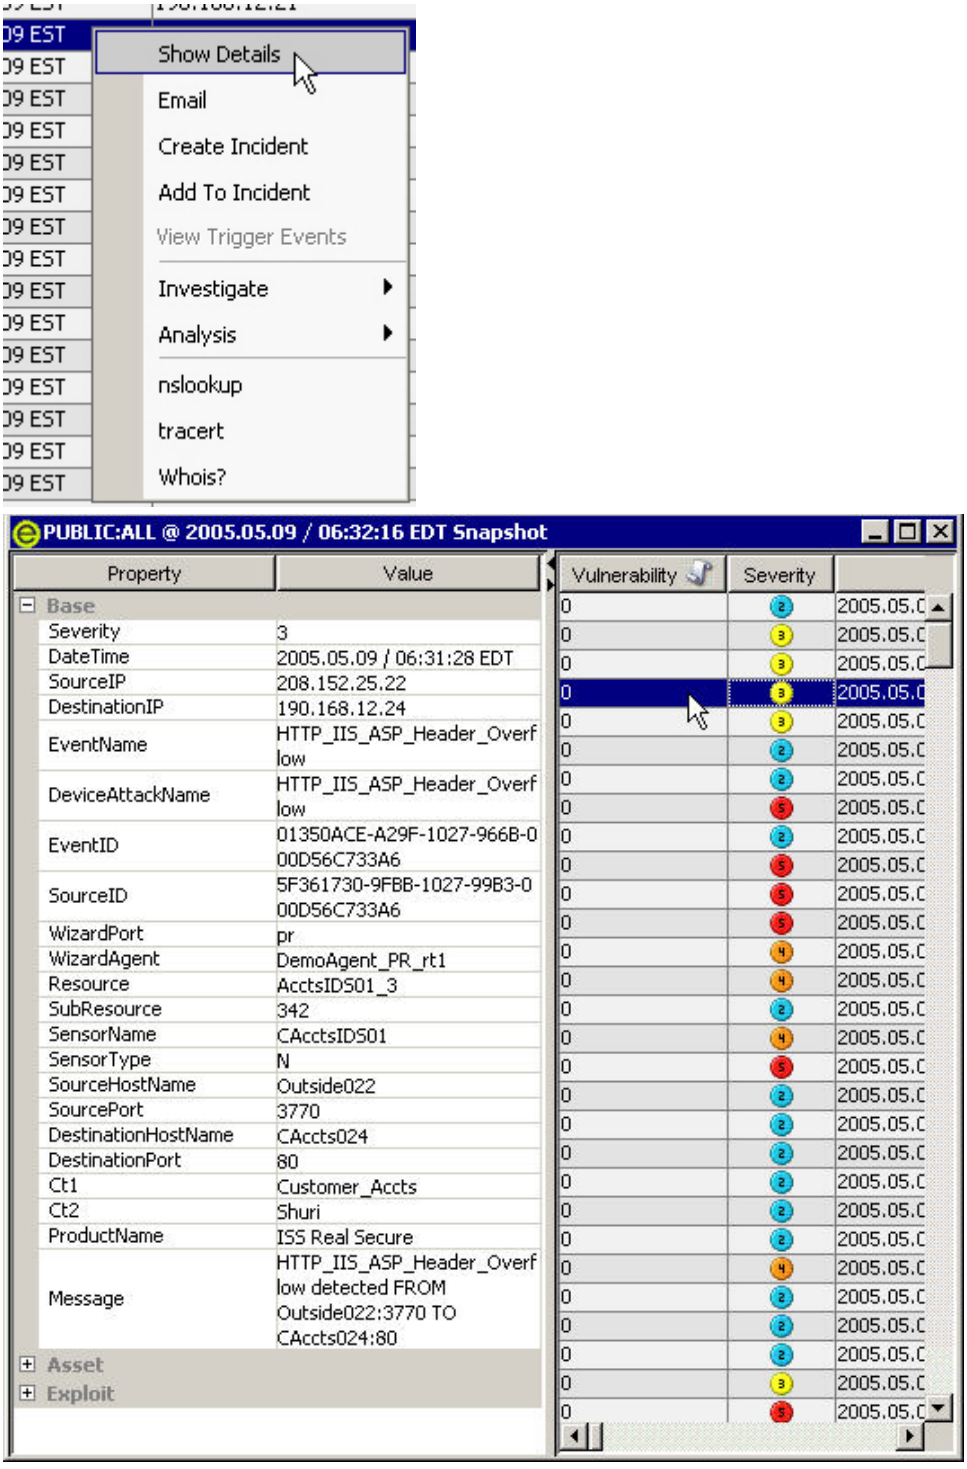


2. If you want details to show the next time you open the Sentinel Control Center, click File > Save Preferences or click the Save User Preference button.

To hide an event detail

1. In an Event Real Time table of the Visual Navigator or Snapshot, with event details displayed in the left panel, right-click an event and click Show Details. The event details window will close.
2. If you don't want details displayed the next time you open the Sentinel Control Center, click File > Save Preferences or click the Save User Preference button.

## Sending Messages about Events and Incidents by e-Mail

Ability to send emails is set in the execution.properties file during installation. This file can be edited after installation. This file is located:

For Windows:

```
%ESEC_HOME%\sentinel\config
```

For UNIX:

```
$ESEC_HOME/sentinel/config
```

For more information, go to Chapter 11 – Utilities, Configuring Sentinel email.

To send an event message by e-mail

1. In an Event Real Time table of the Visual Navigator or Snapshot, select an event or a group of events, right-click and select Email.

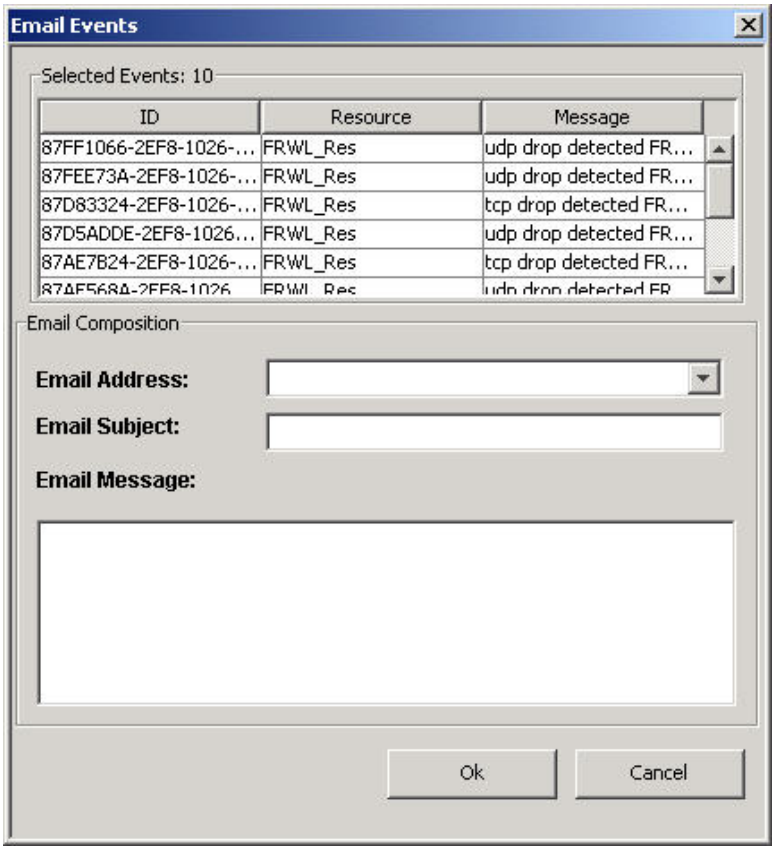


2. Complete the following:
   - Email Address
   - Email Subject
   - Email Message
3. Click OK.

To send an incident message by e-mail

1. After you save your incident, click the Incidents tab > View Incident List.
2. Double-click on an Incident View.
3. Double-click on an Incident.
4. Click the Email Incident button  .
5. Enter:
   - Email Address
   - Email Subject
   - Email Message
6. Click Ok. The e-mail message will have html attachments that address incident details, events, assets, vulnerabilities, advisor information and incident history.

## Creating an Incident

To perform this function you must have user permission Create Incident(s).

This is useful in grouping a set of events together as a whole representing something of interest (group of similar events or set of different events that indicate a pattern of interest such an attack).

**NOTE**: If events aren't initially displayed in a newly created Incident, it is most likely due to a lag in the time between display in the Real Time Events window and insertion into the database. If this occurs, it may take a few minutes for the original events to finally be inserted into the database and display in the incident.

To create an incident

1. In an Event Real Time table of the Visual Navigator or a Snapshot Event Real Time table, select an event or a group of events and right-click and select Create Incident.

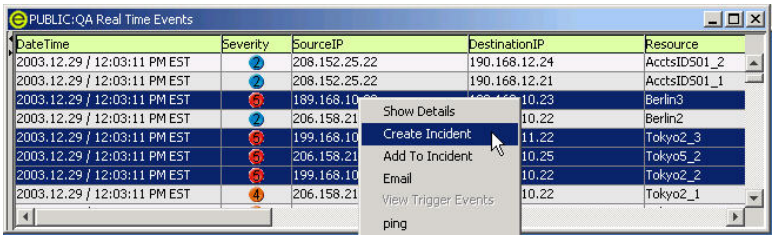


   In the New Incident Window, you have the following tabs:
   - Events – shows which events make up the incident
   - Assets – show affected assets
   - Vulnerability – show related asset vulnerabilities
   - Advisor – Asset attack and alert information
   - Workflow – under this tab, you may assign a WorkFlow (iTrac)
   - History – Incident history
   - Attachments – you may attach any document or text file with pertinent information to this incident

   In the Create Incident dialog box, enter:
   - Title
   - State
   - Severity
   - Priority
   - Category
   - Responsible - the user account assigned to the case
   - Description
   - Resolution
2. Click Save. The incident is added under the Incidents tab of the Sentinel Control Center.

## Viewing Events that Triggered a Correlated Event

You must right-click a correlated event in order to view the events that triggered the correlated event. In the event table from which you are selecting the event, look in the summary display panel on the right for an event that has a property of SensorType with a Value of C (C: correlated event) or W (W: watchlist).

To view events that triggered a correlated event

1. In an Event Real Time table of the Visual Navigator or Snapshot, or an Event Query table, right-click a correlated event and select View Trigger Events. A window opens showing the events that triggered the rule and the name of the Correlation Rule.

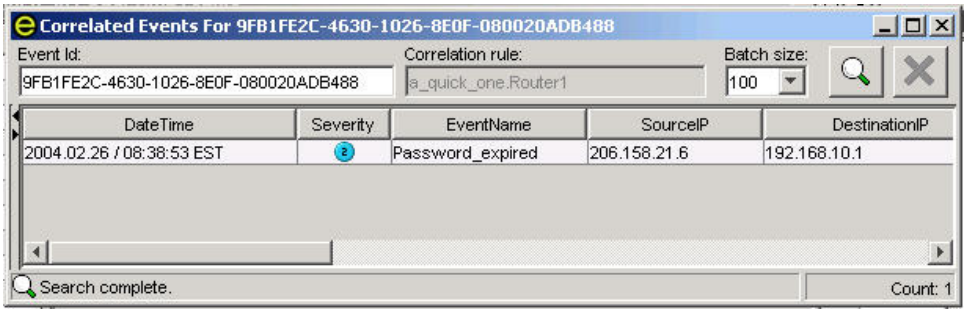


## Investigating an Event or Events

This function allows you:

- Graphically display the source fields (IP, port, event, sensor type, agent name, ...) mapped to destination fields (IP, port, event, sensor type, agent name, ...) of selected events.
- Perform a Event Query for the last hour on a single event for:

**NOTE**: You cannot perform a query on a null (empty) field.

  - Destination IP addresses
  - Source IP addresses
  - Event Name

Below is an illustration of source IP addresses to destination IP addresses.

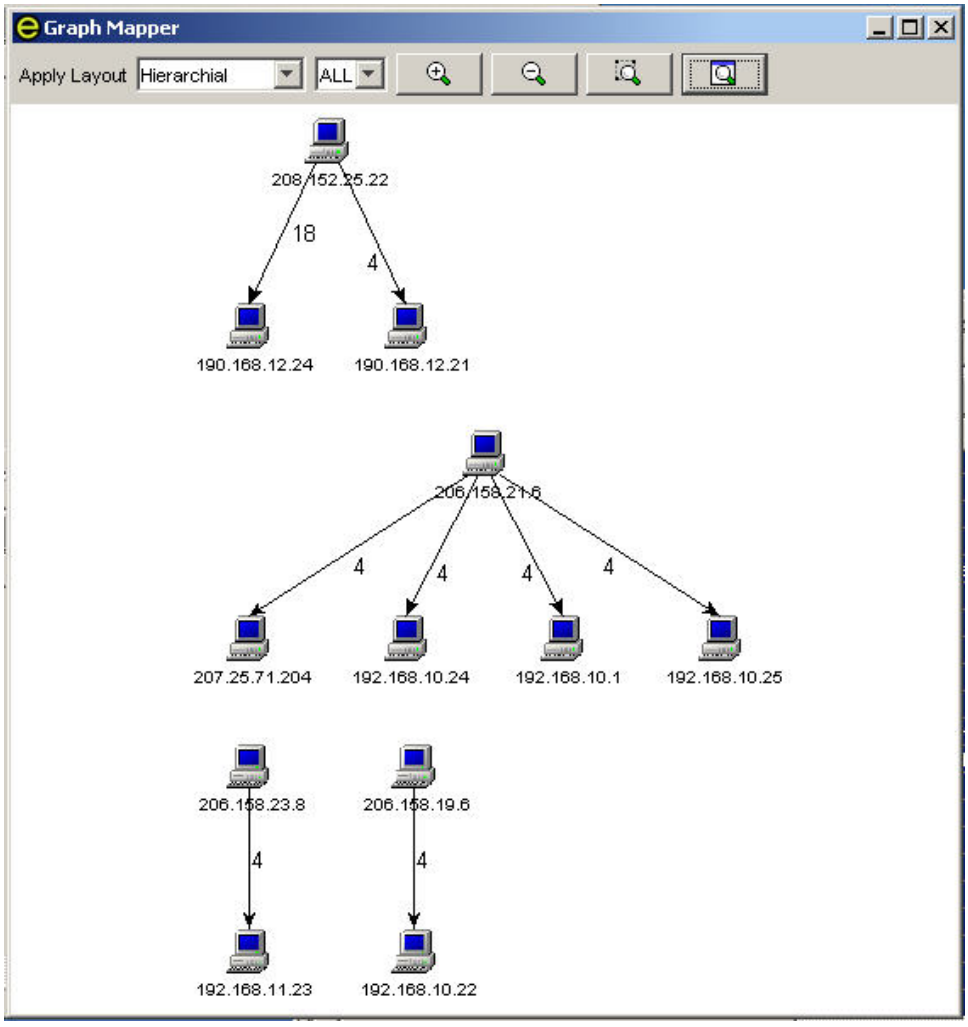


## Investigate – Graph Mapper

To create a graph map

1. In an Event Real Time of the Visual Navigator or Snapshot window, right-click an event or events > Investigate > Visual > Show Graph.

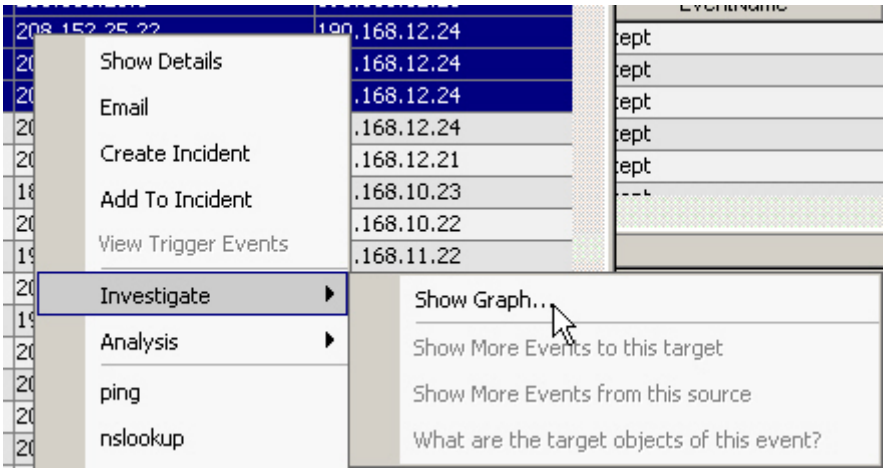

The following is a graphic depiction of Sensor Name to Event Name of severity 5 in an organic format. You can view a graphic mapping in the following formats:

- Circular
- Hierarchial
- Organic
- Orthoganal

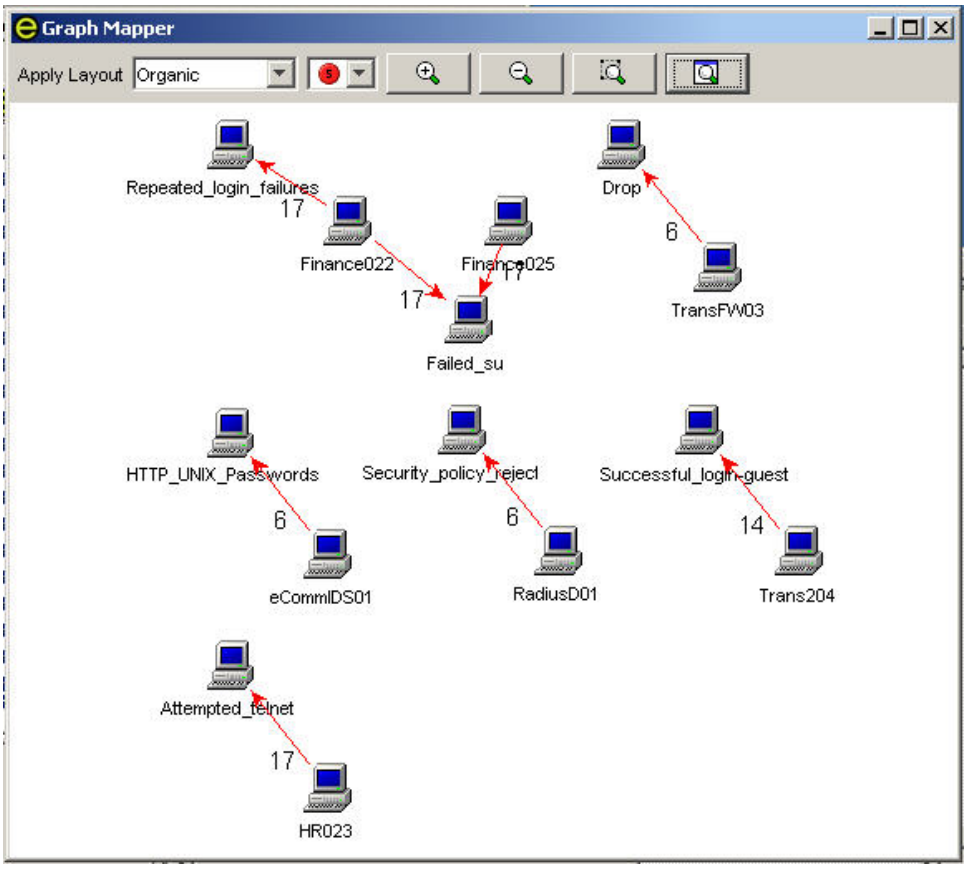


### Investigate – Event Query

This function allows you to Event Query within the last hour.

To perform a Event Query using the Investigate function

1. In a Visual Navigator or Snapshot window, right-click an event > Investigate > <select one of three options below>

| Option | Function |
|---|---|
| Show More Events to this target | Destination IP address |
| Show More Events to this source | Source IP address |
| What are the target objects of this event | Event Name |

## Analysis - Viewing Advisor Data

Advisor provides a cross-reference between real-time IDS attack signatures and Advisor's knowledge base of vulnerabilities. Advisor feed has an alert and attack feed. The alert feed contains information about vulnerabilities and viruses. The attack feed lists the exploits associated with vulnerabilities.

The supported Intrusion Detection Systems are:

- Cisco Secure IDS
- Enterasys Dragon Host Sensor
- Enterasys Dragon Network Sensor
- ISS BlackICE PC Protection
- ISS RealSecure Desktop
- ISS RealSecure Network
- ISS RealSecure Server Sensor
- ISS RealSecure Guard
- Snort/Sourcefire
- Symantec ManHunt
- Symantec Intruder Alert
- McAfee IntruShield

The IDS agent populates the DeviceAttackName (rt1) field of an event. Advisor uses this information to generate attack and vulnerability information. Some example vulnerabilities are:

- FINGER: Cfinger Search Probe
- SMTP: SmartServer3 MAIL FROM Buffer Overflow
- HTTP: Dragon Fire IDS Web Interface Remote Execution
- FTP:MKDIR-DOS
- hp-printer-flood
- wh00t-backdoor
- nt-telnet
- FINGER / execution attempt
- tellurian-tftpdnt-filename-bo
- FTP MKD Stack Overflow

## To View Advisor Data

1. In an Event Real Time table of the Visual Navigator or Snapshot, right-click an event or a series of selected events > Analysis > Advisor Data. If the DeviceAttackName field is properly populated, a report similar to the one below will appear. This example is for a WEB-MISC amazon 1-click cookie theft.

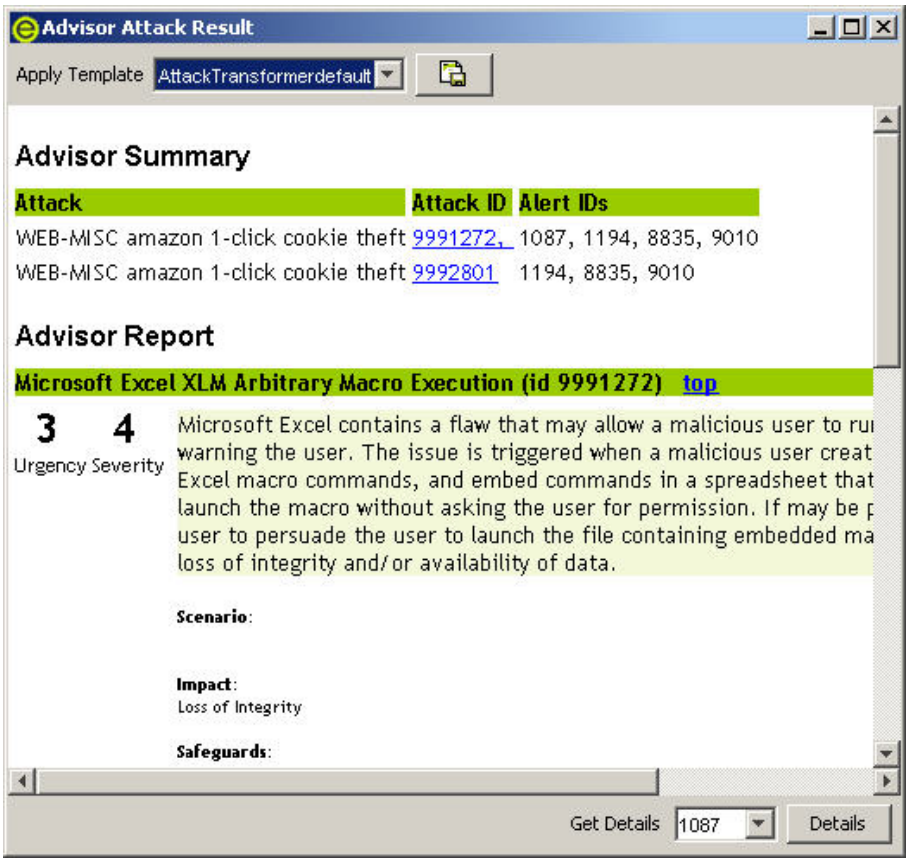

## Analysis - Viewing Asset Data

This function allows you to view and save your view as an HTML file of your Asset Report. You must run your asset management agent to view this data. The available data for viewing are:

**Hardware**
- MAC Address
- Name
- Type
- Vendor
- Product
- Version
- Value
- Criticality
- Sensitivity
- Environment
- Location

**Network**
- IP Address
- Hostname

**Software**
- Name
- Type
- Vendor
- Product
- Version

**Contacts**
- Order
- Name
- Role
- Email
- Phone Number

**Location**
- Room
- Rack
- Address

### To view Asset Data

1. In an Event Real Time table of the Visual Navigator or Snapshot window, right-click an event or events > Analysis > Asset Data. Window similar to the one below will appear.

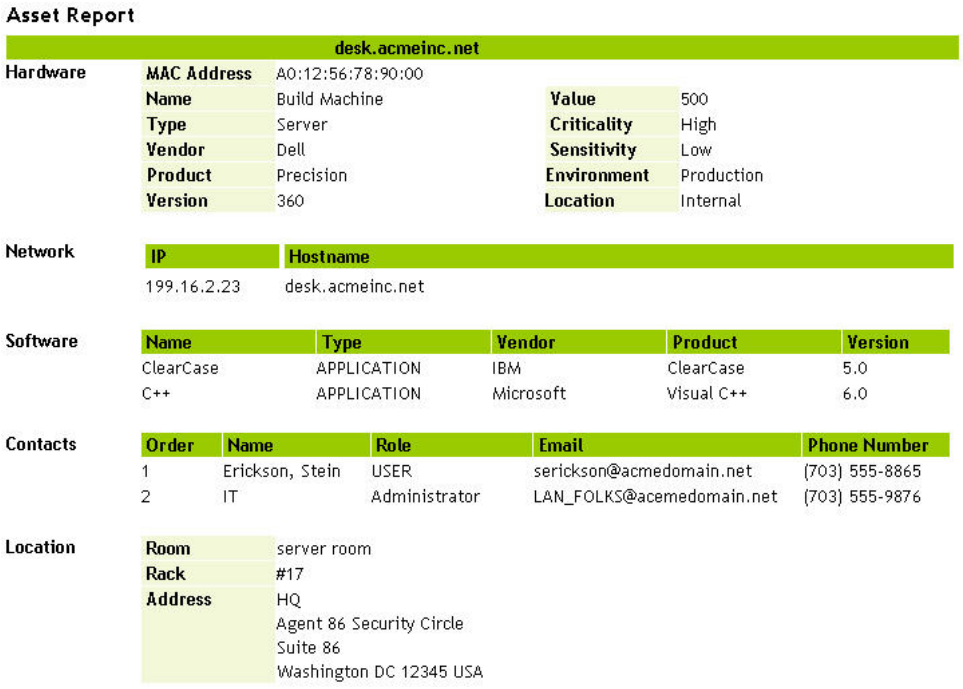

**Asset Report**

| desk.acmeinc.net | | | | |
|---|---|---|---|---|
| **Hardware** | **MAC Address** | A0:12:56:78:90:00 | | |
| | **Name** | Build Machine | **Value** | 500 |
| | **Type** | Server | **Criticality** | High |
| | **Vendor** | Dell | **Sensitivity** | Low |
| | **Product** | Precision | **Environment** | Production |
| | **Version** | 360 | **Location** | Internal |

| **Network** | **IP** | **Hostname** |
|---|---|---|
| | 199.16.2.23 | desk.acmeinc.net |

| **Software** | **Name** | **Type** | **Vendor** | **Product** | **Version** |
|---|---|---|---|---|---|
| | ClearCase | APPLICATION | IBM | ClearCase | 5.0 |
| | C++ | APPLICATION | Microsoft | Visual C++ | 6.0 |

| **Contacts** | **Order** | **Name** | **Role** | **Email** | **Phone Number** |
|---|---|---|---|---|---|
| | 1 | Erickson, Stein | USER | serickson@acmedomain.net | (703) 555-8865 |
| | 2 | IT | Administrator | LAN_FOLKS@acemedomain.net | (703) 555-9876 |

| **Location** | | |
|---|---|---|
| | **Room** | server room |
| | **Rack** | #17 |
| | **Address** | HQ<br>Agent 86 Security Circle<br>Suite 86<br>Washington DC 12345 USA |

# Analysis - Vulnerability Visualization

e-Security has agents available that process vulnerability scans from Nessus, ISS, Foundstone, eEye and Qualys scans. Vulnerability Visualization provides graphical representation of real-time event data against vulnerable systems and is available on an event for current and event time vulnerability.

This feature retrieves and displays vulnerability data for the destination IP's of selected events. For more information see the Agent pdf documentation located in %ESEC_HOME%\wizard\elements\<agent name>\doc. You must have one of the Scanner agents uploaded and running to get any vulnerability data.

**NOTE**: The vulnerability agent is an information agent, not an event agent.

You can view your vulnerability Visualization in either:

- HTML
- graphical
  - circular (organic)
  - hierarchical
  - All
  - Events Mapped Nodes
  - orthogonal

The HTML view is a report type view listing:

- IP
- host
- vulnerability
- port/protocol

Below is an example of a Nessus scan.

**Vulnerability Summary**

| IP | Host | Vulnerabilities | Port/Protocol |
|---|---|---|---|
| 172.16.0.132 | | 18 | 0\|)/TCP, 21(ftp)/TCP, 21(ftp)/TCP, 21(ftp)/TCP, 22(ssh)/TCP, 22(ssh)/TCP, 22\|ssh)/TCP, 23(telnet\|/TCP, 23\|telnet)/TCP, 25(smtp\|/TCP, 25(smtp\|/TCP, 25(smtp\|/TCP, 587(submission)/TCP, 587(submission)/TCP, 587\|submission)/TCP, 1241(nessus)/TCP, 1241\|nessus\|/TCP, 3306(mysql)/TCP |
| 172.16.0.71 | | 49 | 0\|)/TCP, 0()/TCP, 21(ftp)/TCP, 21(ftp)/TCP, 22(ssh)/TCP, 23(telnet)/TCP, 23\|telnet)/TCP, 23(telnet)/TCP, 25\|smtp)/TCP, 25(smtp\|/TCP, 25(smtp\|/TCP, 25(smtp\|/TCP, 25\|smtp)/TCP, 25\|smtp)/TCP, 25(smtp)/TCP, 25(smtp)/TCP, 25(smtp\|/TCP, 25(smtp\|/TCP, 90(dnsix)/TCP, 90(dnsix\|/TCP, 90(dnsix)/TCP, 90(dnsix)/TCP, 90\|dnsix)/TCP, 90(dnsix)/TCP, 90\|dnsix)/TCP, 111(sunrpc)/TCP, 111\|sunrpc)/TCP, 161\|snmp)/UDP, 512(exec\|/TCP, 513(login)/TCP, 514\|shell\|/TCP, 587(submission\|/TCP, 587(submission\|/TCP, 587(submission)/TCP, 587(submission)/TCP, 587\|submission)/TCP, 587\|submission)/TCP, 587(submission\|/TCP, 587(submission\|/TCP, 587(submission)/TCP, 587(submission)/TCP, 6000(x11)/TCP, 7100(font-service)/TCP, 32778(sometimes-rpc19)/TCP, 32779(sometimes-rpc21)/TCP, 32779(sometimes-rpc21)/TCP, 32779(sometimes-rpc21)/TCP |
| 172.16.0.132 | | 18 | 0\|)/TCP, 21(ftp)/TCP, 21(ftp)/TCP, 21(ftp)/TCP, 22(ssh)/TCP, 22(ssh)/TCP, 22\|ssh)/TCP, 23(telnet\|/TCP, 23\|telnet)/TCP, 25(smtp\|/TCP, 25(smtp\|/TCP, 25(smtp\|/TCP, 587(submission)/TCP, 587(submission)/TCP, 587\|submission)/TCP, 1241(nessus)/TCP, 1241\|nessus\|/TCP, 3306(mysql)/TCP |
| 172.16.0.71 | | 49 | 0\|)/TCP, 0()/TCP, 21(ftp)/TCP, 21(ftp)/TCP, 22(ssh)/TCP, 23(telnet)/TCP, 23\|telnet)/TCP, 23(telnet)/TCP, 25\|smtp)/TCP, 25(smtp\|/TCP, 25(smtp\|/TCP, 25(smtp\|/TCP, 25\|smtp)/TCP, 25\|smtp)/TCP, 25\|smtp)/TCP, 25(smtp)/TCP, 25(smtp)/TCP, 25(smtp\|/TCP, 25(smtp\|/TCP, 90(dnsix)/TCP, 90(dnsix\|/TCP, 90(dnsix)/TCP, 90(dnsix)/TCP, 90\|dnsix)/TCP, 90(dnsix)/TCP, 90\|dnsix)/TCP, 111(sunrpc)/TCP, 111\|sunrpc)/TCP, 161\|snmp)/UDP, 512(exec\|/TCP, 513(login)/TCP, 514\|shell\|/TCP, 587(submission\|/TCP, 587(submission\|/TCP, 587(submission)/TCP, 587(submission)/TCP, 587\|submission)/TCP, 587\|submission)/TCP, 587(submission\|/TCP, 587(submission\|/TCP, 587(submission)/TCP, 587(submission)/TCP, 6000(x11)/TCP, 7100(font-service)/TCP, 32778(sometimes-rpc19)/TCP, 32779(sometimes-rpc21)/TCP, 32779(sometimes-rpc21)/TCP, 32779(sometimes-rpc21)/TCP |

The graphical display is a rendering of vulnerabilities that link them to an event through common ports. Below is an example four available views.

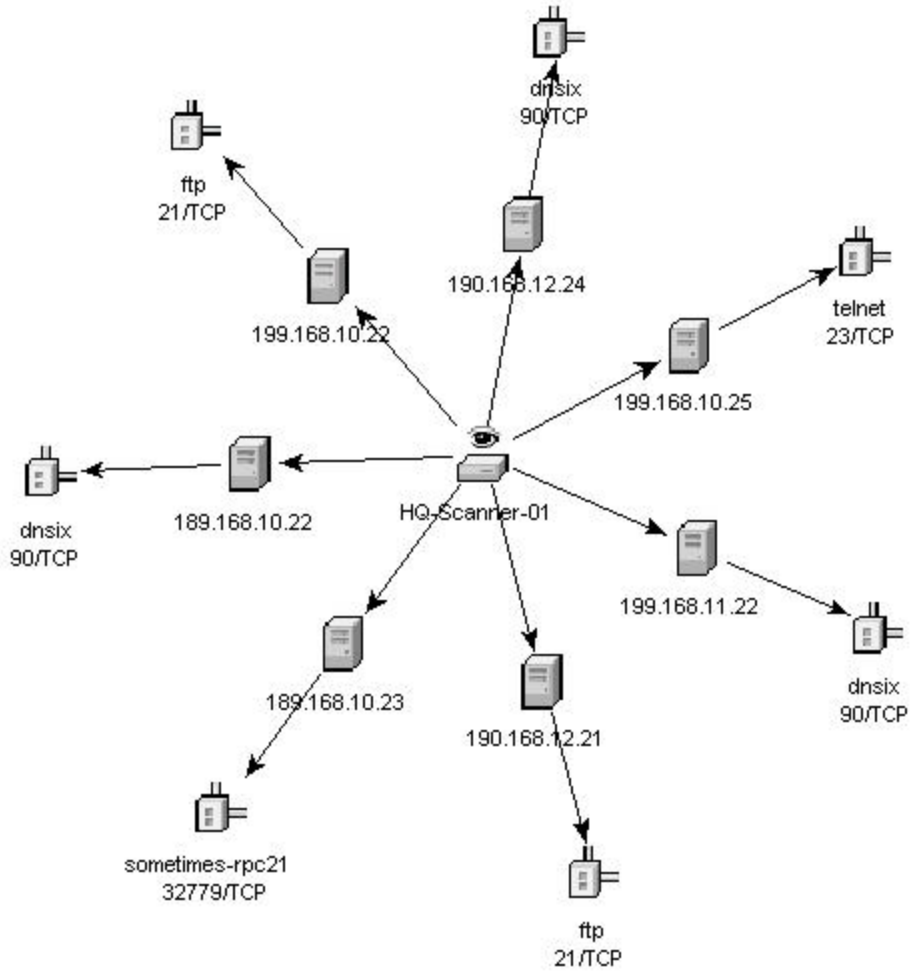


**Organic**

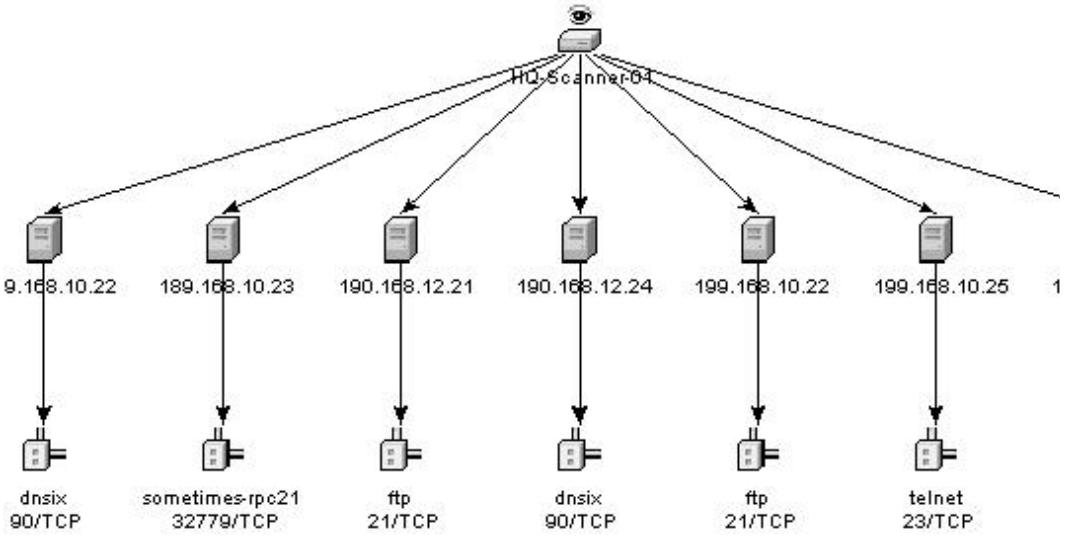


**Hierarchical**

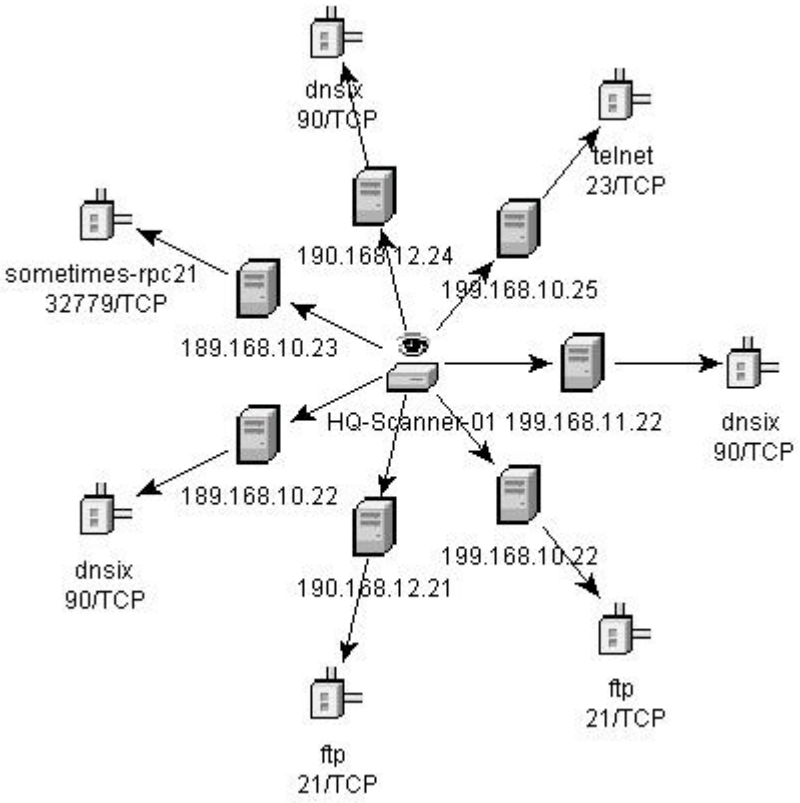


**Circular**

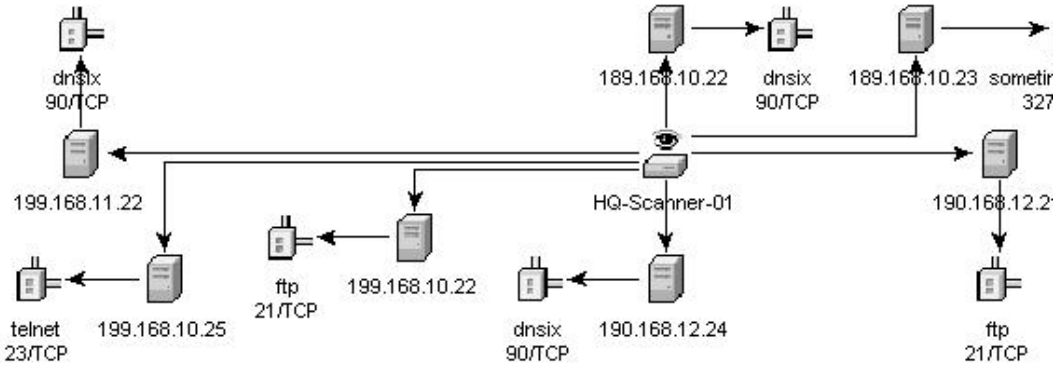


**Orthogonal**

In the graphical display there are four panels. They are:

- graph panel
- tree panel
- control panel
- details/events panel

The graph panel display associates vulnerabilities to a port/protocol combination of a resource (IP address). For example, if a resource has five unique

port/protocol combinations that are vulnerable, there will be five nodes attached to that resource. The resources are grouped together under the scanner that scanned the resources and reported the vulnerabilities. If two different scanners are used (ISS and Nessus), there will be two independent scanner nodes that will have vulnerabilities associated with them.

**NOTE**: Event mapping takes place only between the selected events and the vulnerability data returned.

The tree panel organizes data in same hierarchy as the graph. The tree panel also allows users to hide/show nodes at any level in the hierarchy.

The control panel exposes all the functionality available in the display. This includes:

- four different algorithms to display
- ability to show all or selected nodes which have events mapped to them
- zooming in and out of selected areas of the graph

In the Details/Events panel, you have two tabs. When in the Details tab, clicking on a node will result in displaying node details. When in the Events tab, clicking on an event associated with a node the node will display in tabular form as in a Real Time or Event Query window.

### To run a Vulnerability Visualization

1. In an Event Real Time table of the Visual Navigator or Snapshot, right-click an event or a series of selected events and click:
   - Analysis
     - Current Vulnerability – queries the database for vulnerabilities that are active (effective) at the current date and time.
     - Event Time Vulnerability – queries the database for vulnerabilities that were active (effective) at the date and time of the selected event.

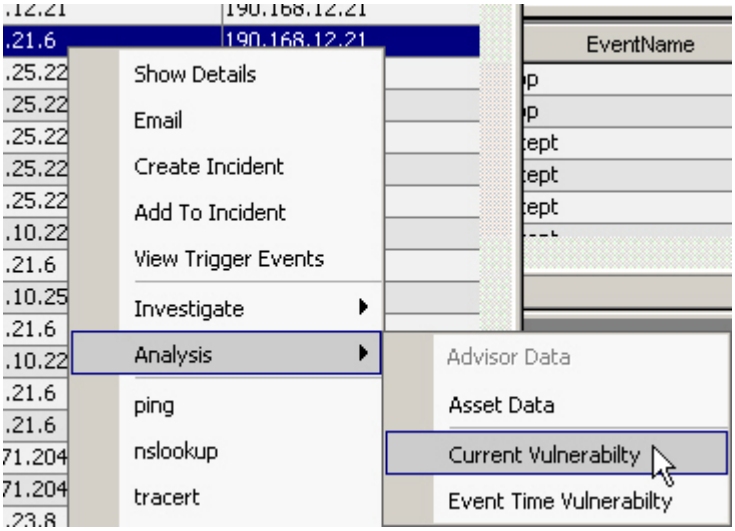


2. At the bottom the vulnerability results window, click on either:
   - Event to Vulnerability Graph
   - Vulnerability Report
3. (For Event to Vulnerability Graph) Within the display, you can:

- move nodes and their labels
- use one of four different layout algorithms to display the graph
- show all nodes or only those nodes that have events mapped to them
- in-line tree filtering in the event that a large number of resources are returned as vulnerable
- zoom in and out of selected areas

## 3rd Party Integration

3rd Party Integration allows you to send events from any display screen including incidents and associated objects to either:

- HP Service Desk
- Remedy

**To send single or multiple events for 3rd Party Software**

1. In an Event Real Time table of the Visual Navigator or Snapshot window, depending on which 3rd Party Integration software you have installed, right-click an event and click Send Event to either:
   - HP Service Desk
   - Remedy

## Using Custom Menu Options with Events

**To use a custom menu option with an event**

1. In an existing Event Real Time table of the Visual Navigator or Snapshot, select an event or a group of events and right-click to and click an option. A dialog box opens with the information the menu option is configured for or enabling you to complete the information needed to perform an action. The default custom menu options are as follows:
   - ping
   - nslookup
   - traceroute
   - Whois?

   You can further assign user permission to View Vulnerability and to perform HP Actions. You can add options using the Menu Configuration window that's available in the Admin tab.

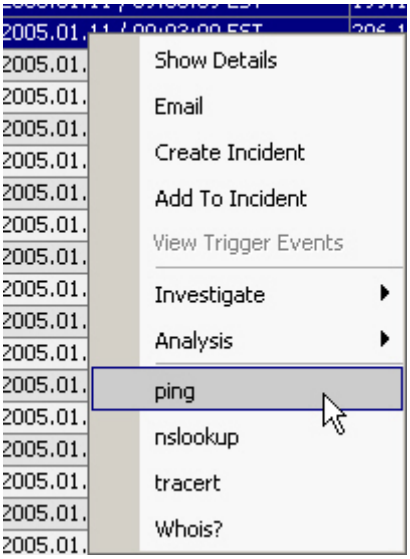


## Managing the Columns in a Snapshot or Visual Navigator Window

To select and arrange columns in a Snapshot or Visual Navigator

1. With a Snapshot or Visual Navigator window open, click Active View > Event Real Time > Manage Columns or click the Manage Columns of Event Real Time Table.

2. Use the Add and Remove buttons to move column titles between the Available Columns list and the Show columns in this order list. The Insert can be used to put an available column item into a specific location. For example, in the illustration below clicking the Insert button will place AttackId above DateTime.

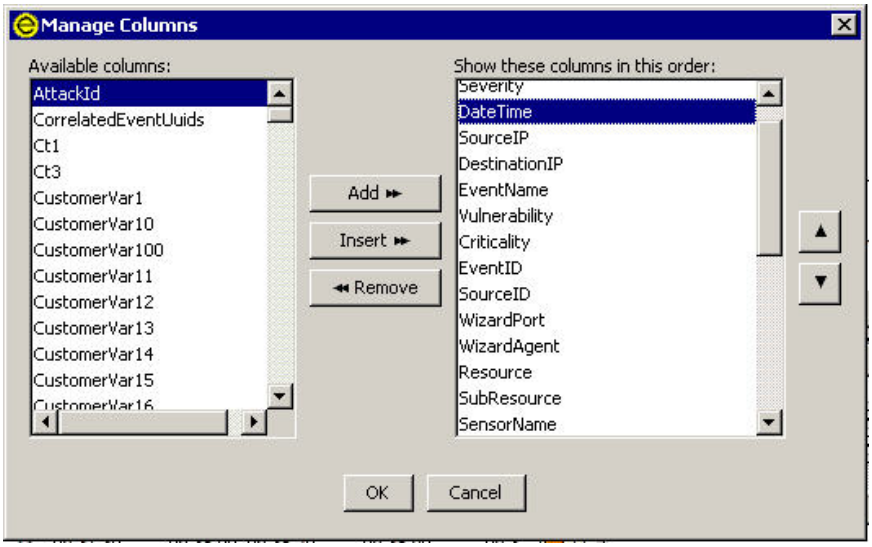


Use the Up and Down arrow buttons to arrange the order of the columns as you want them to display in the Event Real Time table. The top to bottom

order of column titles in the Manage Column dialog box determines the left to right order of the columns in the Event Real Time table.

3. In the Manage Column dialog box, click OK.
4. If your columns to display the next time you open the Sentinel Control Center, click File > Save Preferences or click the Save User Preference button.

## Taking a Snapshot of a Visual Navigator Window

To perform this function you must have user permission Snapshot.

This is useful to study events of interest since the Visual Navigator refreshes automatically and the alert or alerts of interest may scroll off the screen. Also, within a snapshot, you can sort by column.

### To take a snapshot of an Event Real Time table

1. With a Visual Navigator window open, click Active View > Event Real Time > Snapshot or click the Snapshot Event Real Time Table button on the menu bar.

A Snapshot window opens and is added to the Snap Shots folder list under Event Views in the Navigator. The graphical display will not be part of the snapshot.

## Sorting Columns in a Snapshot

### To sort columns in a Snapshot

1. Click any column header once to sort by ascending value and twice to sort by descending value.

## Closing a Snapshot or Visual Navigator

### To close a Snapshot or an Event Real Time table

1. With a Snapshot or a Visual Navigator window open and if you want the table to be available the next time you start the Sentinel Control Center, click file > Save Preferences.
2. Close the table using the Close button (upper right corner in Windows or upper left corner in UNIX).

## Deleting a Snapshot or Visual Navigator

### To delete a Snapshot or Visual Navigator Window

1. With a Snapshot or Visual Navigator open, close by using the Close button (upper right corner in Windows or upper left corner in UNIX).
2. Click File > Save Preferences or click the Save User Preference button.

The view or snapshot will not redisplay when you close and reopen the Sentinel Control Center.

## Adding Events to an Incident

To perform this function you must have user permissions to Modify Incident(s) and Assign Incident(s).

To add events to an incident

1. In an Event Real Time table or a Snapshot, select an event or a group of events and right-click to display and click 'Add To Incident'.
2. In the 'Add To Incident' dialog box, click the Browse button.

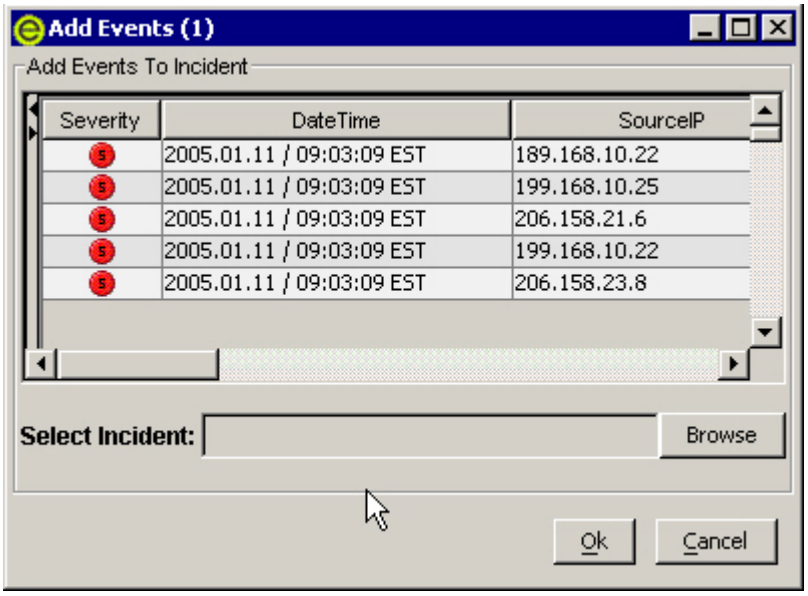


3. Click Search button to list the available incidents.

**NOTE**: You can define your criteria to better search for a particular incident or incidents.

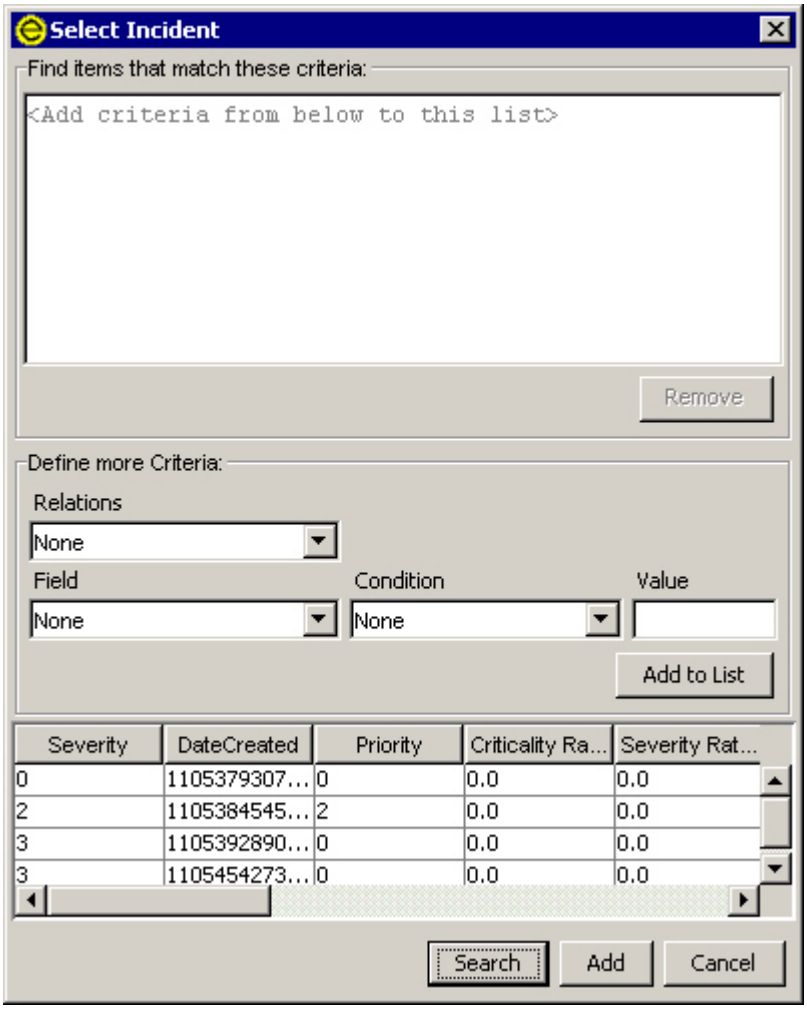


4. Highlight an incident and click the Add button.
5. Click Ok. The event or events selected are added to the incident in the Incidents Navigator.

**NOTE**: If events aren't initially displayed in a newly created Incident, it is most likely due to a lag in the time between display in the Real Time Events window and insertion into the database. If this occurs, it may take a few minutes for the original events to finally be inserted into the database and display in the incident.

# Chapter 4 – Incidents Tab

You must have the proper permission to use View Incidents tab. If this permission is not assigned, none of the other permissions related to actions using this tab will be available.

This chapter discusses incidents. Incidents are groupings of one or more events that are of interest.

Incidents can be created:

- Real Time window, events can be individually selected to create a new incident or added to an existing incident
- Incidents can also be created automatically through correlation rules that are triggered

## Incidents Tab - Description

With incidents, you can:

- Email an Incident
- Modify an Incident
- View an Incident
- Delete an Incident
- Add an Incident View

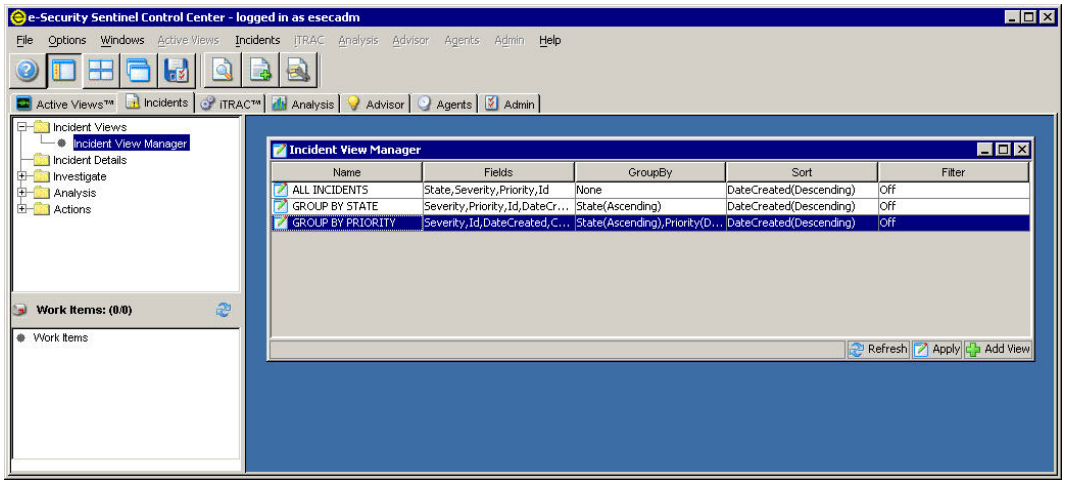


### Relationship between Events and Incidents

An event is an action or occurrence detected by a security device or program. Events are considered to be “stateless”.

An incident is the grouping of one or more events that are deemed to be important (a possible attack). Incidents have “states” in that they require a response and closure.

### Viewing an Incident

You must have the user permission View Incident(s).

To View an Incident

1. Click the Incidents tab.

2. Click Incidents > View Incident List or click the View Incident List button.
 .
3. In the Incident View Manager window you have choice of the following views:
   - All Incidents
   - Group By State
   - Group By Priority

   Double-click on a view name.
4. Right-click > Expand to view the incidents.

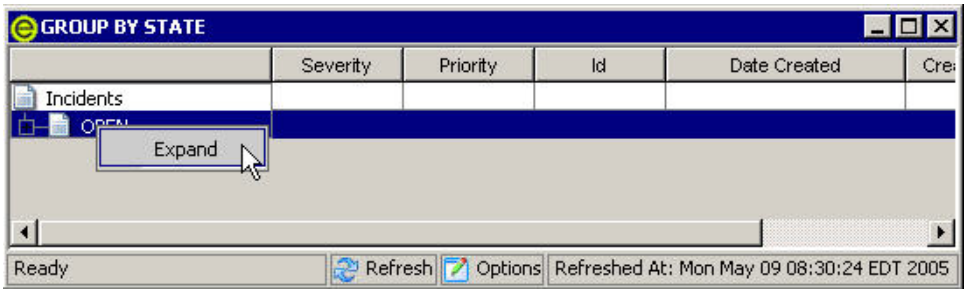


## To set an Incident view option

1. Click the Incidents tab.
2. Click Incidents > View Incident List or click the View Incident List button.
 .
3. In the Incident View Manager window, double-click on a view name.

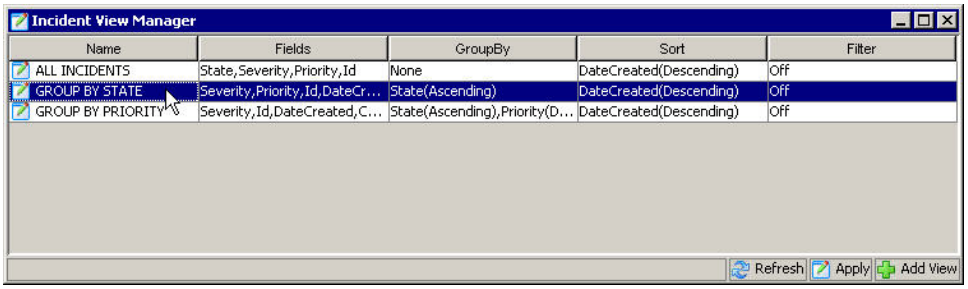


4. Click the Options button.


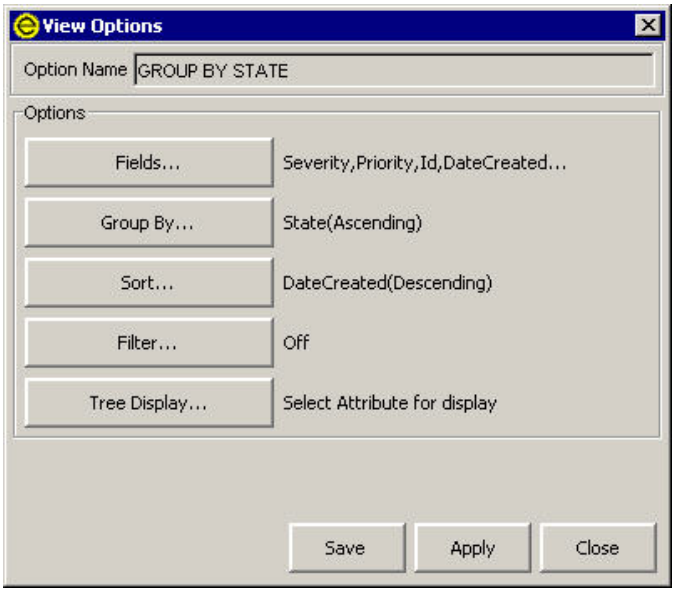


In this window you may also set the following:

- Fields...
- Group by...
- Sort...
- Filter...
- Tree Display

Click Apply and Save.

5. In the View Options Manager window, double-click on a view name.

The following is a default view of the All Incident View window.

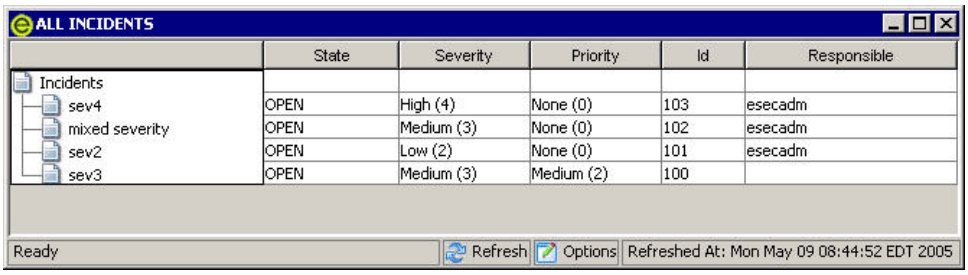


The following is a view sorted by severity, with Fields (column management) for the first four columns set to Severity, Dated Created, Priority and Criticality Rating.

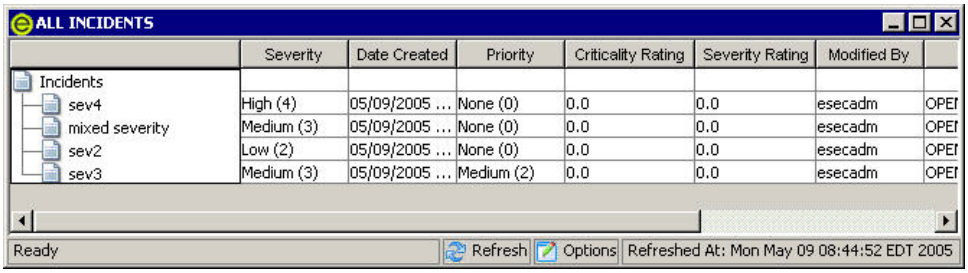


The following is a view grouped by title.

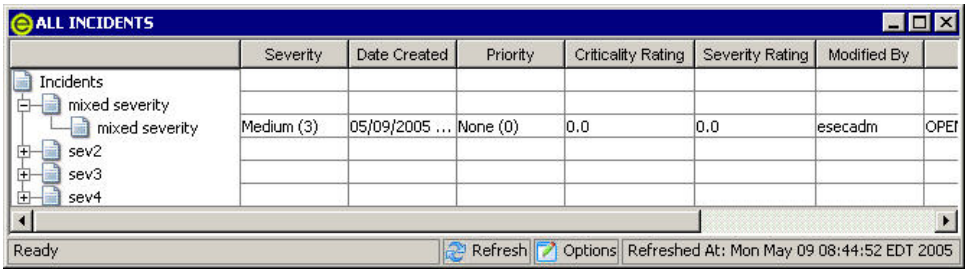


The following is a view tree by date created (DateCreated).

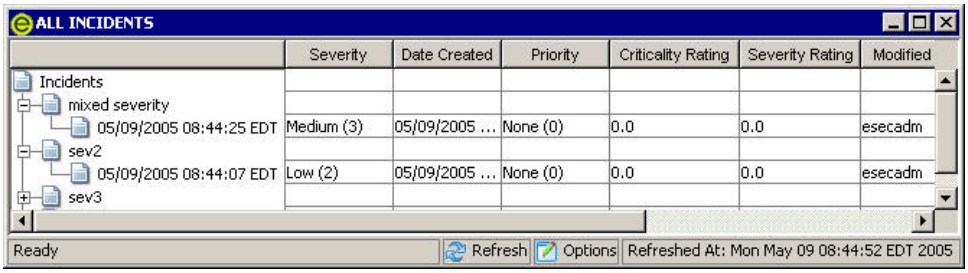


## Adding an Incident View

When adding an Incident View, you have the option to:

- Fields...
- Group by...
- Sort...
- Filter...
- Tree Display

To add an Incident View

1. In the Incident View Manager, click the Add View button.

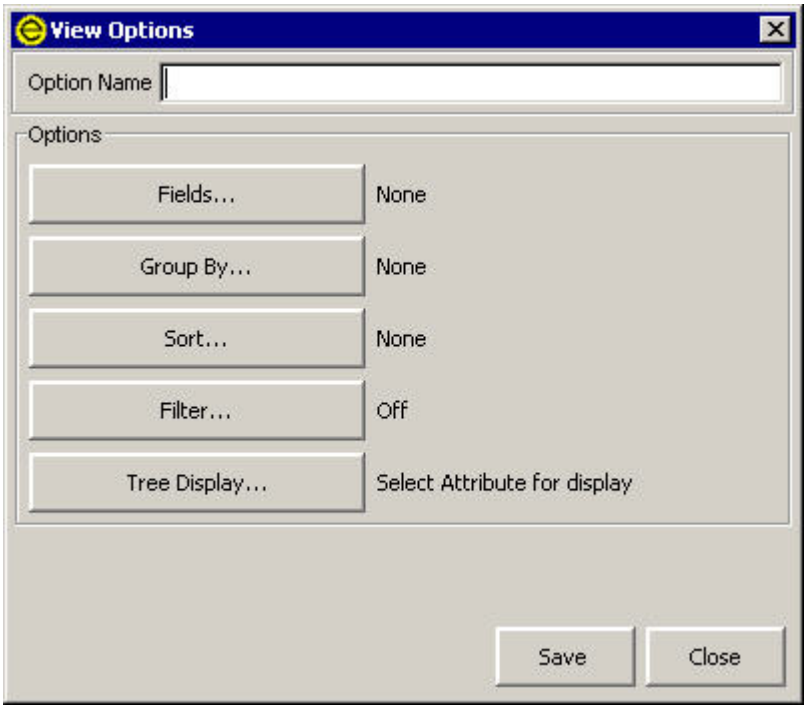

2. Enter an Option Name and select which options you want, click Save.

## Incident Fields and Details

### Incident Fields

- Title – Name of the incident
- State
  - Open
  - Acknowledged
  - Assigned
  - Investigating
  - False Positive
  - Verified
  - Approved
  - Closed
- Severity
  - None (0)
  - Trivial (1)
  - Low (2)
  - Medium (3)
  - High (4)
  - Severe (5)
- Priority
  - Low (1)
  - Medium (2)
  - High (3)
  - Urgent (4)
  - Top (5)
- Category – (optional), text entry that can be used to further identify the incident.
- Responsible - the user account assigned to the case
- Description – text entry
- Resolution – text entry

### Incident Details

- Events – events associated with the incident
- Assets – list of all of the assets associated with the incident
- Vulnerability – displays any vulnerability associated with the incident
- Advisor – displays any attack information associated with the incident
- Workflow – displays workflow associated with the incident. Under this tab, you may assign:
  - None
  - HIPAA Compliance Process
  - SANS Incident Response Process
  - Sarbanes Oxley FTP Compliance Process
  - Automatic Response
- History – incident history (lists all of the actions that were performed on the incident, this includes date/time user action and brief information)
- Attachments – you may attach any pertinent information (text files or documents) to this incident
- External Data

**NOTE**: When events are added to an incident the assets/vulnerability and Advisor tab will be populated with a list of all Asset/Vulnerability/Advisor data corresponding to the DIP/Destination Host names of the associated events.

**NOTE**: The Add and Remove buttons on the Assets/Vulnerability/Advisor tab allows users to manually add or remove assets, vulnerability or advisor data.

## Creating an Incident

Creating an Incident

1. Click the Incident tab.
2. Click Incidents > Create Incident or click the Create a New Incident button.

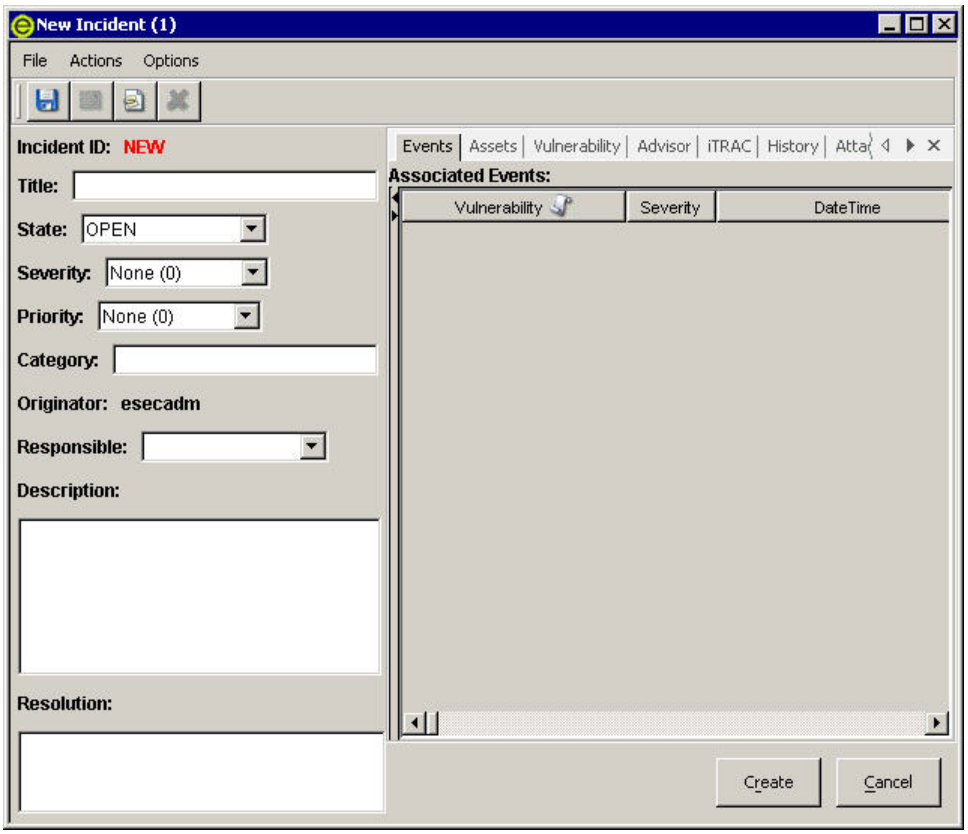


In the Create Incident dialog box, enter your information in the blanks fields.

3. Click Save.

## Viewing and Saving Attachments

To view an attachment

1. Right-clicking on an attachment > View or Save the attachment.

**NOTE**: An attachment viewer has to be configured to view an attachment. If an attachment is not configured to open a file, a prompt will appear as to what program is to open the file. Attachment files are saved to the Sentinel Database.

### Configuring the Attachment Viewer

Configuring the Attachment Viewer

1. Click the Incident tab.

2. Click Incidents > Attachment Viewer Configuration or click the Configure Attachment Viewers button.

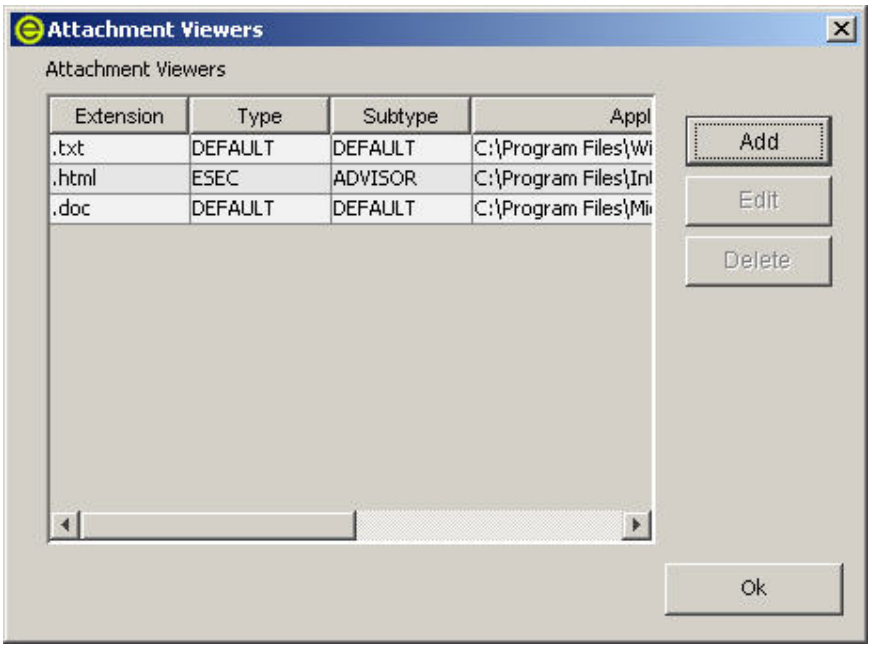


3. Click the Add button.

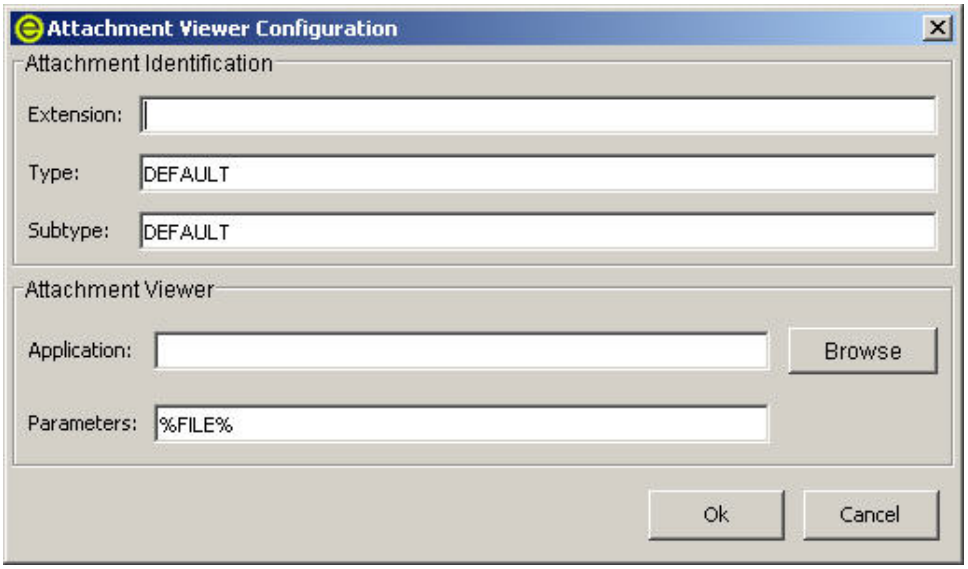


Enter the extension type (such as .doc, .xls, .txt, .html, etc...) and click browse or type in the application program to launch the file type (such as notepad.exe for Notepad).

4. Click Ok.

## Emailing an Incident

Ability to send emails is set in the execution.properties file during installation. To configure this file, see Chapter 11 - Utilities.

emailing an Incident

1. Click the Incidents tab.
2. If available, in the navigator, expand the Incidents folder or Click Incidents > View Incidents List or click the View Incidents List button.

 .

3. Double click on an Incident View name.
4. Double-click on an incident.
5. Click Email Incident button  .
6. Enter:
   - Email Address
   - Email Subject
   - Email Message
7. Click Ok. The e-mail message will have html attachments that address incident details, events, assets, vulnerabilities, advisor information and incident history.

## Modifying an Incident

To modify an incident

1. Click the Incidents tab.
2. Click Incidents > View Incident List or click the View Incident List button.

 .

3. Double-click on an incident view.
4. Double-click on an incident.
5. The Incident details window opens.
6. Optionally, you can edit the following fields in an Incident:
   - Title
   - State
   - Severity
   - Priority
   - Category
   - Responsible
   - Description
   - Resolution
7. Under the Attachments tab, you may add or remove attachments.
8. Click Save or Cancel.

## Deleting an Incident

**NOTE**: To delete an incident that is attached to a WorkFlow (iTRAC), you will need to terminate the iTRAC Process.

To delete an incident

1. Click the Incidents tab.

2. Click Incidents > View Incident List or click the View Incident List button.  .
3. Double-click on an incident view.
4. In the Incidents View Window, right-click on an incident > Delete.

**NOTE**: To delete an incident that is attached to a WorkFlow (iTRAC), you will need to terminate the iTRAC Process. An iTRAC Process can be terminated using the Process View Manager under the iTRAC tab. For more information, see Chapter 5 – iTRAC Tab.

5. In the confirmation window click Yes. Click No to cancel.

# Chapter 5 – iTRAC™ Tab

iTRAC (workflow) involves the automation of procedures, the ability to respond to incidents. e-Security provides a iTRAC management system that provides procedural automation of processes. Tied to iTRAC is e-Security's activity framework. The Activity framework provides the activities that could be performed automatically at each stage in the iTRAC process.

Templates (Process Definition) and Process Execution together constitute the workflow management system.

## Templates (Process Definition)

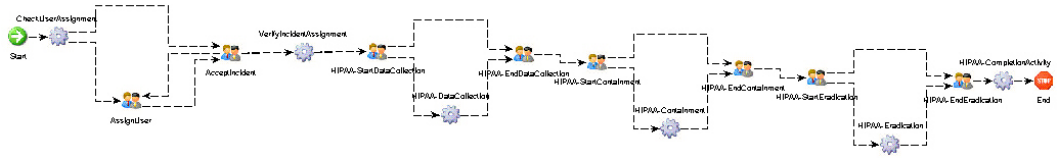


The template is the design that controls the flow of execution in iTRAC. The template consists of a network of activities and their relationships, criteria for transition between activities and information about individual activities. Templates have attributes that can be modified by the user.

iTRAC allows users to set timeout attributes on a iTRAC template.

An activity is a logical, self-contained unit of work within the iTRAC process. An activity represents work, which will be processed either by users/roles (manual activity) or computer applications (automatic activities).

Activities also have timeouts and users may enable/disable timeouts on all manual or automatic activities.

Manual Activities in addition to the timeout attributes allow users to configure the resource attribute that determines the user/role performing that activity.

Automatic Activities in addition to the timeout attributes allow users to configure the automatic activity from the e-security activity framework to be executed.

### Template Manager

iTRAC allows users to create new templates, manipulate process and activity attributes in an existing template and delete templates using the template manager window in the iTRAC Tab.

The template manager may be accessed by clicking on the Template Manager node on the navigator tree in the iTRAC tab.

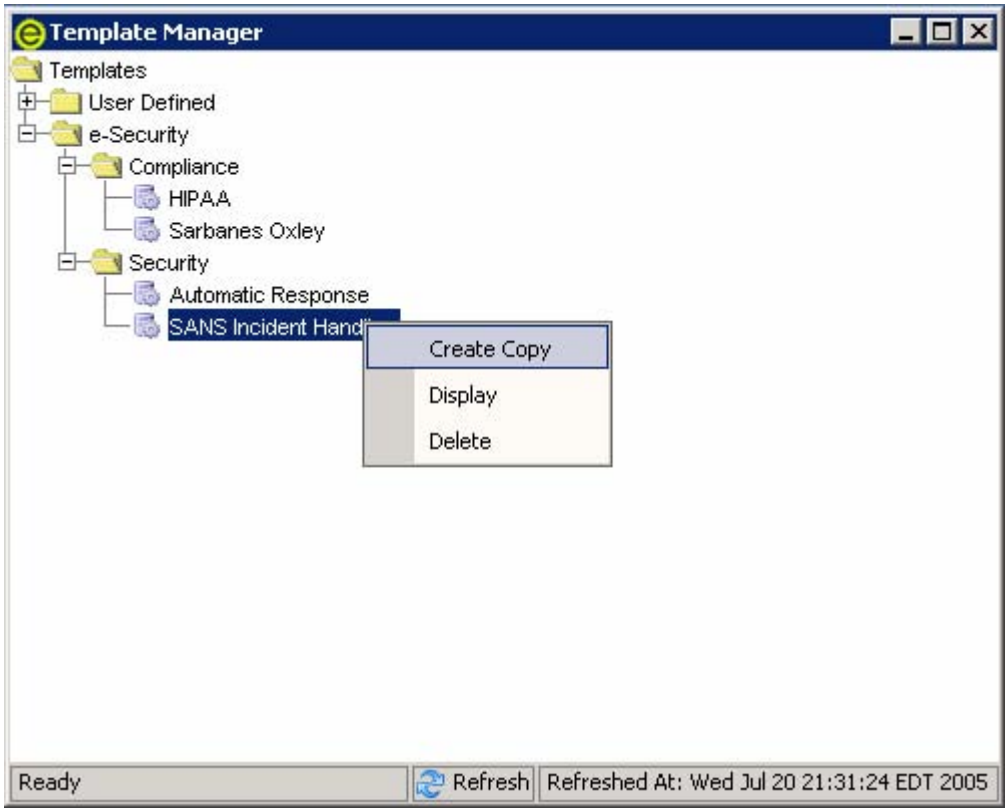


## Default Templates

iTRAC ships with four default templates which consists of automatic and manual activities. The process and activity attributes for these templates have been set to some pre-defined values, users may modify these to suit their requirements. The default templates are:

- HIPAA
- Sarbanes Oxley
- SANS Incident Handling
- Automatic Response

### Creating New Templates

1. Click the iTRAC tab.
2. In the Navigator, click iTRAC Administration > Template Manager.
3. Highlight an existing process (HIPAA, Sarbanes-Oxley, SANS or a user defined process), right-click > Create Copy.
4. Enter a Name.
5. If you select a time out, you must enter an email address and a time. Time is in whole numbers. You can select minutes, seconds, hours or days.
6. Enter a description. See Modifying Existing Templates for changing process and activity attributes.
7. In the Template Customizer, click Save

## Modifying Existing Templates

When modifying a process, you can modify process attributes or attributes of the activities within the process:

The following process attributes may be modified

- name
- time out period or disable the timeout period
- description

## Modifying Process Attributes

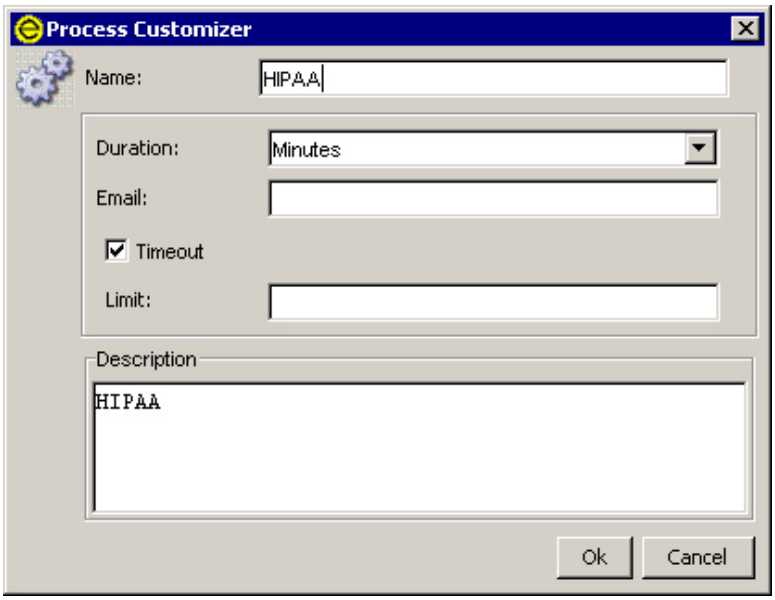


1. Click the iTRAC tab.
2. In the Navigator, click iTRAC Administration > Template Manager.
3. Highlight an existing template, right-click > Display.
   On the template window click the process details button.
   
4. In the Process Customizer dialog, you can edit the following:
   - Name
   - Duration (minutes, seconds, hours or days)
   - Timeout (if enabled you will be required to enter an email address and time)
   - Description

## Modifying Manual Activities

You can edit the resource (user/role), Timeout and Description of manual activities

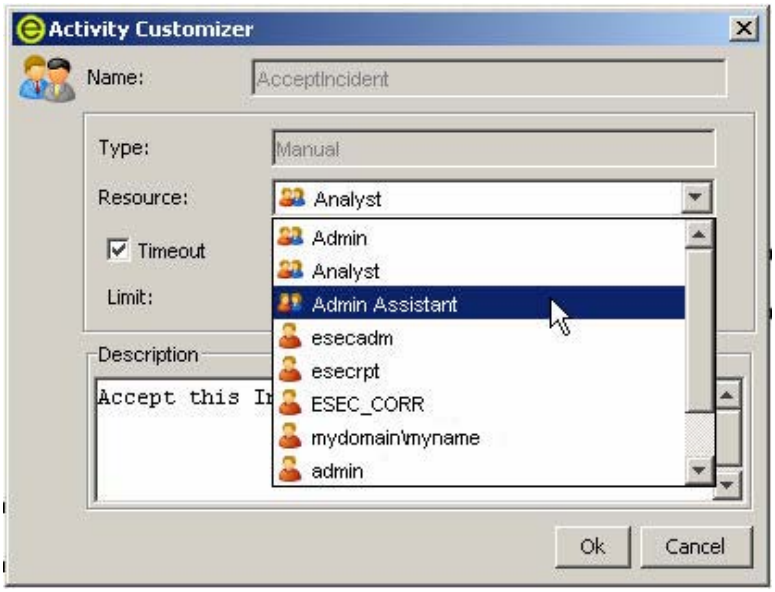


1. Click the iTRAC tab.
2. In the Navigator, click iTRAC Administration > Template Manager.
3. Highlight an existing template, right-click > Display.
4. The Template is displayed on a separate window.
5. To edit, double-click on any of the manual icons on the template and make your changes.

Note: the following manual activities in the existing templates may be modified this way.

- AssignUser
- AcceptIncident
- ConfirmStartDataCollection
- ConfirmEndDataCollection
- ConfirmStartContainment
- ConfirmEndContainment
- ConfirmStartEradication
- ConfirmEndEradication

## Modifying Automatic Activities

You can edit the activity, Timeout and Description of an automatic activity.

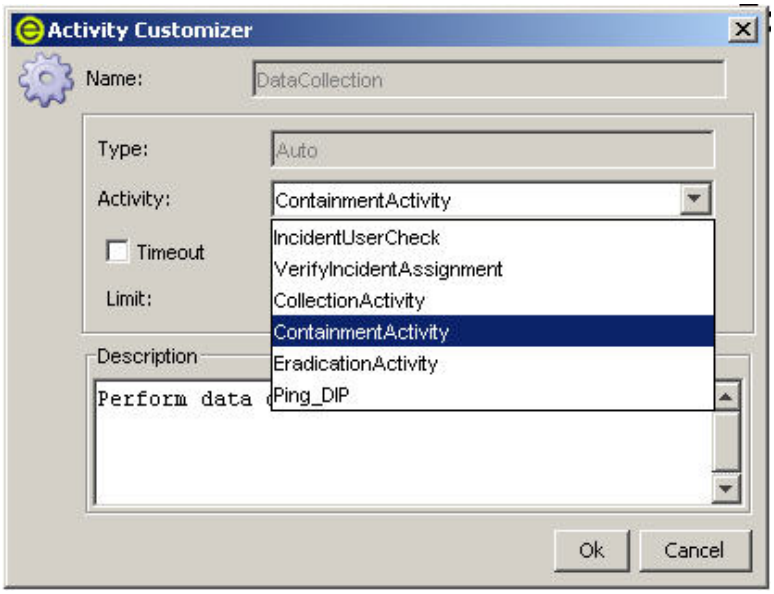


1. To edit, double-click on any of the automatic activity icons on the template and make your changes.
2. The drop down on the activity customizer dialogs displays the list of activities that may be used as automatic activities.The activities in the list are activities created using the activity framework.

**NOTE**: The following automatic activities in the existing templates may be modified this way.

- DataCollection
- Containment
- Eradication

### Deleting Templates

1. Click the iTRAC tab.
2. In the Navigator, click iTRAC Administration > Template Manager.
3. Highlight an existing template, right-click > Delete.
4. Click yes on the delete template popup.

## Process Execution

Process Execution is the time period during which the process is operational, with process instances being created and managed.

When an iTRAC process is executed or instantiated in the iTRAC server, a process instance is created, managed and eventually terminated by the iTRAC server in accordance with the process definition. As the process progresses towards completion or termination it executes various activities defined in the workflow template based on the criteria for the transitions between them. The iTRAC workflow server processes manual and automatic activities differently.

An iTRAC process is dependent on an e-security incident, a process instance cannot exists if there is exists no incident that is related to it. On the other hand an incident may exist without being related to the workflow server. Only 1 incident may be associated to a iTRAC process instance.

### Instantiating a Process

An iTRAC process may be instantiated in the iTRAC server by associating an incident to a iTRAC process by the following 3 methods

- Associate an iTRAC process to the incident at the time of incident creation
- Associate an iTRAC process to incident after an incident has been created
- Associate an iTRAC process to an incident via correlation

Please refer to the chapter on incidents tab for more details on associating a process to an incident.

### Automatic Activity Execution

When the process instance executes an automatic activity, it executes the associated activity defined in the template. The associated activity is an activity created using the activity framework. The iTRAC server executes the activity; stores the result in process variables and transitions to the next activity in the iTRAC template.

For e.g. the activity from the activity framework may be set to ping a server and attach the results to the associated incident.

### Manual Activity Execution

On encountering a manual activity, the iTRAC server sends out notifications in the form of a workitems to the assigned resource. If the assigned resource is a user, then the workitem will be sent only to that user. If the activity was assigned to a role then a worktem will be sent to all users within the role. The iTRAC server then waits for the user to complete the workitem before proceeding to the next activity.

### Work Lists

The workitems are presented to the user via the work list which maintains details of all the workitems allocated to that user. It is to-do list for the user.


The work list is viewable from any tab in the sentinel UI. Workitems are grouped together by process and activity to which they belong. Workitems in bold indicate those workitems that have not yet been accepted by the user.

The worklist allows users to interact with the individual workitems.

- User's may double click or right click > Details to view workitem details.
- Users may right click and accept unaccepted workitems
- Users may right click and view the associated incident details

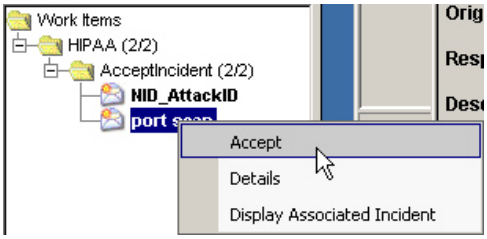


## Workitems

A work item constitutes the task to be undertaken by the user for the currently executing manual activity in an iTRAC process. The control and progression of the workitem rests with the user.

The iTRAC server waits for the user to complete the task before it proceeds to the next activity within the process instance

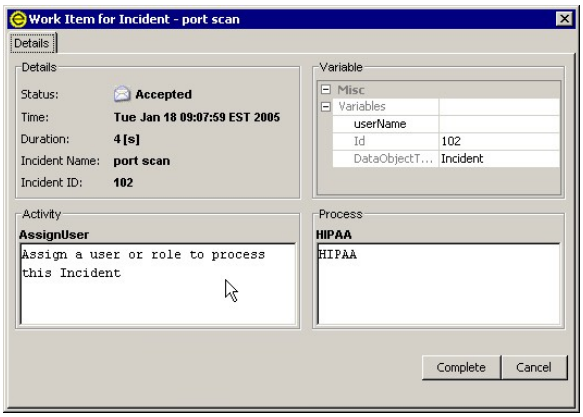


The workitem details dialog shown above has the following information

- Workitem details
- Workitem variables
- Activity description
- Process description

There are three steps in involved in interacting with a workitem, they are:

- Accepting a workitem
- Updating variables in the workitem
- Completing the workitem

## Accepting a workitem

A workitem may be assigned to all users within a role or to just a single user. A workitem has to be accepted by the user before performing any other action on the workitem. Accepting the workitem makes the user the owner of the workitem, the workitem is removed from the work list of all other assigned users.

### Accepting Workitems

1. Within the Worklist, you can right click on a work item and perform the following:

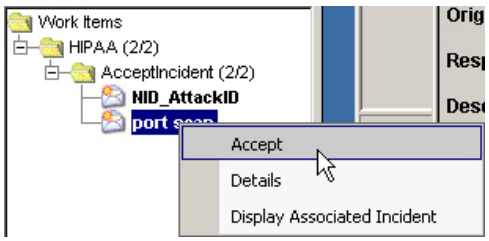


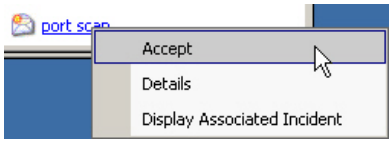


   - Accept (when the process is in an Accept step)
   - Alternatively you can bring up the details window and click on the Accept button.

## Updating variables in the workitem

The iTRAC server uses workitems to get information from users in the form of workitem variables to determine the next activity with a process. The variables may accessed by the user only after accepting the workitem.

iTRAC supports read-only variable and updateable variables, read-only variables are used to inform the user for e.g. status of an activity, id of an incident etc.

Updateable variables are used to accept input from the users. Currently in iTRAC there are two kinds of updateable variables, User list and Boolean list.

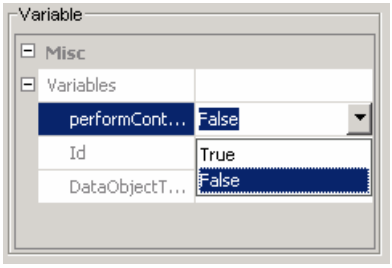


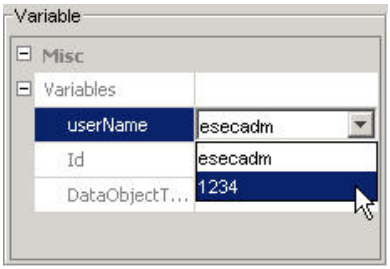


Updating Variables

1. Right click or double click on the workitem to view the details dialog
2. Only updateable variables are in edit mode, read other variables cannot be edited.
3. Click on the combo box and select the appropriate value.

### Completing the workitem

Completing the workitem signals the completion of the task to the iTRAC server. The updateable variables from the workitem are processed by the server to move to the next activity based on some criteria. The workitem is removed from the user's worklist. A workitem has to be accepted before it can be completed.

Completing Workitems

1. Right click or double click on the workitem to view the details dialog
2. Click on the complete button on the dialog

## Process Management

Process Management allows you to:

- Display Status of your Process (Process Monitor)
- Start your Process
- Terminate your Process

### Process Monitor

The Process Monitor function is to monitor the progress of a process. As the process instance progresses from one activity the user may track the progress visually by clicking on the refresh button, the process monitor also provides an audit trail of all the actions performed by the iTRAC server while executing the process.

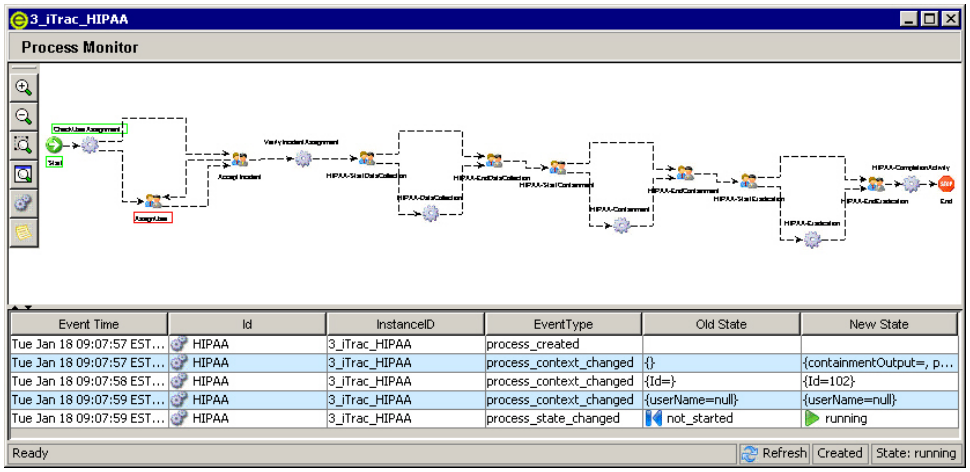


Activities that have completed by the process are displayed with a green border while the activity that is currently executing is displayed with a red border.

## Accessing Process Monitoring

1. Click the iTRAC tab.
2. Click the View Options Manager button.

3. Double-click on one of the default views or create a new view. Default views are:
   - All Processes
   - Processes By Incident
   - Processes By Status
4. In the Active Process Manager, highlight and double-click on a process.

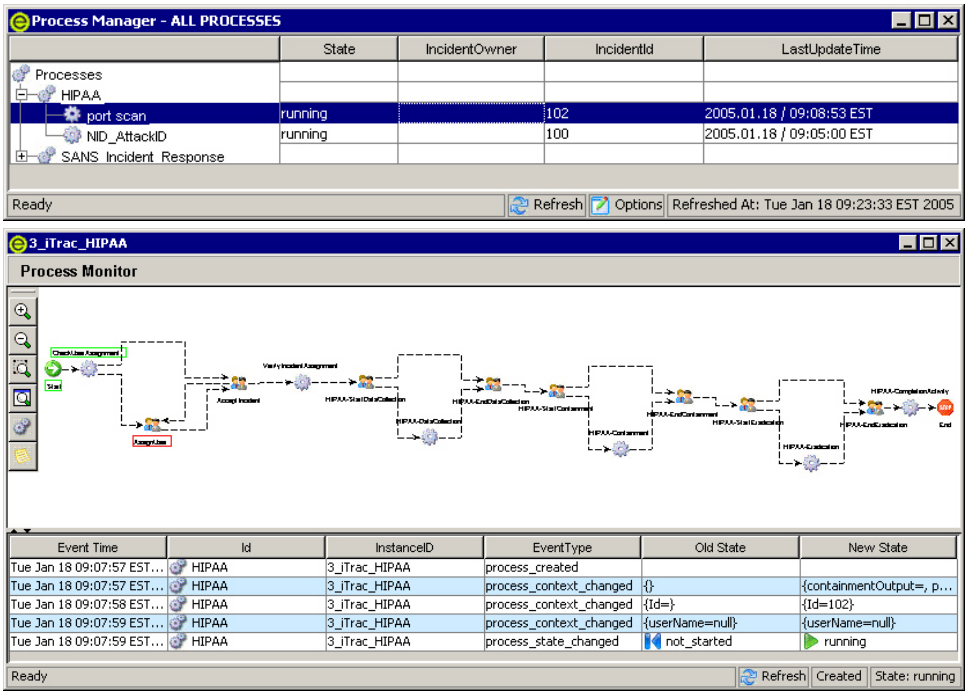

To set an Process Manager option

1. Click the iTRAC tab.
2. Double-click on any of the processes.
3. Click the Options button. In this window you may also set your:
   - Fields...
   - Group by...
   - Sort...
   - Filter...
   - Tree Display
4. Click Apply and Save.

   The following is view with Tree Display set to Status (running and not started).

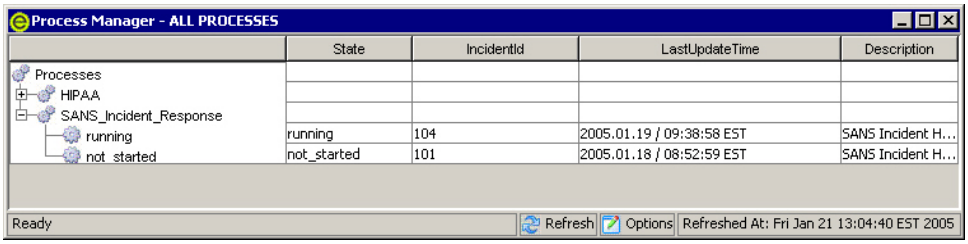


### Starting or Terminating a Process

Starting or Terminating a Process

1. Click the iTRAC tab.
2. Click the View Options Manager button.

3. Double-click on one of the default views or create a new view. Default views are:
   - All Processes
   - Processes By Incident
   - Processes By Status
4. In the Active Process Manager, highlight a process, right-click and select either 'Start Process' or 'Terminate Process'.

## Creating an Activity Using the Activity Framework

Creating an Activity

1. Click the iTRAC tab.
2. In the Navigator, click iTRAC Administration > Activity Manager.
3. Right-click > New Activity.
4. Select one of the following:


- Incident Command Activity – launch a specific command with or without arguments.

  Incident Output gives the following arguments:

  - DIP
  - DIP:Port
  - incident
  - RT1 (DeviceAttackName)
  - SIP
  - SIP:Port
  - Text

  Custom allows you to enter your own custom arguments.

  For this activity you can also set it to email the output and/or attach the output to the incident.

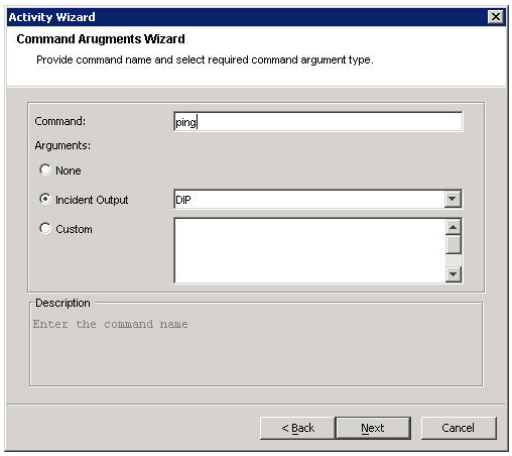


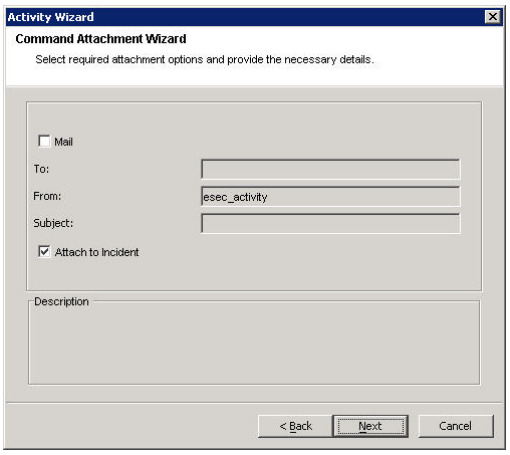


- Incident Internal Activity – allows you mail and/or attach information about:

  - Vulnerability for (SIP or DIP)
  - Advisor Data
  - Asset

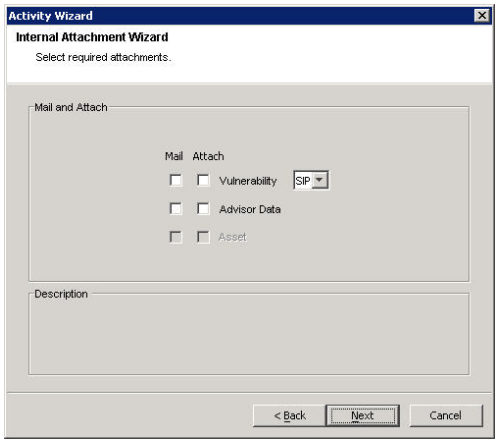


- Incident Composite Activity – allows you to create an activity by combining one or more existing activities.

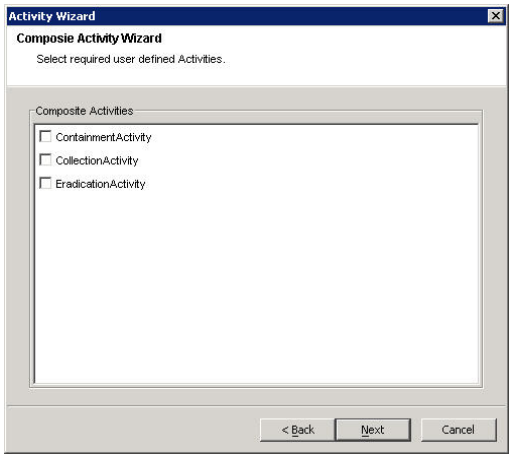


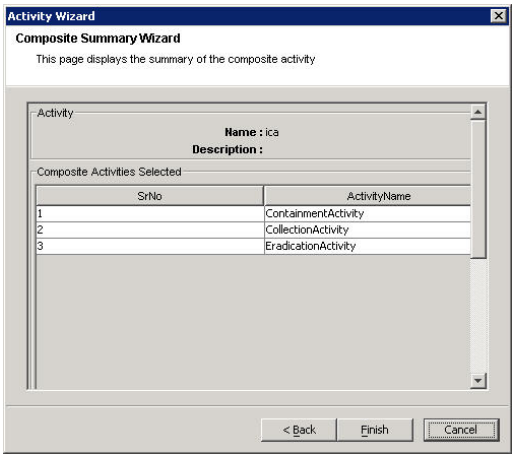


## Modifying an Activity

### Modifying an Activity

1. Click the iTRAC tab.
2. In the Navigator, click iTRAC Administration > Activity Manager > iTRAC Activities.
3. Double-click on an iTRAC activity. Edit and click Ok.

## Importing/Exporting an Activity

Activities are exported as xml files. These files can be imported from one system to another.

### Exporting an Activity

1. Click the iTRAC tab.
2. In the Navigator, click iTRAC Administration > Activity Manager.
3. Right-click on iTRAC Activities > Import/Export Activity.
4. Select Export Activity and click the Explore button.
5. Navigate to where you want save your exported file.
6. Name your file and click Export.

7. Click Next.
8. Select one or more activities to be exported.
9. Click Next and click Finish.

### Importing an Activity

1. Click the iTRAC tab.
2. In the Navigator, click iTRAC Administration > Activity Manager.
3. Right-click on iTRAC Activities > Import/Export Activity.
4. Select Import Activity and click the Explore button.
5. Navigate to your import file. Click Import.
6. Click Next.
7. Click Next and click Finish.

# Chapter 6 – Analysis Tab

You must have the proper permission to use Analysis tab. If this permission is not assigned, none of the other permissions related to actions using this tab will be available.

## Description

The Analysis tab allows for historical reporting. Historical and vulnerability reports are published on a web server, these run directly against the database and they appear on the Analysis and Advisor tabs on the Navigator bar.

**NOTE**: Sentinel is integrated with Crystal Reports® to generate and display reports. The administrator must configure the location of the Crystal Enterprise Server that publishes reports in the General Options window of the Admin tab. In the navigator window is a list of available reports.

In order to run the report templates, you need to have Crystal Reports Enterprise Edition installed and have your Sentinel Control Center configured to access that server. For more information, see Sentinel™ 5 Installation Guide.

### Top Ten Reports

To run any Top 10 reports, aggregation must be enabled and EventFileRedirectService in DAS_Binary.xml must be set to on. For information on how to enable aggregation, see Sentinel User's Guide, Chapter 10 – Sentinel Data Manager, section Reporting Data Tab.

**Enabling EventFileRedirectService for Sentinel Top 10 Reports**

Enabling EventFileRedirectService

1. At your DAS machine, using text editor, open:
   For UNIX:
   ```
   $ESEC_HOME/sentinel/config/das_binary.xml
   ```
   For Windows:
   ```
   %ESEC_HOME%\sentinel\config\das_binary.xml
   ```
2. For EventFileRedirectService, change the status to on.
   ```
   <property name="status">on</property>
   ```
3. For Windows, restart the eSecurity service. For UNIX, reboot the DAS machine.

### Running a Report from Crystal Reports

To create a report from a Crystal Reports template

1. Click the Analysis tab.
2. In the Analysis Navigator, click a report from the available reports.

**NOTE**: To run any Top 10 reports, aggregation must be enabled and EventFileRedirectService in DAS_Binary.xml must be set to on. For information on how to enable aggregation, see Sentinel User's Guide, Chapter 10 – Sentinel Data Manager, section Reporting Data Tab.

3. Click Analysis > Create Report or click the Create Report button.

4. Complete the information in the template and click View Report. The report displays.

## Running a Event Query Report

To create a Event Query report

1. Click the Analysis tab.
2. In the Analysis Navigator, open the Historical Reports folder.
3. Click Event Query.
4. Click Analysis > Create Report or click the Create Report button.

An Event Query window will open.

5. Set the following:
   - time frame
   - filter
   - severity level
   - batch size (this is the number of events to view – events display from oldest events to newer events)
6. Click the Refresh Query button.
7. To view the next batch of events, click the More button.
8. Rearrange the columns by dragging and dropping them and arrange the sort order by clicking in the column heading.
9. When your query is complete, it is added to the list of quick queries in the Navigator.

## Running a Correlated Events Report

To create a correlated events report

1. Click the Analysis tab.
2. In the Analysis Navigator, open the Historical Reports folder.
3. Click Correlated Events.
4. Click Analysis > Create Report or click the Create Report button.

A Correlated Events Report window will open.

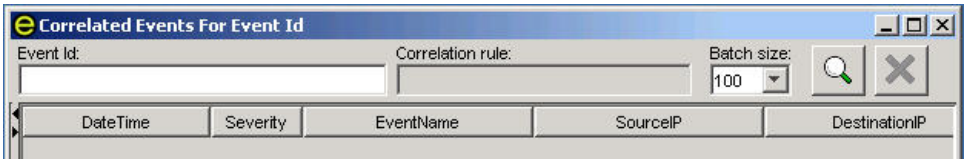


5. In the Correlation ID field, enter either:
   - Event ID number
   - CorrelatedEventUUID

**NOTE**: CorrelatedEventUUID is only available from a real time event table.

6. To view the next batch of events, click the More button.

# Chapter 7 – Advisor Tab

You must have the proper permission to use the Advisor tab. If this permission is not assigned, none of the other permissions related to actions using this tab will be available.

Advisor is an optional module. If you do not have an Advisor license, when you click on the Advisor tab you will get a notification screen indicating as such.

e-Security Advisor is powered by SecurityNexus. Advisor provides real-time intelligence into enterprise vulnerabilities, expert advice and recommended steps toward remediation. Advisor provides a cross-reference between real-time IDS attack signatures and Advisor's knowledge base of vulnerabilities. Visit http://www.esecurity.net/Software/Products/Advisor.asp to find out more information.

The Advisor data feed contains two parts:

- Alert Data: information relating to known security vulnerabilities and threats
- Attack Data: normalization of intrusion detection signatures and vulnerability scanning plug-ins

**NOTE**: During install and up until the initial datafeed from SecurityNexus, the right-click function on an event (with the rt1 field populated) for Advisor data will not be fully functional.

## Running Advisor Reports

To create an Advisor report

1. Click the Advisor tab.
2. In the Advisor Navigator, click a report template.
3. Click Advisor > Create Report.
4. Complete the information in the template and click View Report.

## Standalone Installation – Advisor Manual Updating

Manual Advisor Feed Updating

1. Go to url //advisor.esecurityinc.com/advisordata/.
2. Enter your username and password.
3. Go to the latest month under the attack and alert folders and download the zip files.
4. Place the new alert and attack feed data files (the files will be in a zip format) on your computer.

**NOTE**: Do not place the zip files in the attack and alert directories.

5. Unzip the attack feed zip files to:

   For Windows:

   ```
   <location specified during install for Advisor data
      files>\attack
   ```

or

For UNIX:

```
<location specified during install for Advisor data
   files>/attack
```

6. Unzip the alert feed zip files to:
   For Windows:

```
<location specified during install for Advisor data
   files>\alert
```

   or

   For UNIX:

```
<location specified during install for Advisor data
   files>/alert
```

7. Go to:
   For Windows:

```
%ESEC_HOME%\sentinel\bin
```

   For UNIX:

```
$ESEC_HOME/sentinel/bin
```

8. Run the following command:
   For Windows:

```
advisor.bat
```

   For UNIX:

```
./advisor.sh
```

**NOTE:** advisor.sh and advisor.bat will update the database and then delete the attack and alert files that were unzipped into the attack and alert directories.

## Direct Internet Download – Advisor Manual Updating

### Manual Advisor Feed Updating

1. Go to:
   For Windows:

```
%ESEC_HOME%\sentinel\bin
```

   For UNIX:

```
$ESEC_HOME/sentinel/bin
```

2. Run the following command:
   For Windows:

```
advisor.bat
```

   For UNIX:

```
./advisor.sh
```

**NOTE:** advisor.sh and advisor.bat will update the database and then delete the attack and alert files that were unzipped into the attack and alert directories.

## Changing Your Advisor Server Password and email Configuration

### Changing Your Advisor Server Password (standalone)

This procedure is not applicable for standalone configurations.

### Changing Your Advisor Server Password (Direct Download)

To change your Advisor server password (Direct Download)

1. Submit a password change to e-Security Support Desk.
2. After being informed of the password change from e-Security, for UNIX login as esecadm or for Windows login with administrative rights.
3. cd to:

   For UNIX:

   ```
   $ESEC_HOME/sentinel/bin
   ```

   For Windows:

   ```
   %ESEC_HOME%\sentinel\bin
   ```

4. Enter the following commands:

   For UNIX:

   ```
   ./adv_change_passwd.sh <oldpassword> <newpassword>
   ```

   For Windows:

   ```
   adv_change_passwd.bat <oldpassword> <newpassword>
   ```

### Changing Your Advisor Server email Configuration

To change your Advisor server email configuration

1. For UNIX login as esecadm or for Windows login with administrative rights.
2. cd to:

   For UNIX:

   ```
   $ESEC_HOME/sentinel/config
   ```

   For Windows:

   ```
   %ESEC_HOME%\sentinel\config
   ```

3. Using a text editor open alertcontainer.xml and alertcontainer.xml. Make your editorial changes in the grayed area.

   ```
   <property
      name="advisor.mail.from">fromNAME@domain.com</pro
      perty>
   ```

```
<property
   name="advisor.mailto.list">toNAME@domain.com</pro
   perty>
```

**NOTE**: To more than one mail to address, enter email addresses as comma separated without spaces.

## Changing Your Datafeed Time

By default, the datafeed times are:

- Six Hour: 01:00, 07:00, 13:00 and 19:00
- Twelve Hour: 02:00 and 14:00

To change the datafeed times

1. Login to your Advisor machine (for UNIX login as esecadm).
2. To edit your datafeed times:

   For UNIX: use the 'crontab' command

   For Windows: use the 'at' command

# Chapter 8 – Agents Tab

You must have the proper permission to use Agents Tab. The Agents Tab allows limited Wizard functionality. For full Wizard functionality, use the e-Security Agent Builder. The Agents Tab allows you to:

- monitor a Wizard Host
- monitor an Agent
- start and stop Agents (Agent Manager) for a selected host

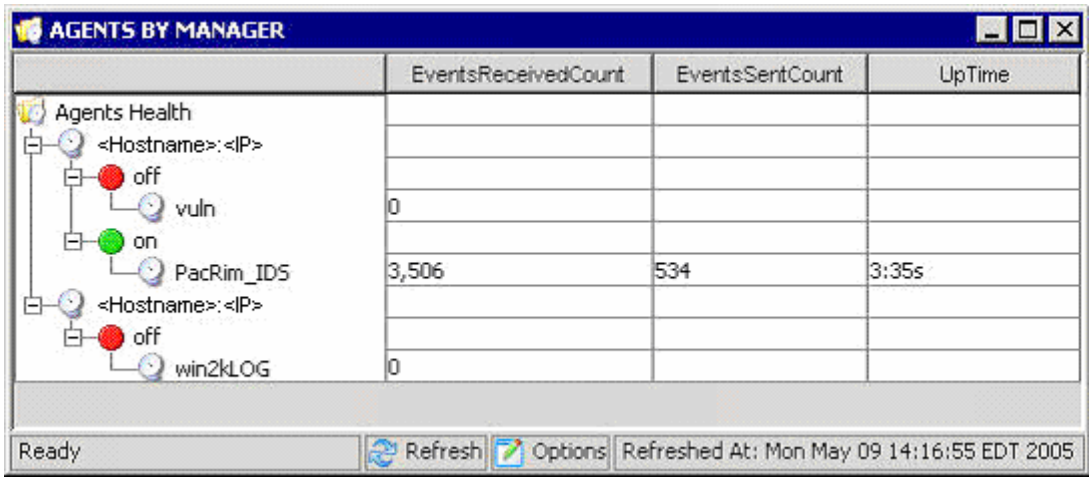


## Layout

The left panel in the Agents tab contains a tree of views. By default, the root of the true has two children: Agent Manager Views and Agent View. The right panel displays views in tables. Each view in the right panel has an entry in the tree on the left.

There are four views displayed in the right panel:

- Agent View
  - Agent View Manager
- Agent Manager View
  - Agent Manager View Manager

The Agent View displays information about agents and the Agent Manager View displays information about agent managers. Each view is displayed as a tree table: the object is grouped by one or more of their attributes. The configuration of the view is adjustable. The options of a view can be changed and new view types can be added. The view configuration is displayed in a View Manager (Agent View Manager or Agent Manager View Manager).

When the tab is first displayed the tree in the left panel is populated with the two view managers and the agent view manager is displayed in the right panel.

The Agent View Manager has 3 pre-configured view options by default; new ones can be created. The three are: All Agents, Agents by Manager and Agents by Status.

The All Agents view displays all the agents grouped by the manager in which they are running.

The Agents by Manager view groups all the agents by their manager and further groups them by their status (on or off) within each manager.

The Agents by Status view groups all the agents by status (On or Off) and then within each status they are grouped by manager.

There is one default view for viewing Agent Managers and that is the All Managers view. It displays all the active agent managers in the system with no grouping.

## Monitoring an Agent

In the Wizard Host Window, by default you can monitor the following:

**Agent Manager View Manager**

- StartTime — Time when Agent Manager was started, given in mm/dd/yy hh:mm:ss and time zone
- UpTime — Length of time that Agent Manager has been running, given in days, hours, minutes and seconds.
- EventReceivedCount — Number of events received from all agents by Agent Manager since Agent Manager started.
- EventReceivedRate — Average event rate per second that Agent Manager received in the last minute.

**Agents View Manager**

- Status — on or off
- EventsReceivedRate — Average event rate per second that the Agent Port received in the last minute.
- EventsReceivedCount — Number of events received by the Agent Port since the Agent Port started.
- UpTime — Length of time that the Agent Port has been running, given in hours, minutes and seconds.

You can create your own views with that have additional or less fields.

### Monitoring a Wizard Host

Monitoring a Wizard Host

1. Click the Agents tab.
2. Click the Agent Manager View Manager button.

3. Select a view option by double-clicking on a view or create a new view. A Wizard Host window will display.

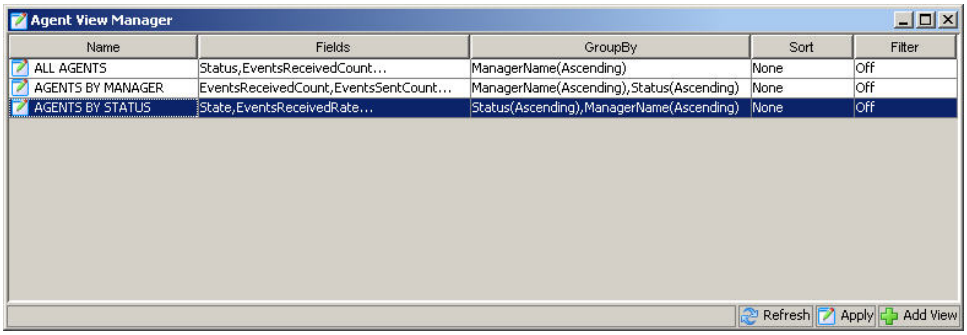


## Creating an Agent View

### Creating an Agent View

1. Click the Agents tab.
2. Click the Agent Manager View Manager button.

3. To create a new view, click the Add View button.
   - Enter your Option Name
   - To arrange which fields you want shown, click the Fields button
   - To group different titles, click the Group button
   - To sort by title, click the Sort button
   - To filter, click the filter button

   Below is a view set with Group set to ManagerUUID and by Port name.

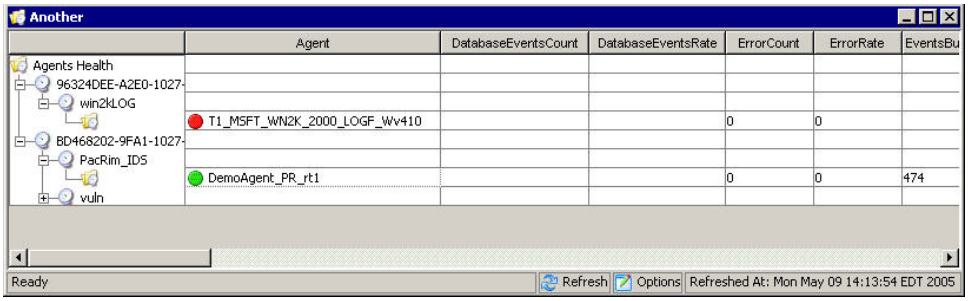


## Modifying an Agent View

### Modifying an Agent View

1. Click the Agents tab.
2. Double-click on any of the names.
3. Click the Options button. In this window you may also set your:
   - Fields...
   - Group by...
   - Sort...
   - Filter...
   - Tree Display
4. Click Apply and Save.

The following is view with Tree Display set to Manager Type.

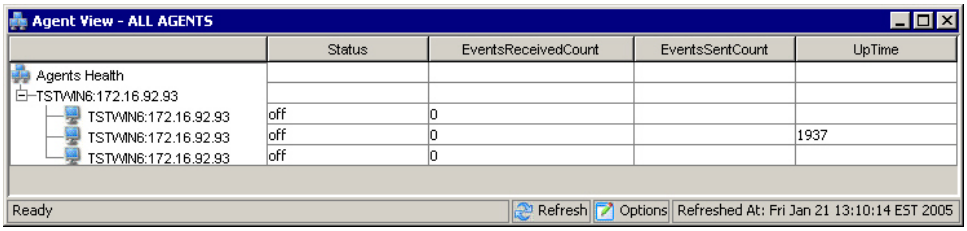


## Stopping/Starting/Details Agents

### Stopping/Starting Agents

1. Click the Agents tab.
2. Open an Agent View Manager.
3. To stop/start/Show Detail a single agent, right-click on an Agent > Actions > Start or Stop.

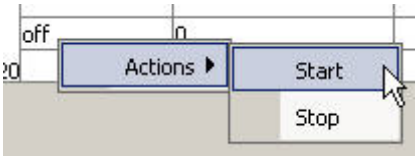

# Chapter 9 – Admin Tab

To use this feature, you must have the proper permission. If this permission is not assigned, none of the other permissions related to actions using this tab will be available.

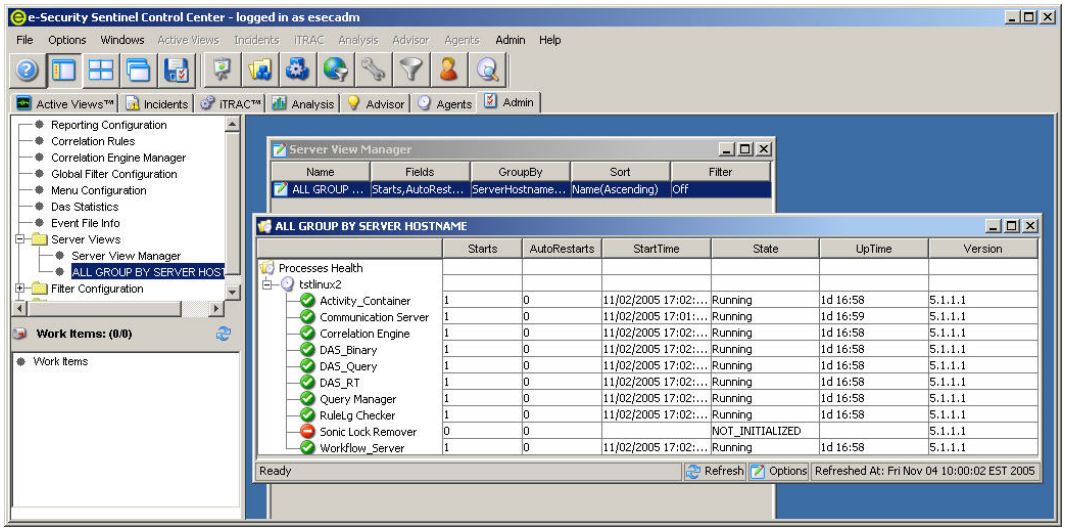


## Admin Tab - Description

The Admin Tab allows you access to:

- Reporting Configuration for the Analysis and Advisor Reports
- Manage filters
- Working with Sentinel Correlation Rules
- Configure Menu Configuration Menu
- DAS Statistics
- Event File Information
- Server Views (Linux Only)
- Configure user accounts

## Reporting Configuration Options for Analysis and Advisor Reports

### To configure the URL for Analysis and Advisor Reports

1. Click the Admin tab.
2. In the Admin Navigator, click Reporting Configuration.
3. In the Reporting Configuration window, click Modify.
   - In the Analysis URL box, enter the URL for the Crystal Enterprise Server and click Refresh.

     ```
     http://<IP>/GetReports.asp?APS=<IP>&user=Guest&pass
        word=&tab=Analysis
     ```

**NOTE**: <IP> is the IP address of the Crystal Enterprise Server.

   - In the Advisor URL box, enter the URL for the Crystal Enterprise Server and click Refresh.

```
http://<IP>/GetReports.asp?ASP=<IP>&user=Guest&pass
   word=&tab=Advisor
```

**NOTE**: <IP> is the IP address of the Crystal Enterprise Server.

For more information, see the Installation Guide.

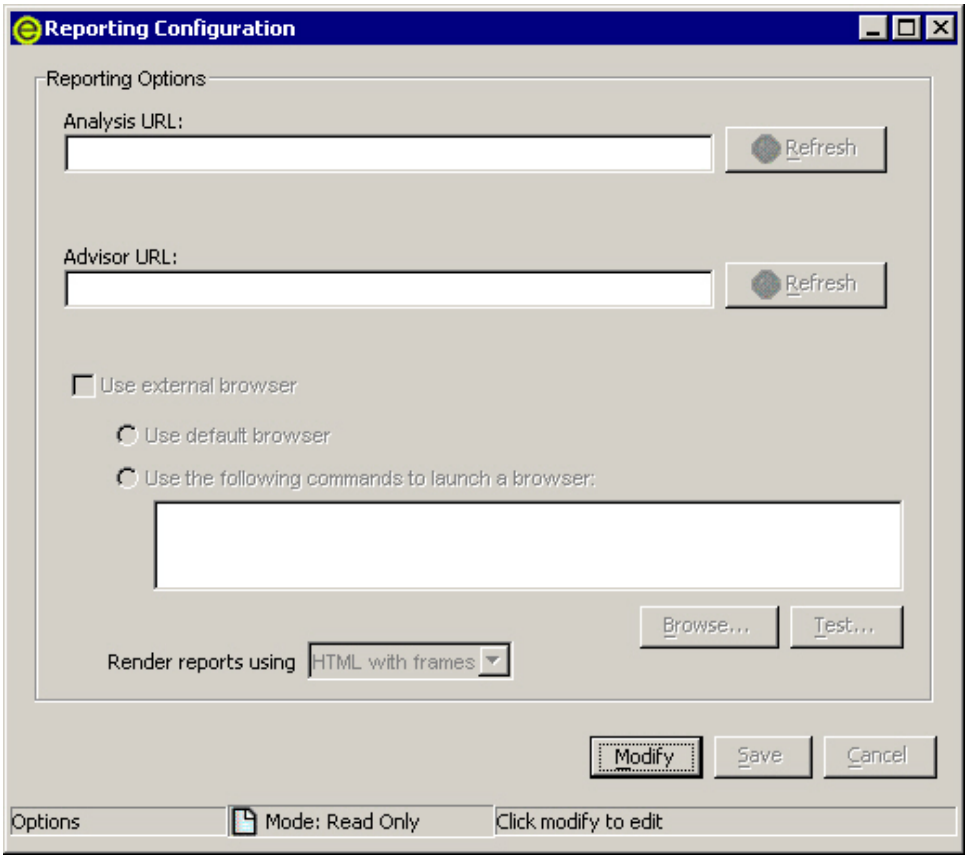


The external browser option allows you to use your default or another browser. When using a browser other than the default browser, your command line must be followed by a %URL%. For example:

```
C:\Program Files\Internet Explorer\IEXPLORE.EXE
   %URL%
```

4. Wait for the Refresh button to turn green and click Save. You will have to logout of the Sentinel Control Center and log back in.

## Sentinel Correlation Rules

Correlation adds intelligence to security event management by enabling you to automate analysis of the incoming event stream to find patterns of interest. Correlation allows you to define rules that identify critical threats and complex attack patterns so that you can prioritize events and initiate effective incident management and response.

Rule Folders are the logical grouping of correlation rules. Grouping correlation rules into Rule Folders also allows you to have a set of rules that runs during the business day or a set that runs at night and another set that runs during the weekend. In essence, watching for different activities based on time of day.

For example, you can enable all of the daytime correlation rules at once at 8 A.M. on Monday through Friday and also to disable the nighttime correlation rules all at the same time. Specifically, If you don't need to group correlation rules into Rule Folders, you can create only one Rule Folder and then create all of your correlation rules under that Rule Folder.

There is no limit to the number of users that can access Correlation Rules. When more than one user is editing the same rule, the last person to save will overwrite all previous saves.

This section discusses:

- Rule Folders and Rules
- Correlation Rule Types
- Correlation Engine Rule Deployment
- Importing and Exporting of Correlation Rules
- Role of the Database in Storing Correlation Rules
- Logical Conditions

**NOTE**: You cannot correlate on a null (empty) value.

## Rule Folders and Rules

The following defines the relationship between Rule Folders and Rules. Rule Folders and Rules are displayed in hierarchical fashion in the Correlation Rules window.

- A Rule Folder may contain zero or more rules
- The number of Rule Folders and Rules is only limited by available disk (storage) space
- Double-clicking a Rule Folder displays the Rule Editor for that type of correlation rule
- The maximum length of Rule Folder names is 255 characters for the folder path and rule names is 255 characters
- Rule Folder and Rule descriptions can be up to 1024 characters

## Correlation Rule Types

There are four Correlation Rule types that you can choose when defining rules. They are:

- Watchlist
- Basic Correlation
- Advanced Correlation
- Free Form RuleLg

**CAUTION:** You should be familiar with the RuleLg correlation rule definition language before using this correlation rule type. In addition, if you renamed a tag, do not use the original name when creating a correlation rule using RuleLg.

### Watchlist

There are four different Filter Types to choose from. They are:

- Allow All - Lets all events through.
- Pattern - Any regular expression with a grep-like syntax.

- Filter Manager - A drop-down list that displays the Filter Manager to select or create a new filter.

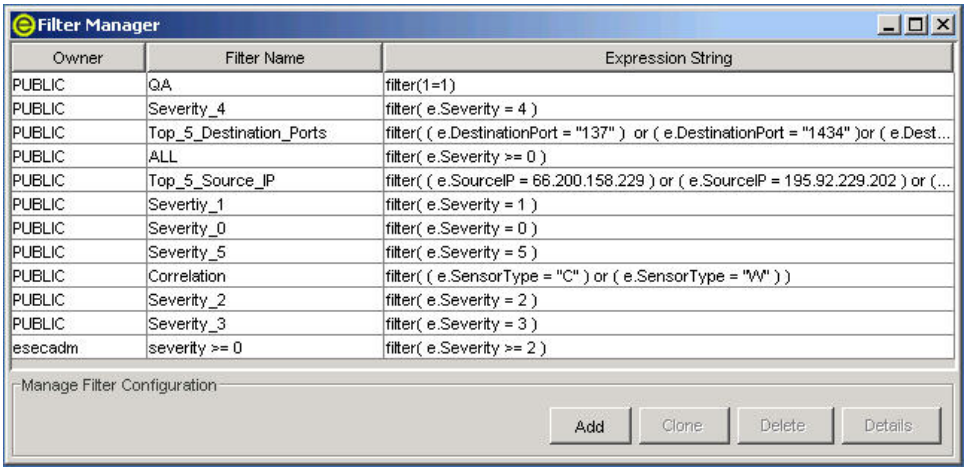


- Builder - Create criteria for inclusion and exclusion of events based on boolean algebra. Two panes are available (include and exclude). Enter your values here, for example:

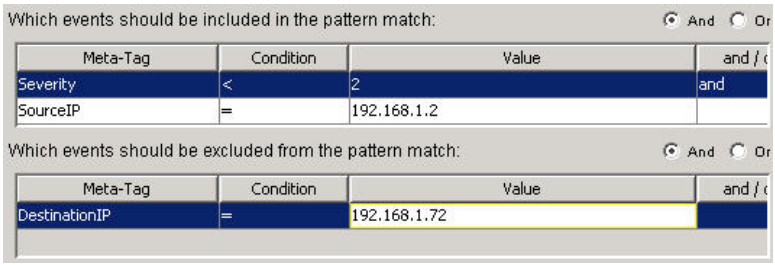


### Basic Correlation

There are four different Filter Types to choose from. They are:

- Allow All - Lets all events through.
- Pattern - Any regular expression with a grep-like syntax.
- Filter Manager - A drop-down list that displays the Filter Manager to select or create a new filter.
- Builder - Create criteria for inclusion and exclusion of events based on boolean algebra.

This rule allows you to count the number of times certain conditions are met within a specific timeframe.

For example, a Basic Correlation rule can look for the same source IP address reported five times in five minutes, even if the events are reported from different devices, such as an intrusion detection system (IDS) and a firewall.

### Advanced Correlation

There are four different Filter Types to choose from. They are:

- Allow All - Lets all events through.
- Pattern - Any regular expression with a grep-like syntax.
- Filter Manager - A drop-down list that displays the Filter Manager to select or create a new filter.

- Builder - Create criteria for inclusion and exclusion of events based on boolean algebra.

This rule allows you to:

- Count the number of times certain conditions are met within a specific timeframe.
- Incorporate all of the features of the simple correlation rule, as well as evaluate events against past events.

For example, an Advanced Correlation rule can look for events from the same source IP address to the same destination IP address with the same event name that occur both inside and outside a firewall (meaning the attack may have made it through the firewall).

**Free Form RuleLg Correlation**

RuleLg correlation rule definition language allows you to have complete control in defining correlation rules. Before you use this correlation rule type, you should be familiar with the RuleLg correlation rule definition language.

### Correlation Engine Rule Deployment

To use this feature, you must have the user permission Start/Stop Correlation Engine. The Correlation Engine is in one of two states, activated or deactivated. The current state is displayed in the icon.

- Activated - 
- Deactivated - 

When the Correlation Engine is activated it is processing active correlation Rules Folders.

When the Correlation Engine is deactivated, all of its in-memory data is preserved and no new correlation events are generated. This state is equivalent to deactivating all of the Rule Folders. Deactivating the Correlation Engine does not affect other parts of the system. Incoming events still go through, populating the Sentinel database.

### Importing and Exporting of Correlation Rules

The exporting capability allows e-Security to create and export “canned” correlation rules and make them available to you for importing into your system. These XML documents are formatted specifically for the Correlation Engine. These pre-packaged rules are built by e-Security and are available on the Customer Portal at http://www.esecurityinc.com.

The ability to export rules as XML documents assists you when you need e-Security’s help in troubleshooting your correlation rules. Exporting is also beneficial when you have an “in production” Sentinel and a “development” Sentinel. You can develop and test correlation rules in the development environment and then export them to the production environment. The file extension for exported correlation rules is .crf.

### Role of the Database in Storing Correlation Rules

When you activate the Correlation Engine (a Sentinel Server process) in Sentinel Control Center, it requests the deployment information and rules from the

database. When you modify correlation rules and then save them, they are sent to the database for storage. The changes in the rule will not be reflected in the Correlation Engine, unless one of the following is met:

- the deployed rule is disabled and then enabled
- the rule is freshly deployed

When you modify deployment rules and then save them, they are sent to the database for storage and to the Correlation Engine where they are put into use.

### Logical Conditions for Correlation Rules

The following are logical conditions used when creating correlation rules. For more information about Meta-tags, see the e-Security User's Reference Guide.

| Condition | Type Field | Description |
|---|---|---|
| = | numeric string | The content of the meta-tag selected is equal to the value entered. |
| != | numeric string | The content of the meta-tag selected is not equal to the value entered. |
| < | numeric | The content of the property selected is less than the value entered. |
| > | numeric | The content of the meta-tag selected is greater than the value entered. |
| <= | numeric | The content of the meta-tag selected is less than or equal to the value entered. |
| >= | numeric | The content of the meta-tag selected is greater than or equal to the value entered. |
| =Meta-Tag | numeric string | The contents of the meta-tag selected in the drop-down list on the left is equal to the contents of the meta-tag selected on the right of the expression. |
| !=Meta-Tag | numeric string | The contents of the meta-tag selected in the drop-down list on the left is not equal to the contents of the meta-tag selected on the right of the expression. |
| <Meta-Tag | numeric | The contents of the meta-tag selected in the drop-down list on the left is less than the contents of the meta-tag selected on the right of the expression. |
| >Meta-Tag | numeric | The contents of the meta-tag selected in the drop-down list on the left is greater than the contents of the meta-tag selected on the right of the expression. |
| <=Meta-Tag | numeric | The contents of the meta-tag selected in the drop-down list on the left is less than or equal to the contents of the meta-tag selected on the right of the expression. |
| >=Meta-Tag | numeric | The contents of the meta-tag selected in the drop-down list on the left is greater than or equal to the contents of the meta-tag selected on the right of the expression. |
| =Regex | numeric string | Use a period (.) and asterisk (*) with string for the value. |

| Condition | Type Field | Description |
|---|---|---|
| Subnet | numeric string | A match subnet operation will match if the IP address being compared is in the same subnet as specified in the match subnet operation. |

## Opening Correlation Rules Window

The Correlation Rules window allows you to:

- New Folder – to create a new Rule Folder
- New Rule - create a Rule for a Rule Folder
- Copy a Rule Folder – this allows you to modify copied Rule Folders or Rules while saving the original Rule Folder or Rule
- Delete a Rule Folder or Rule – you cannot recover a deleted Rule Folder or Rule after you confirm deletion
- Rename – rename a Rule or Rule Folder
- Import a Rule Folder – a browser window will open
- Export a Rule Folder – a browser window will open exporting the Rule Folder as an xml file.
- Edit – allows editing and previewing of rules and folder properties

### Opening the Correlation Rules Window

1. Click the Admin tab.
2. In the Admin Navigator, click Correlation Rules.

## Copying and Creating a Rule Folder or Rule

### Creating a Rule Folder

1. Open the Correlation Rules window.
2. Select the parent folder to contain the new folder.
3. Right-click > New Folder.
4. Type in the Rule Folder name, limited to 255 case sensitive characters with no periods.
5. (Optional) Type in the Description of the rule, limited to 1024 characters.
6. Click OK.

### Creating a Rule

1. Select the parent folder to contain the new rule.
2. Right-click > New Rule.
3. Rule Wizard opens, select one of the following rule types:
   - Watchlist
   - Basic Correlation
   - Advanced Correlation
   - Free Form

**NOTE:** For descriptions of the rule types, go to section Correlation Rule Types.

4. Click Finish.

## Deleting a Correlation Rule Folder or Rules

Deleting a Correlation Rule Folder or Rule

1. Open the Correlation Rules window.
2. Select the Rule Folder or Rule you wish to delete.
3. Right-click > Delete.
4. A confirmation box will appear:
   - Yes – when deleting a Rule Folder, the Rules within that Rule Folder will also be deleted. You cannot recover a deleted Rule after you click OK.
   - No – will return you to the Correlation Rule window

## Importing or Exporting a Correlation Rule Folder

Importing or Exporting a Correlation Rule Folder

1. Open the Correlation Rules window.
2. Select a Rule Folder.
3. Right-click > [Import Rule Folder or Export Rule Folder]
   - Import – A file browser will open, browse to the Rule Folder to import and click OK.
   - Export – A file browser will open, browse to the target device to write the Rule Folder to and click OK. The Rule Folder will be exported as a crf file.

## Editing in the Correlation Window

Editing in the Correlation Window

1. Open the Correlation Rules window.
2. Right-click > Edit.
3. Edit the rule and click Finish.

## Activating or Deactivating a Correlation Engine

Activating or Deactivating a Correlation Engine

1. Open the Correlation Engine Manager window.
2. Highlight and right-click on a Correlation Engine > Activate or Deactivate Engine.

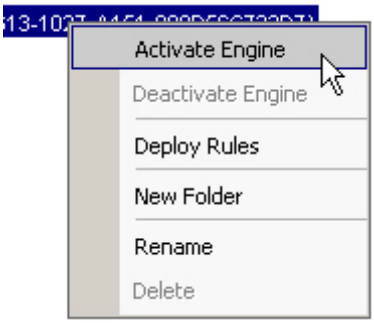

## Deploying Correlation Rules

### Deploying Correlation Rules

1. Open the Correlation Engine Manager window.

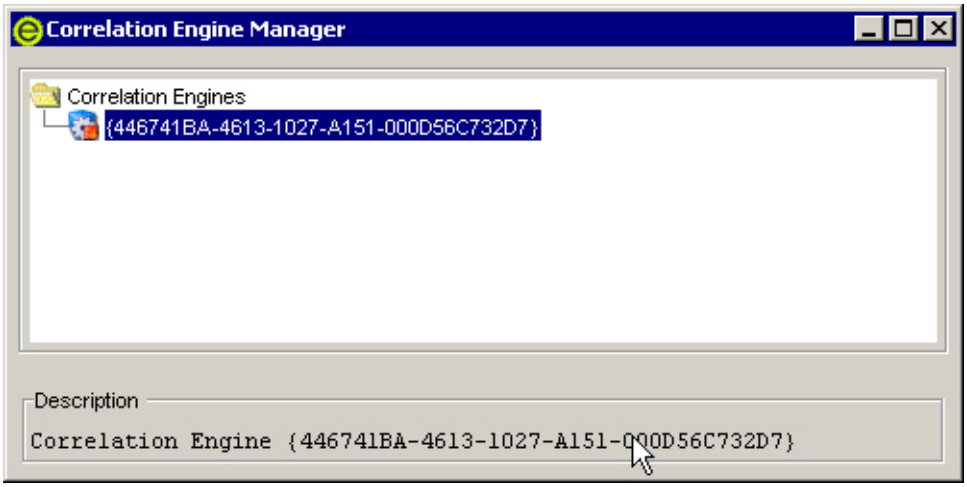


2. Right-click (any folder in the window or highlight the engine to have the rule deploy to there) > Deploy Rules.
3. Place a check mark next to the rules you want to deploy. Click OK.

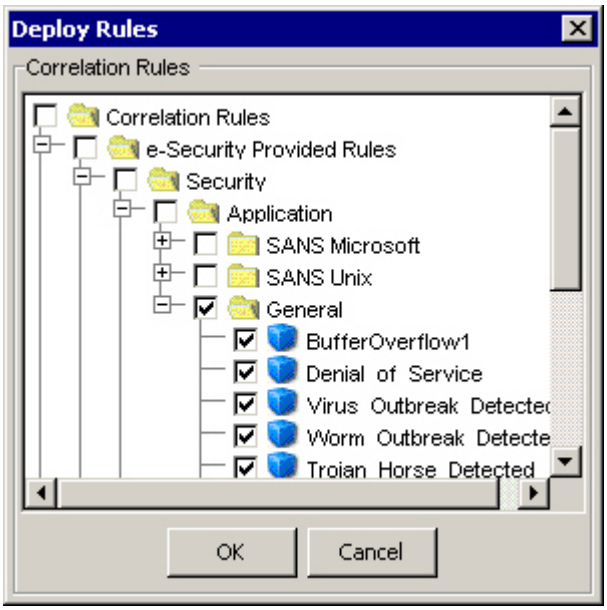


4. To start your rule, you must move the rule to under a correlation engine.

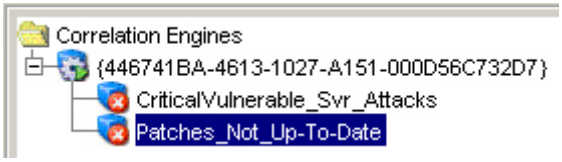


**NOTE**: Rules are deployed enabled.

5. Under the correlation engine, highlight your rule and right-click > Enable Rule.

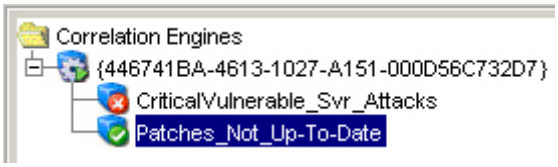


## Server Views (Linux)

For the Linux Sentinel Servers, Server Views allows you to:

- monitor the status of all Sentinel Server processes across the system
  - Activity Container
  - Communication Server (the server cannot be stopped using this feature)
  - Correlation Engine
  - DAS_Binary
  - DAS_Query
  - DAS_RT
  - Query Manager
  - RuleLg Checker
  - Workflow_Server

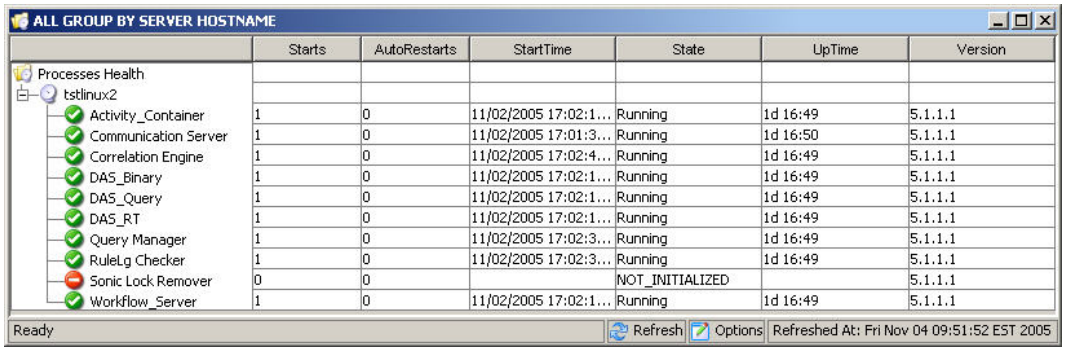


- start, stop or restart processes

The terms Starts and AutoRestarts, in the context of the Server View, are defined as follows:

- Starts – The number of times the process was started, for whatever reason. This includes starts initiated by the user via the GUI or done automatically.
- AutoRestarts – The number of times the process was automatically restarted. Since this only applies to purely automatic restart scenarios, it does not apply to restarts initiated by a user. This field is helpful for determining if the process exited (e.g. – due to an error) and was automatically restarted by Sentinel Watchdog.

### Monitoring a Process (Linux)

Monitoring a Process

1. Click the Admin tab.
2. Click the Server View button.

3. Double-click on a view. A view will appear.

4. Expand the server view. All the processes will list.

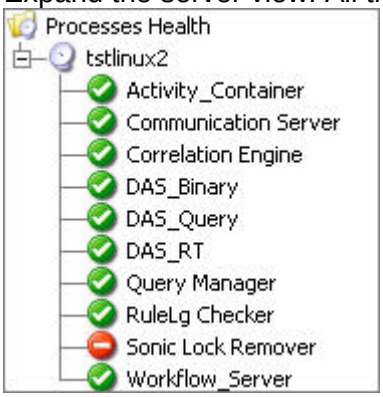


## Creating a Server View (Linux)

Creating a Server View

1. Click the Admin tab.
2. Click the Server View button.

3. To create a new view, click the Add View button.
   - Enter your Option Name
   - To arrange which fields you want shown, click the Fields button
   - To group different titles, click the Group button
   - To sort by title, click the Sort button
   - To filter, click the filter button
4. Click Ok and then click Save.

## Starting, Stopping and Restarting Processes (Linux)

You cannot stop the Communication Server using this feature.

Starting, Stopping and Restarting Processes

1. Click the Admin tab.
2. Click the Server View button.

3. Double-click on a view. A view will appear.

4. Expand the server view. All the processes will list.

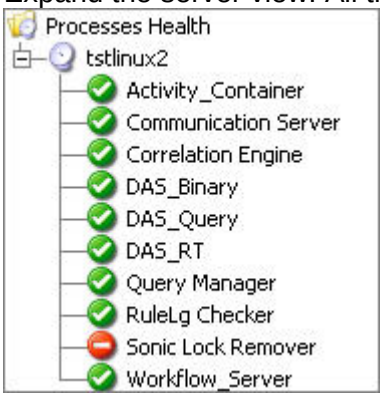


5. Select a process, right-click > Actions > select a function (Start, Restart or Stop).

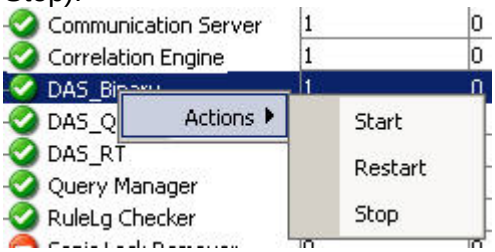


## Filters

Filters allow you to process data based on specific criteria for events in real-time and for users of the system. Filters enable you to manage data seen in the Sentinel Control Center. The Filter Engine drives the Real Time Event windows by maintaining the data structure for each security filter. Filters prevent users from viewing unauthorized events and drop events that users don't wish to see. Filters are created in the Admin tab of the Sentinel Control Center.

**NOTE**: The following are invalid filter name characters: $ # . * & : < >.

There are three types of filters:

- Public Filters
- Private Filters
- Global Filters

### Public Filters

Public filters are system-owned. Public filters can be used as security filters or display filters. Security filters are based on user permissions. Display filters determine which events are depicted in the real time event tables, charts and graphs.

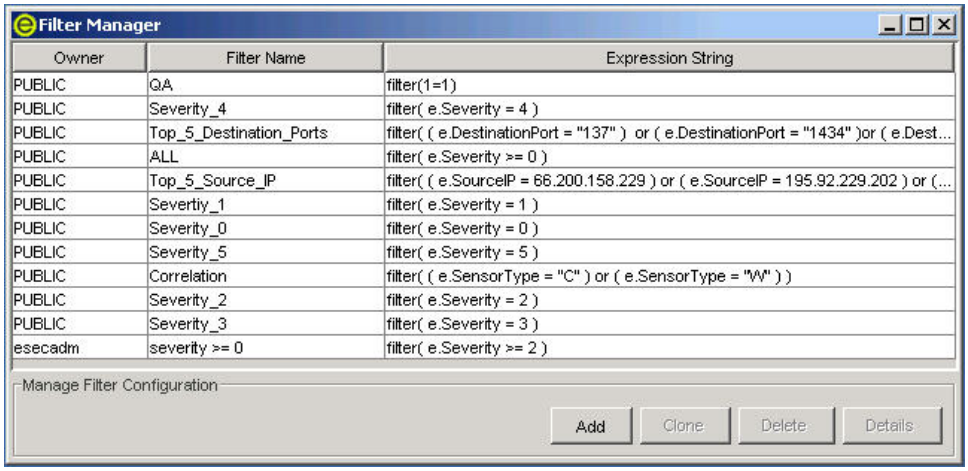


## Private Filters

Private filters are user-owned. Private filters are display filters and are shareable if you have the View Private Filters permission.

## Global Filters

Global filters are classified as Public Filters. Global filters are processed at the Agent Manager sequentially for each event until a match is found. Global filter evaluation stops for that event and the matched global filter action is taken for that event. The order of evaluation of global filters is top to bottom, as shown in the Console. They can be enabled or disabled as needed.

Global filters do the following:

- Enable a global action on events, such dropping events, routing events to the database only or routing events to the database and the Sentinel Control Center
- Are processed by Wizard's Agent Manager
- Are configured in the Admin tab under the Global Filter Configuration option where they can be enabled and disabled
- Drop events
- Can route events to the database only
- Can route events to the database and to the Sentinel Control Center

Through the Global Configuration window, you can:

- create Global Filter
- rearrange a Global Filter
- delete a Global Filter

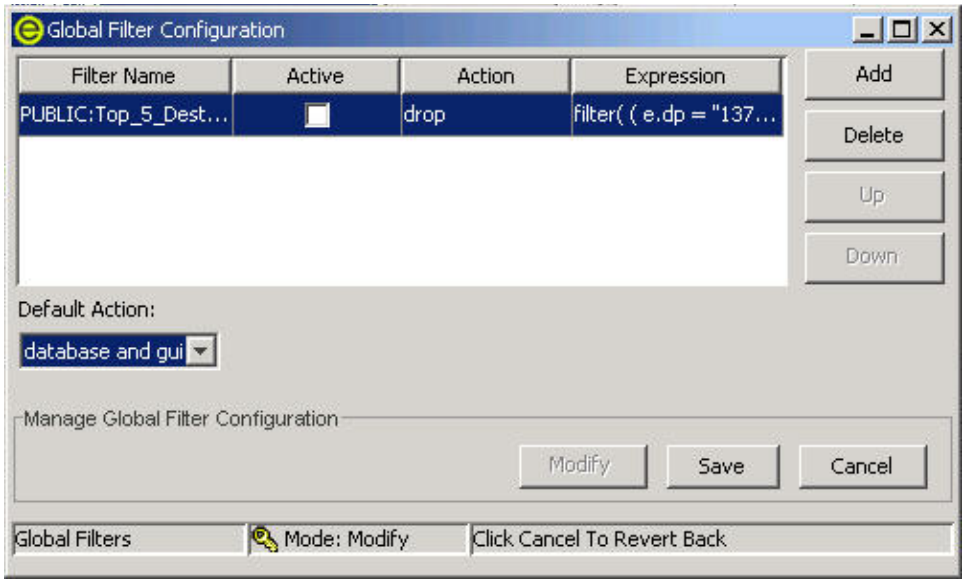


## Creating a Global Filter

### Creating a Global Filter

1. Click the Admin tab.
2. Click Admin > Global Filter Configuration or select Global Filter Configuration in the navigation tree.
3. In the Global Configuration window, click Modify and click Add.
4. In the new blank row, click the Filter Name column.
5. Select a filter and click Select or Add (if you need to create a filter).
6. In the Active column, click the Active box.
7. In the Action column, select the action that the global filter will have on events that pass this global filter. If an event does not meet any of the active global filters, then the default action determines how the event is handled.

   You can set the Default Action box to one of the following:

   - drop – events will not go to the Sentinel Control Center or the Sentinel Server database
   - database - events will be sent directly to the database, bypassing the Sentinel Control Center
   - database and GUI - events will be sent to the Sentinel Control Center and Sentinel Server database
8. Continue adding filters until you are finished.
9. Click Save.

## Rearranging Global Filters

### Rearranging Global Filters

1. In the Global Configuration window, click Modify.
2. Select a filter and click Up or Down to move it to a different location on the list.

3. Click Save.

### Deleting a Global Filter

**NOTE:** When deleting a Global Filter, you will not get a confirmation message.

To delete a global filter

1. In the Global Configuration window, click Modify.
2. Select a filter from the list and click Delete.
3. Click Save.

## Configuring Public and Private Filters

Configuring Public and Private filters allows you to:

- Add a Filter
- Clone a Filter
- Modify a Filter
- View the Details of a Filter
- Delete a Filter

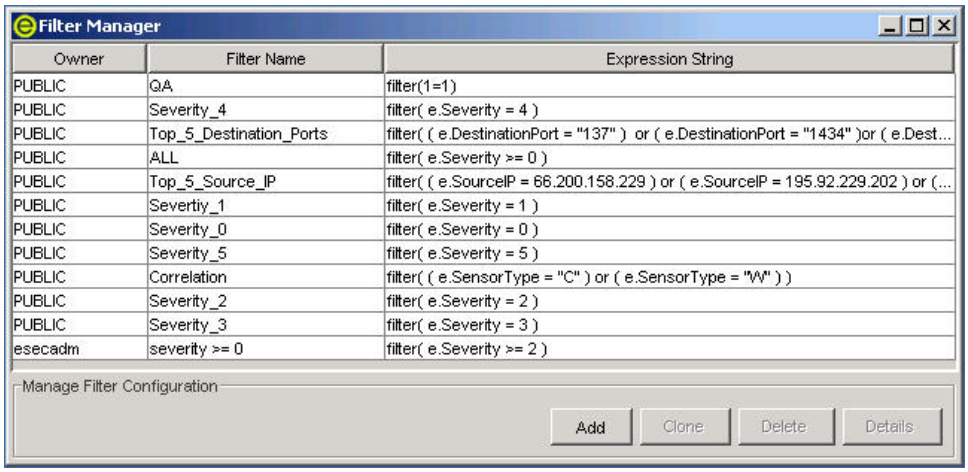


### Adding a Filter

To add a public and private filter

1. Click the Admin tab.
2. Click Admin > Filter Manager or select File Manager under the Filter Configuration folder in the navigator.
3. Open the Filter Manager window.
4. Click Add.
5. Select an Owner ID (public or private [user owned]).

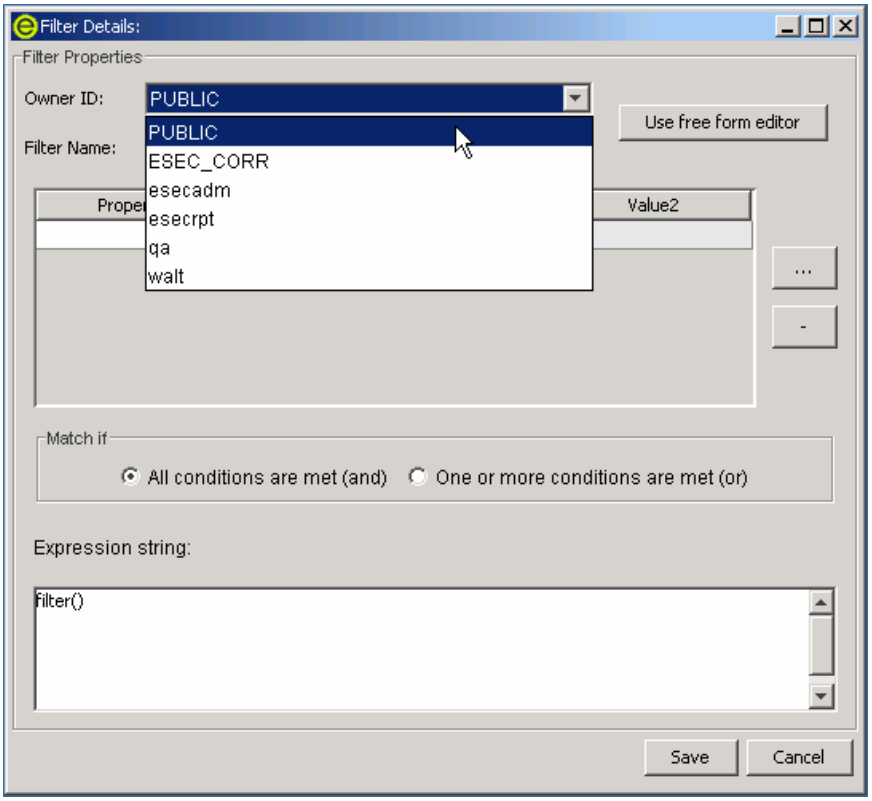


6. Enter a Filter Name.
7. The table editor is the default selection for editing the contents.

**NOTE**: Optionally, you can click Use free form editor to display a free form editor. The free form editor allows you to create complex expressions not possible with the table editor. However, once the expression is modified with the free form editor, the table editor cannot be used with the expression.

8. Select the criteria for the following columns:
   - Property
   - Operator
   - Value columns.

   Your choices display in the Expression string box.
9. In the Match if box, click either:
   - All conditions are met (and)
   - One or more conditions are met (or)
10. To create another filter expression, click the Create a New Filter Expression button (+) to add another row to the filter expression table.
11. To remove a filter expression, select a filter expression from the table and click the Remove the Selected Expression button (-).
12. Click Save.

**To clone a Public and Private filter**

Cloning is a convenient way to duplicate a filter to assure consistency of criteria among a group of filters or users.

To clone a public and private filter

1. Open the Filter Manager window.
2. Click Clone.
3. Enter a new filter name.
4. Change any the original filter's criteria.
5. Click Save.

**Modifying a Public and Private Filter**

To modify a Public and Private filter

1. Open the Filter Manager.
2. Select a filter and click Details.
3. Change any of the criteria as desired. You will not be able to change the Owner ID and the Filter Name.
4. Click Save.

**Viewing the Details of a Public and Private Filter**

To view a public or private filter

1. Open the Filter Manager window.
2. Select a filter and click Details.

**Deleting a Public and Private Filter**

To delete a Public and Private filter

1. Open the Filter Manager window.
2. Select a filter and click Delete.
3. A confirmation window will open.

## Configure Menu Configuration

To use this feature, you must have the user permission Menu Configuration.

Use the Menu Configuration window to create the menu items that appear on the Event menu, which displays on any table displaying an event (e.g., Event Real Time window, Snapshot window, Incidents Events window, etc…) when you select one or more events and right click. Sentinel has the following default Menu Configuration items that you can clone, activate or deactivate:

- ping - Ping the destination IP of the selected event
- nslookup - Perform an nslookup on the Source IP of the selected event
- traceroute (tracert on MS SQL) - Perform a traceroute from the Source IP of the selected event to the Sentinel Server
- Whois? - Perform an ARIN Whois? lookup on the Source IP of the selected event

Menu Configuration allows you to:

- Adding an Option to the Menu Configuration Menu
- Cloning an Menu Configuration Option
- Modifying an Menu Configuration Option

- Viewing an Menu Configuration Option's Parameters
- Activating or Deactivating an Menu Configuration Option
- Rearranging Event Menu Options
- Deleting a Menu Configuration Option
- Add a browser feature to your Menu Configuration Option

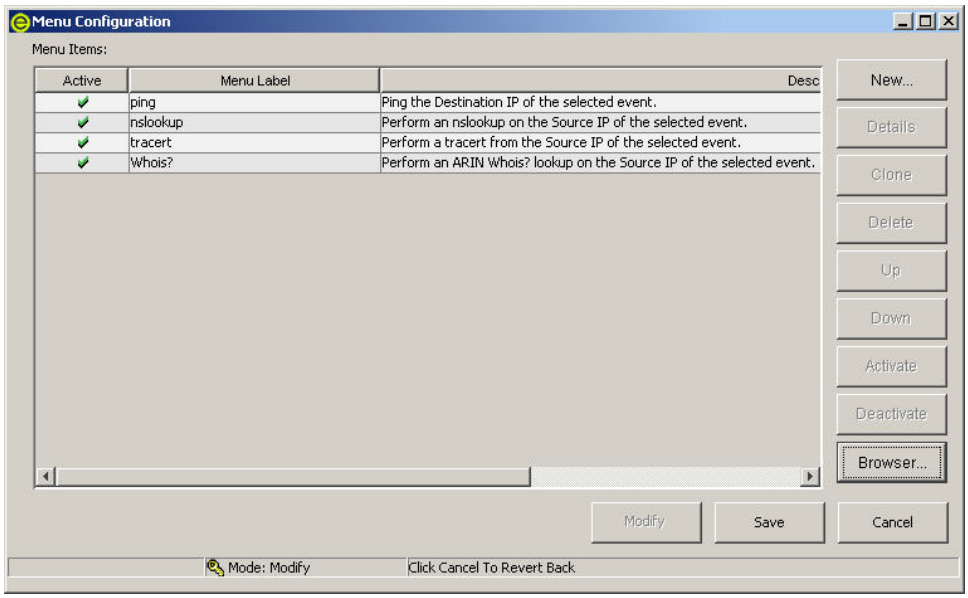


Menu Configuration for Windows and Solaris

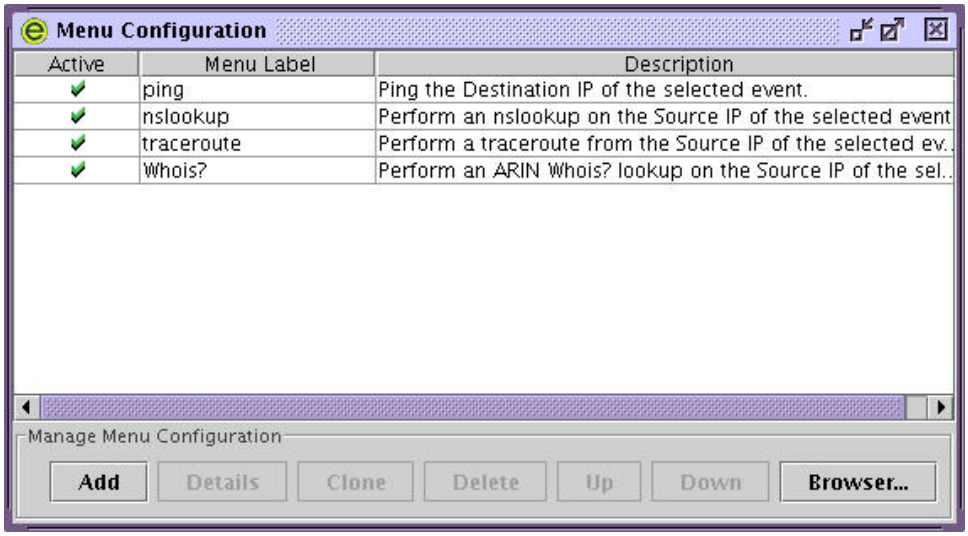


Menu Configuration for Linux

## Adding an Option to the Menu Configuration Menu

**NOTE**: If you renamed a tag, such as renaming CustomerVar24 to PolicyName, you must use the new name when setting parameters.

To add an option to the Menu Configuration menu

1. Click the Admin tab.
2. In the Admin Navigator, click Admin > Menu Configuration.

3. (For Windows and Solaris) In the Menu Configuration window, click Modify and then New.
4. In the Menu Configuration dialog box, enter:
   - Name
   - Description
   - Action – either execute a command or launch a browser
   - Use browser - If you chose the Action “Execute Command” and your Browser settings are setup to “Use External Browser” (see Editing Your Menu Configuration Browser Settings for editing Browser settings), you have the option to select Use browser. Selecting this option will cause the output of your command to be displayed using the Menu Configuration Browser settings for your Sentinel Control Center.
   - File Type - If you chose the Action “Execute Command”, your Browser settings are setup to “Use External Browser”, and you selected the option “Use browser”, you have the option of setting the File Type for the output of this command.
   - Command line/URL

**NOTE**: For UNIX, the script/application or symbolic link to the script/application must be located in the $ESEC_HOME\sentinel\exec directory. For any script, application or symbolic link, only enter the command. Any path entered will be ignored.

**NOTE**: For Windows (correlation), the script/application must be located in a one of the directories listed in your Windows Environmental Variables. Any path entered will be ignored.

**NOTE**: For Windows (non-correlation), entering a path is optional. Entering a command without a path will default to %ESEC_HOME%\sentinel\bin and all other paths specified in your environmental variables.

   - Parameters – must be enclosed by the percent sign (e.g., %EventName%)

**NOTE:** For a list of available tags you can use when specifying parameters, click Help on the Menu Configuration dialog box or go to the Meta-tag chapter in the e-Security User's Reference Guide.

5. Click OK. The new option is added to the list of menu items in the Menu Configuration window.
6. (For Windows and Solaris) Click Save.

**NOTE:** For an example, highlight any of the default menu items and click Details. The following is an nslookup configuration:

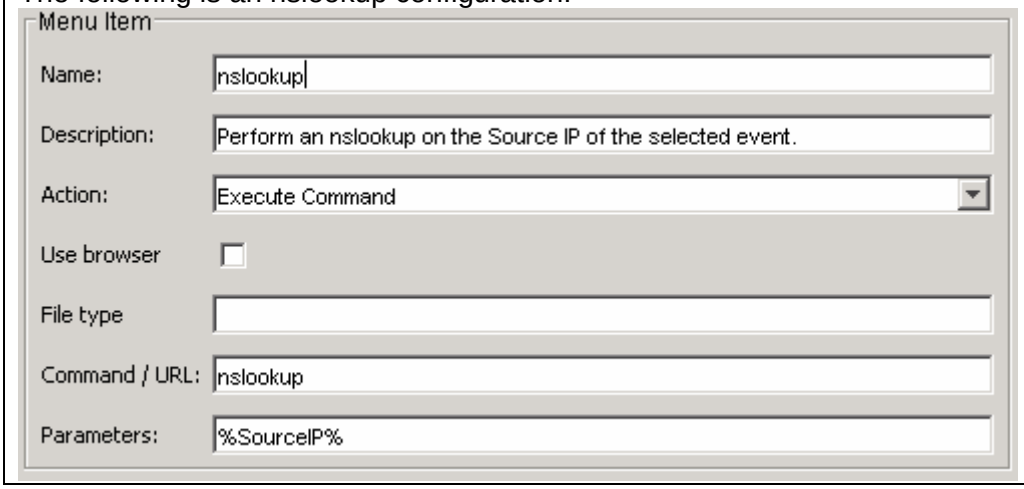


## Cloning an Menu Configuration Menu Option

To clone an Menu Configuration menu option

1. Open the Menu Configuration window.
2. (For Windows and Solaris) Click Modify.
3. Select a menu item from the table and click Clone.
4. In the Menu Configuration dialog box, edit:
   - Name
   - Description
   - Action
   - To use a browser or not. For information, see Add a browser feature to your Menu Configuration Option.
   - Command line/URL
   - Parameters
   - Select an action:
     - Execute Command
     - Launch Web Browser.

**NOTE:** For a list of available tags you can use when specifying parameters, click Help on the Menu Configuration dialog box or go to the Meta-tag chapter in the e-Security User's Reference Guide.

5. Click OK. The new option is added the list of menu items in the Menu Configuration window.
6. (For Windows and Solaris) Click Save.

## Modifying an Menu Configuration Menu Option

To modify an Menu Configuration menu option

1. Open the Menu Configuration window.

2. (For Windows and Solaris) In the Menu Configuration window, click Modify.
3. Double-click a menu option.
4. Type your desired changes and click OK.
5. (For Windows and Solaris) Click Save.

### Viewing Menu Configuration Option Parameters

To view the parameters for an Menu Configuration menu option

1. Open the Menu Configuration window.
2. (For Windows and Solaris) Click Modify
3. Highlight a menu item and click details.

### Activating or Deactivating an Menu Configuration Menu Option

To activate or deactivate an Menu Configuration menu option

1. Open the Menu Configuration window.
2. (For Windows and Solaris) Click Modify.
3. (For Windows and Solaris) Select a menu option and click Activate or Deactivate button.
4. (For Linux) Select a menu option, right-click and select either Activate or Deactivate.

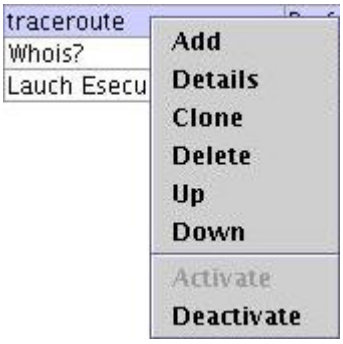


5. (For Windows and Solaris) Click Save.

### Rearranging Event Menu Options

To move an Event menu option up or down

1. Open the Menu Configuration window.
2. (For Windows and Solaris) Click Modify.
3. Select a menu option and click Up or Down.
4. (For Windows and Solaris) Click Save.

### Deleting a Menu Configuration Menu Option

To delete an Menu Configuration menu option

1. Open the Menu Configuration window.
2. (For Windows and Solaris) Click Modify.

3. Select a menu option and click Delete.
   - Click Yes to delete the menu option
   - Click No to retain the menu option
4. (For Windows and Solaris) Click Save.

## Editing Your Menu Configuration Browser Settings

This option allows you to send your Menu Configuration Option output to an external browser. The external browser can be any application. It is not restricted to Internet Browsers. By changing the file extension you can launch whatever application is associated with that extension. For example, txt is usually associated with Notepad. You can also choose to launch a specific program, for example you can have a txt file be opened by wordpad or other editor.

Editing Your Menu Configuration Browser Settings

1. Open the Menu Configuration window.
2. (For Windows and Solaris) Click Modify.
3. Click Browser.

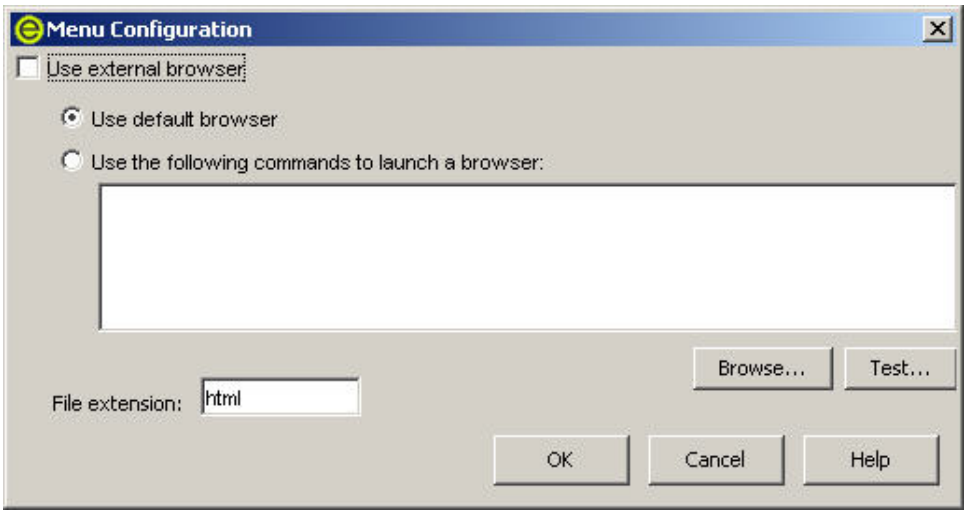


If you select 'Use Browser' when setting up a Menu Configuration Option with the Browser Feature set to the default setting (as above), the Menu Configuration Option will react as if the 'User Browser' box was not checked.

If you check the 'Use external browser' box, you have the option to do one of the following:

- 'Use default browser' - uses the default browser (application) that associated with the file extension set in the File extension field.
- 'Use the following commands to a launch a browser' – allows you to specify a specific application to launch. When using a browser other than the default browser, your command line must be followed by a %URL%. For example:

  ```
  C:\Program Files\Internet Explorer\IEXPLORE.EXE
     %URL%
  ```

The following is an example where the output of the Menu Option will launch into notepad.

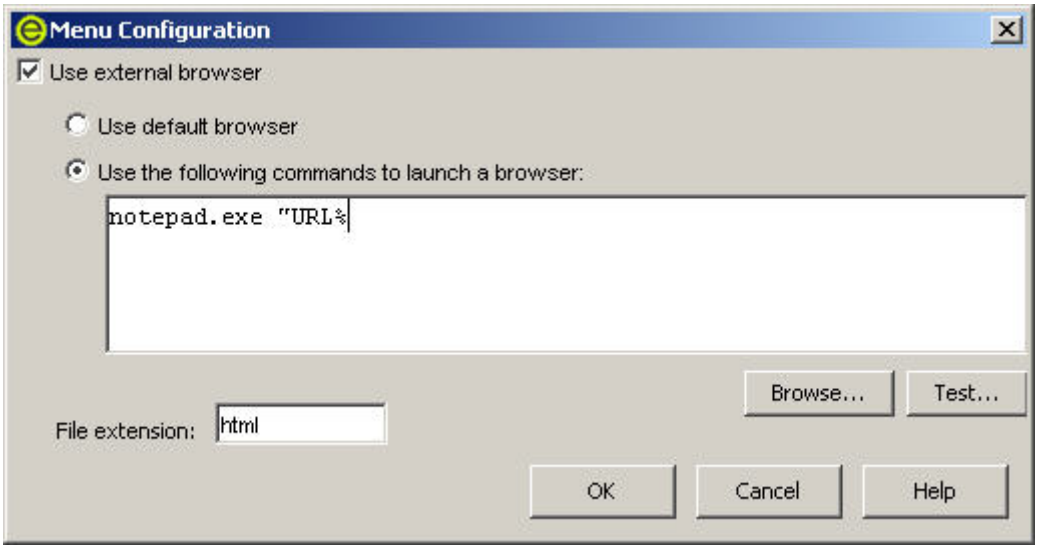


4. After you set your configuration, click OK.

## DAS Statistics

This feature is for internal monitoring of your system. It is not intended for the average user. DAS Statistics monitors the following:

- DAS_Binary
- DAS_Query
- DAS_rt

Statistics are broken down as follows:

- Service – name of service such as DAS_Query
- Time – Time since the last update
- num - number of requests processed for this entry
- WaitTime - average wait time in seconds for a request before its processing starts
- Runtime - average time to process a request (in seconds)
- #wait - Average size of the wait queue
- #run - Average size of the run queue

The information is divided into 3 sections:

- Requests
- Services
- ThreadPools

Under Requests it keeps all the requests by channel (such as services.CorrelationService). Under services it does the same by service. Sometimes it provides a breakdown by appending "<category>" under the name, such as Services.CorrelationService or Services.RemoteObjectService.EMap.getMapPK.

Under Services, all the remote method calls from user defined services (your XML services) are all under services.RemoteObjectService. Under that it puts the name of the service (EMap) in the above example and if asked, the name of the method (getMapPK in the above).

When a request is received by a server, such as DAS Query, a task is created and scheduled. The task is then assigned to a thread pool for execution. There can be more than one thread pool and a thread pool can service multiple services. For that reason, a request may have to wait for an available thread even if the service is not heavily used. If the statistics indicate that the wait time for a request is large and the number of requests for that service is low, check the information about the thread pools.

The numbers next to an entry are the sum for all its children. So requests 15 means that there are 15 requests for all requests method calls. Under that, requests.configurations 1 means that 1 of the 15 are to configurations, requests.esecurity.correlation.config 2 means that 2 of the 15 are to esecurity.correlation.config and so on.

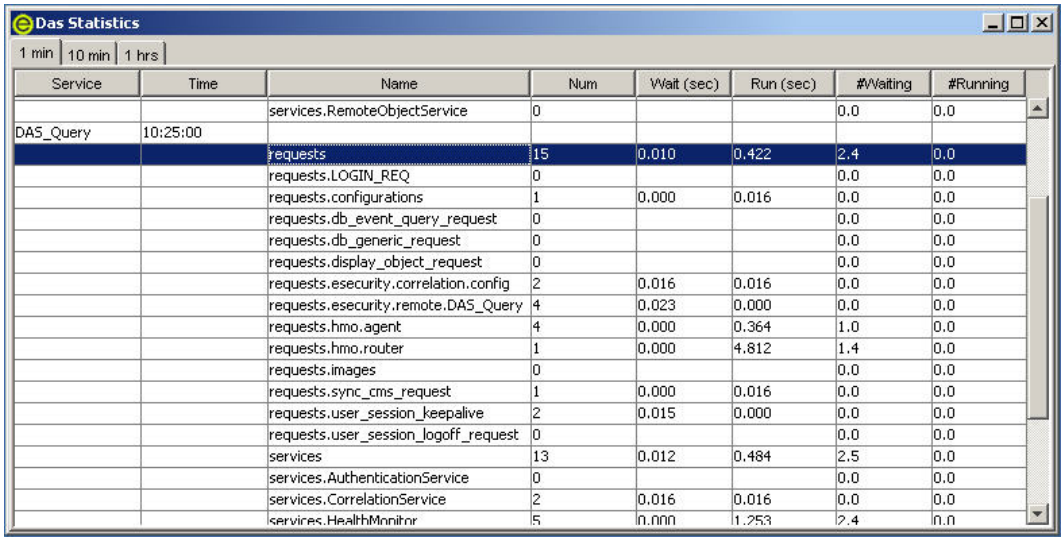


| Service | Time | Name | Num | Wait (sec) | Run (sec) | #Waiting | #Running |
|---|---|---|---|---|---|---|---|
| | | services.RemoteObjectService | 0 | | | 0.0 | 0.0 |
| DAS_Query | 10:25:00 | | | | | | |
| | | requests | 15 | 0.010 | 0.422 | 2.4 | 0.0 |
| | | requests.LOGIN_REQ | 0 | | | 0.0 | 0.0 |
| | | requests.configurations | 1 | 0.000 | 0.016 | 0.0 | 0.0 |
| | | requests.db_event_query_request | 0 | | | 0.0 | 0.0 |
| | | requests.db_generic_request | 0 | | | 0.0 | 0.0 |
| | | requests.display_object_request | 0 | | | 0.0 | 0.0 |
| | | requests.esecurity.correlation.config | 2 | 0.016 | 0.016 | 0.0 | 0.0 |
| | | requests.esecurity.remote.DAS_Query | 4 | 0.023 | 0.000 | 0.0 | 0.0 |
| | | requests.hmo.agent | 4 | 0.000 | 0.364 | 1.0 | 0.0 |
| | | requests.hmo.router | 1 | 0.000 | 4.812 | 1.4 | 0.0 |
| | | requests.images | 0 | | | 0.0 | 0.0 |
| | | requests.sync_cms_request | 1 | 0.000 | 0.016 | 0.0 | 0.0 |
| | | requests.user_session_keepalive | 2 | 0.015 | 0.000 | 0.0 | 0.0 |
| | | requests.user_session_logoff_request | 0 | | | 0.0 | 0.0 |
| | | services | 13 | 0.012 | 0.484 | 2.5 | 0.0 |
| | | services.AuthenticationService | 0 | | | 0.0 | 0.0 |
| | | services.CorrelationService | 2 | 0.016 | 0.016 | 0.0 | 0.0 |
| | | services.HealthMonitor | 5 | 0.000 | 1.253 | 2.4 | 0.0 |

The information can be useful because it shows what is going on. The number of requests is especially useful, you can see where they are all going or concentrated. The #waiting is useful because it shows how busy the server is. That number should be small. If it is large, new requests (even for simple tasks) will have to wait for potentially slow ones. This is not a good situation. The average run time is very important because it shows which requests are actually taking all the time, as opposed to waiting for others.

## Event File Information

The top pane shows the Status information for each event file. Status is of event files when the window was open. The pane will not show status of any past event file status. Gives file_id (which is the arch_id in the events table), file name and statistics about the file (if it is complete, start and end time of writing to the file, the minimum and maximum time of events contained in the file, etc).

When you hi-lite a file from the top pane, the bottom pane will show the summary status for that event file. The bottom pane displays the summary name, start and end time of processing, number of events processed and if there were any error messages.

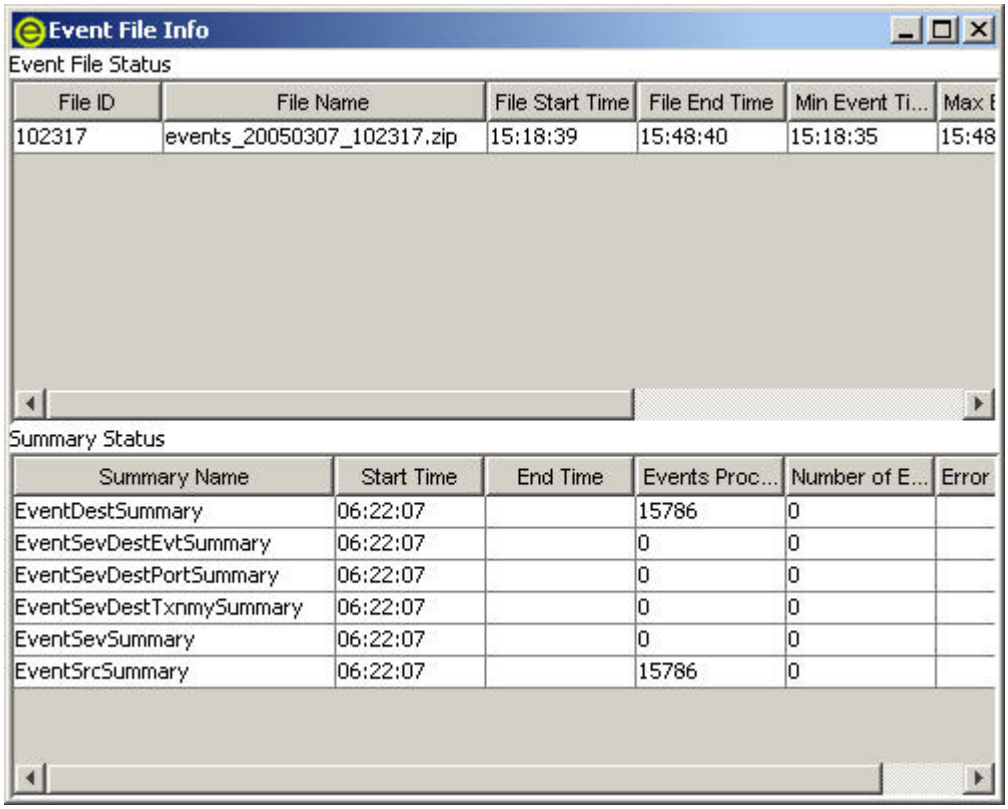


## User Configurations

To use this feature, you must have the user permission User Configuration in order to work in the User Configuration window.

User configuration allows you to:

- Create a User Account
- Modify a User Account
- View Details of a User Account
- Clone a User Account
- Delete a User Account
- Terminating an Active Session
- Add a iTRAC Role
- Delete iTRAC Role
- Details of an iTRAC Role

The installer will create the following default users on the Sentinel Server:

**Oracle and MS SQL Authentication**:

- esecdba - Schema owner (configurable at install time).
- esecadm - Sentinel administrator user (configurable at install time).

**NOTE**: For UNIX, the Installer also creates the operating system user with the same user name and password.

- esecrpt - Reporter user, password as the admin user.
- ESEC_CORR - Correlation Engine users, used to create incidents.
- esecapp - e-Security application username for connecting to the database.

**Windows Authentication**:

- e-Security DB Administrator - Schema owner (configurable at install time).

- e-Security Administrator - Sentinel administrator user (configurable at install time).
- e-Security Report User - Reporter user, password as the admin user.
- e-Security Application DB User - e-Security application username for connecting to the database

## Opening the User Manager Window

To open the User Manager window

1. Click the Admin tab.
2. Click Admin > User Configuration.

## Creating a User Account

**NOTE**: In order to meet stringent security configurations required by Common Criteria Certification, e-Security requires a strong password with the following characteristics:

1. Choose passwords of at least 8 with characters in length that includes at least one UPPER CASE, one lower case, one special symbol (!@#$%^&*()_+), and one numeric (0-9).

2. Your password may not contain your e-mail name or any part of your full name.

3. Your password should not be a "common" word (for example, it should not be a word in the dictionary or slang in common use).

4. Your password should not contain words from any language, because numerous password-cracking programs exist that can run through millions of possible word combinations in seconds.

5. You should choose a password you can remember and yet is complex. For example, Msi5!YOld (My Son is 5 years old) OR IhliCf5#yN (I have lived in California for 5 years now).

To use this feature, you must have the user permission Create User Account. User permissions are fairly detailed, see e-Security User's Reference Manual – User Permissions for information.

**NOTE**: The esecrpt user password must be changed directly in the database. Enterprise Manager can be used to do this.

To create a user account

1. Open the User Manager window.
2. Click the Add a new User button,


or high-light any user, right-click > Add User.

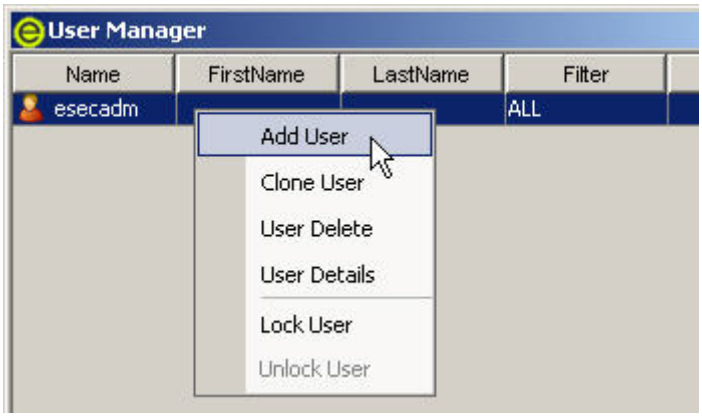


3. Under Authorization, enter:
   - User Name
   - Password
   - Confirm Password
   - Security Filter - To select a filter, click in the down arrow. The Filter Selection window opens. Highlight a filter or click the Add button to create a filter for this user account.

**NOTE**: After assigning a security filter to a user, you cannot delete that filter.

   - Click Select

**NOTE**: e-Security recommends as a best practice a minimum password length of 8 characters that includes alphanumerics.

   (Optional) Under Details, enter:
   - First Name
   - Last Name
   - Department
   - Phone
   - Email
4. Click the Permissions tab and assign user permissions.
5. Click the Roles tab and select the role for the user.
6. Click Ok.

**NOTE:** Oracle does not allow the creation of users named the same as one of the Oracle Reserved words. Also, e-Security does not allow you to use these names.

### Modifying a User Account

To use this feature, you must have the user permission Modify Existing User Account.

**NOTE**: The esecrpt user password must be changed directly in the database. Enterprise Manager can be used to do this.

To modify a user account

1. Open the User Manager window.
2. Double-click on a user account or right-click > User Details.
3. Modify the account.
4. Click Ok.

### Viewing Details of a User Account

To use this feature, you must have the user permission Use/View User Account.

To view user account details

1. Open the User Manager window.
2. Double-click on a user account or right-click > User Details.
3. Review the details of the user account and close the window.

### Cloning a User Account

To clone a user account

1. Open the User Manager window.
2. Select a user account ID, right-click > Clone User.
3. Change the user information and the user permissions.
4. Click Save.

### Deleting a User Account

To use this feature, you must have the user permission Delete User Account.

**NOTE**: When a user is deleted, that user cannot be created again. For instance if you create a user called Joe and later delete Joe, you will not be able to re-create a user called Joe.

To delete a user account

1. Open the User Manager window.
2. Select a user account ID, right-click > Delete User.

### Terminating an Active Session

Terminating an active session

1. Open the Active User Sessions window.
2. Highlight an active session you wish to terminate.
3. Right click > Kill Session.
4. You will be prompted for a termination message. This provided so that you can inform the user why you are killing the session.

### Adding a iTRAC Role

To add a iTRAC Role

1. Open the Role Manager window.
2. Click the Add a new Role button,


or right-click > Add New Role.

## Deleting a iTRAC Role

To delete a iTRAC Role

1. Open the Role Manager window.
2. Select a role, right-click > Delete Role.

## Viewing Details of a Role

To view role details

1. Open the Role Manager window.
2. Select a role, right-click > Role Details.

# Chapter 10 – Sentinel Data Manager

The Sentinel Data Manager (SDM) is a tool by which users can manage the Sentinel Database. The SDM allows users to perform the following operations:

- Monitor Database Space Utilization
- View and Manage Database Partitions
- Manage Database Archives
- Import Data into the Database
- Configure Data Mapping
- Configure Event Tag Names
- Configure Summary Report Settings

## Installing the SDM

The SDM can be installed directly from the Sentinel 5 InstallShield Wizard by selecting the “Sentinel Data Manager” component on the Sentinel 5 Feature Selection screen.

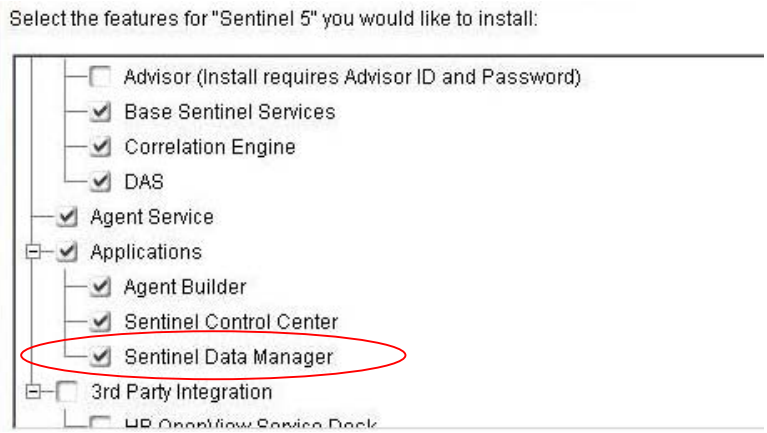


(Oracle only) Note that for the SDM to communicate with Oracle Databases, you must also manually download the Oracle 9.2.0.4 or 9.2.0.5 JDBC driver and copy the downloaded .jar file to the $ESEC_HOME/lib directory on the same box where you installed the SDM or %ESEC_HOME%\lib if installing the SDM on Windows. You can download the JDBC driver from the following URL:

http://otn.oracle.com/software/tech/java/sqlj_jdbc/index.html

The typical name of this jar file is ojdbc14.jar.

> **NOTE:** As of the date of publication of this guide, the above mentioned website was correct.

> **NOTE**: SDM for Oracle requires that Oracle Enterprise with partitioning be installed.

## Starting the SDM GUI

**NOTE**: In order to use the SDM GUI, your configuration.xml file must be pointing to a Communication Server that also has DAS_Binary and DAS_Query connected to it. This will normally be the case, by default, as long as the Communication Server and DAS processes are running.

For UNIX: starting SDM GUI

1. Login to the UNIX box as a member of the esec group (for example: esecadm).
2. cd to $ESEC_HOME/sdm
3. Enter the following command line:

```
./sdm
```

For Windows: starting SDM GUI

1. Click Start > Program Files > e-Security > Sentinel Data Manager

**NOTE**: To run the SDM from the command line, see the SDM Command Line section of this document.

## Connecting To Database

When the SDM starts up, you will need to establish a connection to your database. In the “Connect to Database” dialog, enter the appropriate values for each field.

Connecting to the Database

1. Start the SDM GUI.
2. Select your database type as Oracle or MSSQL.
3. Specify your Database instance name (Such as ESEC).
4. Specify your Database Host (Use the hostname or IP address).
5. For the port, use the default port of 1521 for Oracle or the default port of 1433 for MSSQL.
6. For the username and password, use your e-Security Database Administrator username and password. (Such as esecdba).

**NOTE**: For Windows and MS SQL, if you installed MS SQL in mixed mode, you can login using Windows Authentication OR SQL Server Authentication. If you installed MS SQL in Windows Authentication Only mode, you must login using Windows Authentication. If you choose to use Windows Authentication, you will be authenticated with the MS SQL database as the user you are currently logged into Windows as (i.e.- single sign-on).

For Oracle:

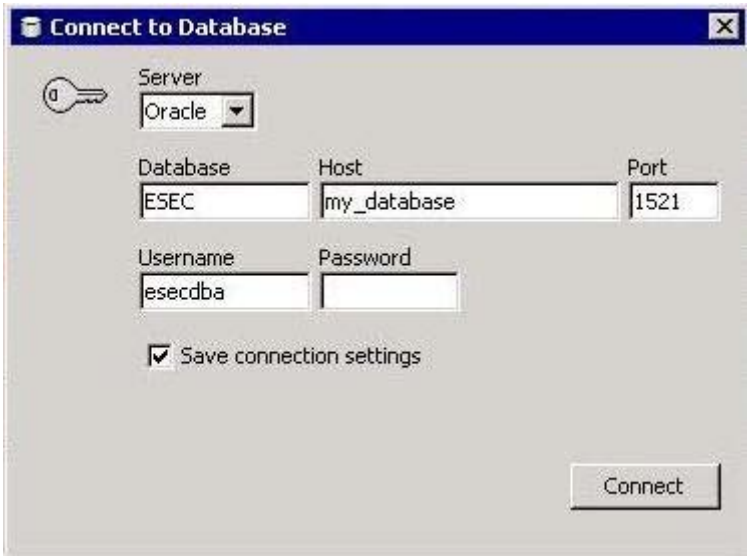


For Windows:

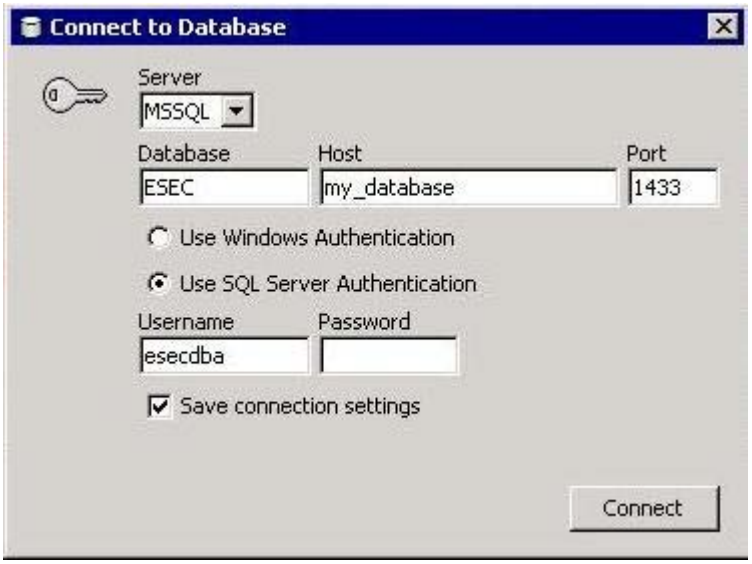


**NOTE:** If you select to save your connection settings, the settings are saved to the local sdm.connect file. Next time you start the GUI, the connection settings will be re-populated from the sdm.connect file. This file can be used when running SDM from the command line.

7. Click Connect.

## Partitions

The Partitions tab in the SDM allows users to view and manage database partitions.

To view partitions in the GUI

1. Click the Partitions tab.
2. Select the table in the dropdown list you would like to see.

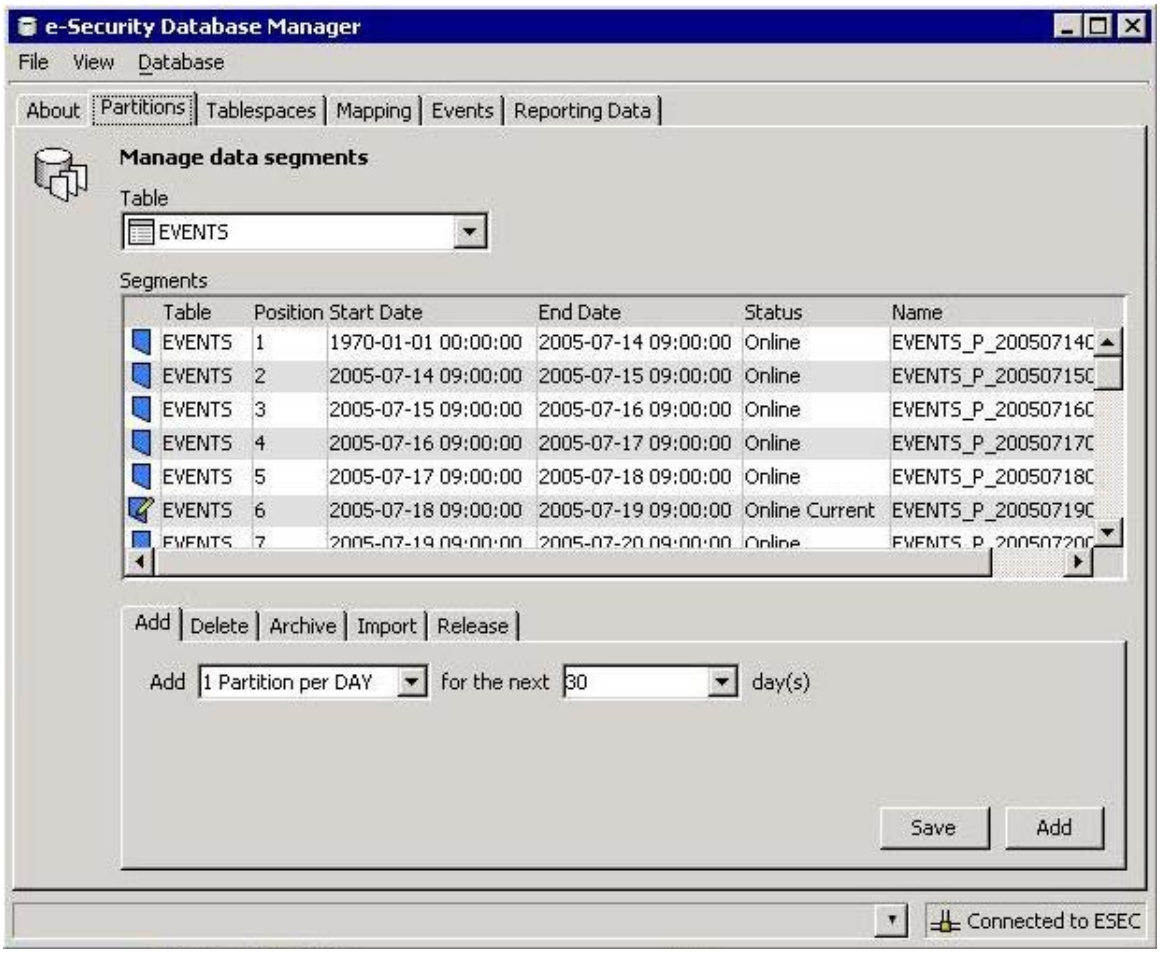


The Segments table displays the partitions of the currently selected Database Table.

Each row in the Segments table displays the related Database Table, Time Range, Status and Name of the partition.

**NOTE**: The first date in your Partitions tab will always have a Start Date of 1970-01-01 00:00:00.

| Table | Position | Start Date | End Date | Status | Name |
|---|---|---|---|---|---|
| EVENTS | 1 | 1970-01-01 00:00:00 | 2004-04-30 09:00:00 | Online | EVENTS_P_20040430090000 |
| EVENTS | 2 | 2004-04-30 09:00:00 | 2004-05-01 09:00:00 | Online | EVENTS_P_20040501090000 |

The Status of each of the partitions shown in the Segments table will have one of the following states:

| | |
|---|---|
| Online | data in an online partition is available for access |
| Online Current | an online partition where rows are currently getting inserted into |
| Online Archived | partition whose data is archived but the data is still accessible due to one of the following reasons:<br>▪ partition not yet dropped<br>▪ partition is imported back |
| Offline | data in an offline partition is not available for access because the partition is dropped and not imported |
| Offline Archived | partition that is archived and dropped |

### To manage partitions

1. Click the Partitions tab.
2. Select the table in the dropdown list.
3. Select the tab in the bottom of the window that relates to the operation that you would like to perform – Add, Delete, Archive, Import or Release.

### To add partitions

1. Select the “Add” partitions tab.
2. Specify the number of partitions to add and the number of days over which to add the partitions.
3. Press the “Add” button.

### To delete partitions

1. Select the “Delete” partitions tab.
2. Specify the number of days for which older partitions will be deleted.
3. Press the “Delete” button.

### To archive partitions

**NOTE**: Aggregation tables are not archived.

1. Select the “Archive” partitions tab.
2. Specify the number of days for which older partitions will be archived and the directory into which the archive will be stored.

**NOTE**: For UNIX, partitions cannot be archived to /root.

3. Press the “Archive” button.

**NOTE**: When archiving, be sure to enter a valid path on the database server with the correct permissions.

**NOTE**: The Archive tab is different for MSSQL and Oracle. For Oracle, Oracle allows you to specify the maximum size of your archive file.

Oracle Archive Partitions Tab:

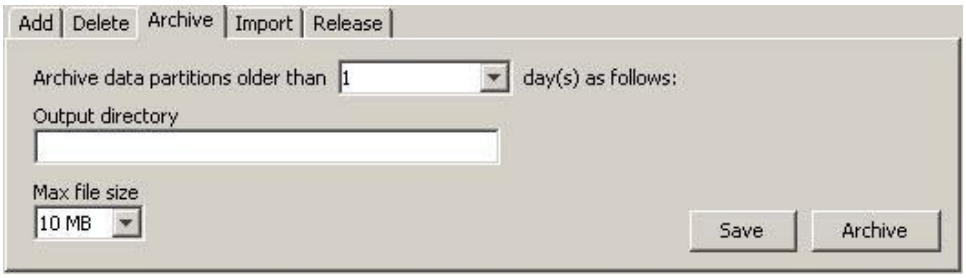


MSSQL Archive Partitions Tab:

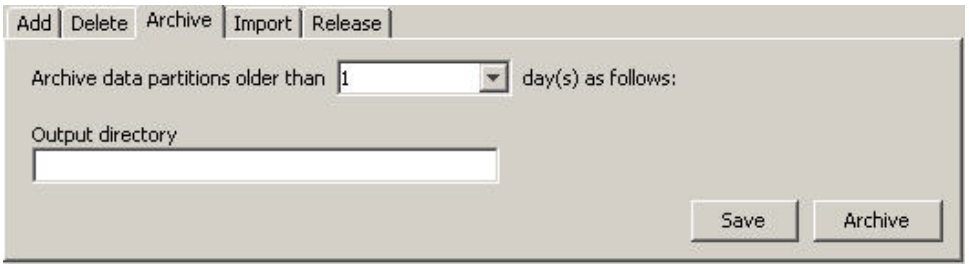


To import partitions

1. Select the “Import” partitions tab.
2. Select the partition in the Segment table into which the data will be imported.
3. Specify the input directory from which the archived data will be read.
4. Press the “Import” button.

To release imported partitions

1. Select the “Release” partitions tab.
2. Select the partition in the Segment table that will be released.
3. Press the “Release” button.

## Tablespaces

The Tablespaces tab in the SDM allows users to view the current database space utilization.

To view the tablespaces in the GUI

1. Click the Tablespaces tab.

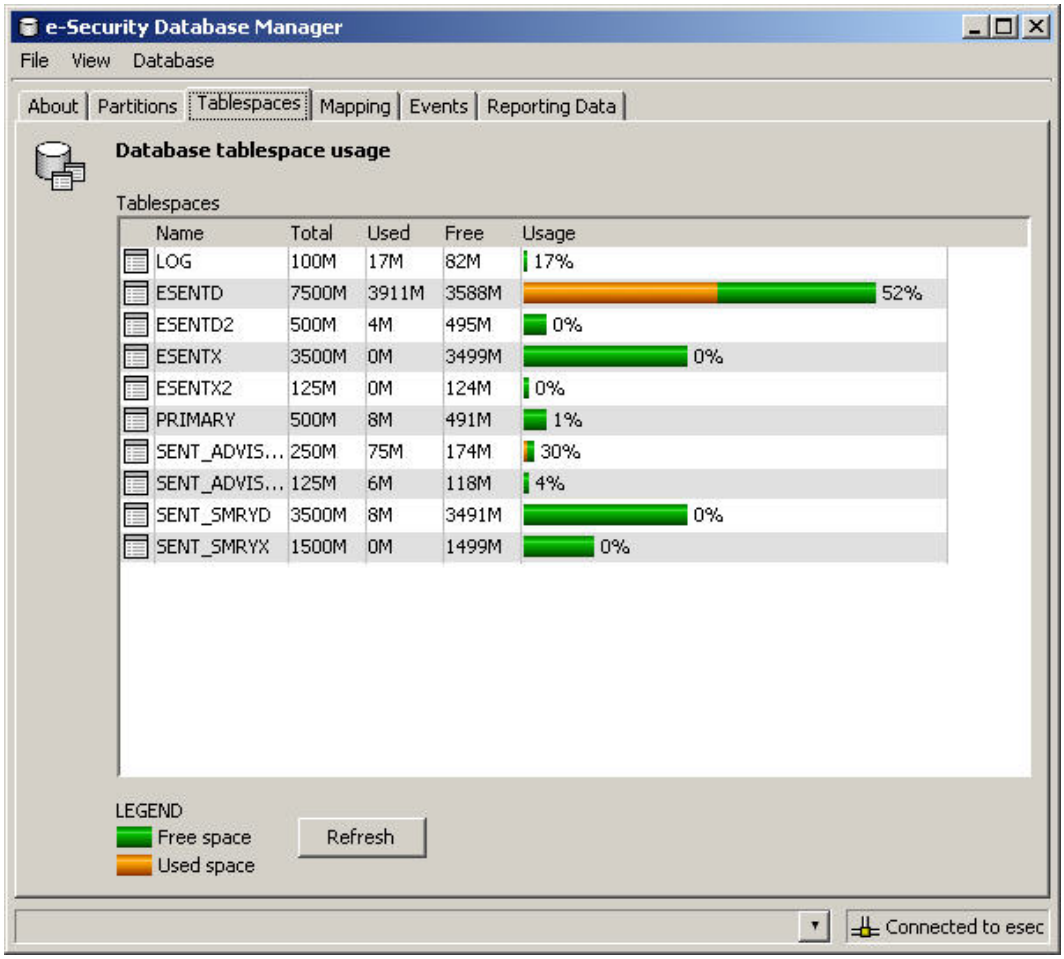


The Tablespace Usage table displays the total space allocated for each tablespace, how much memory has been used by each tablespace and how much memory is still available (free) for each tablespace. Color coded bar graphs help to visualize the total space allocated for each tablespace and the percent used of each tablespace.

> **NOTE**: On MS SQL, there is no such thing as tablespaces therefore filegroups are used.

## Mapping Tab

> **NOTE**: In order to use the Mapping Tab, your configuration.xml file must be pointing to a Communication Server that also has DAS_Binary and DAS_Query connected to it. This will normally be the case, by default, as long as your Communication Server and DAS processes are running.

The Mapping tab allows you to:

- Add new map definitions
- Edit map definitions
- Delete map definitions
- Update map data

Mapping works together with the "Referenced from Map" Data Source option under the Events tab.

To view maps in the GUI

1. Click the Mapping tab.

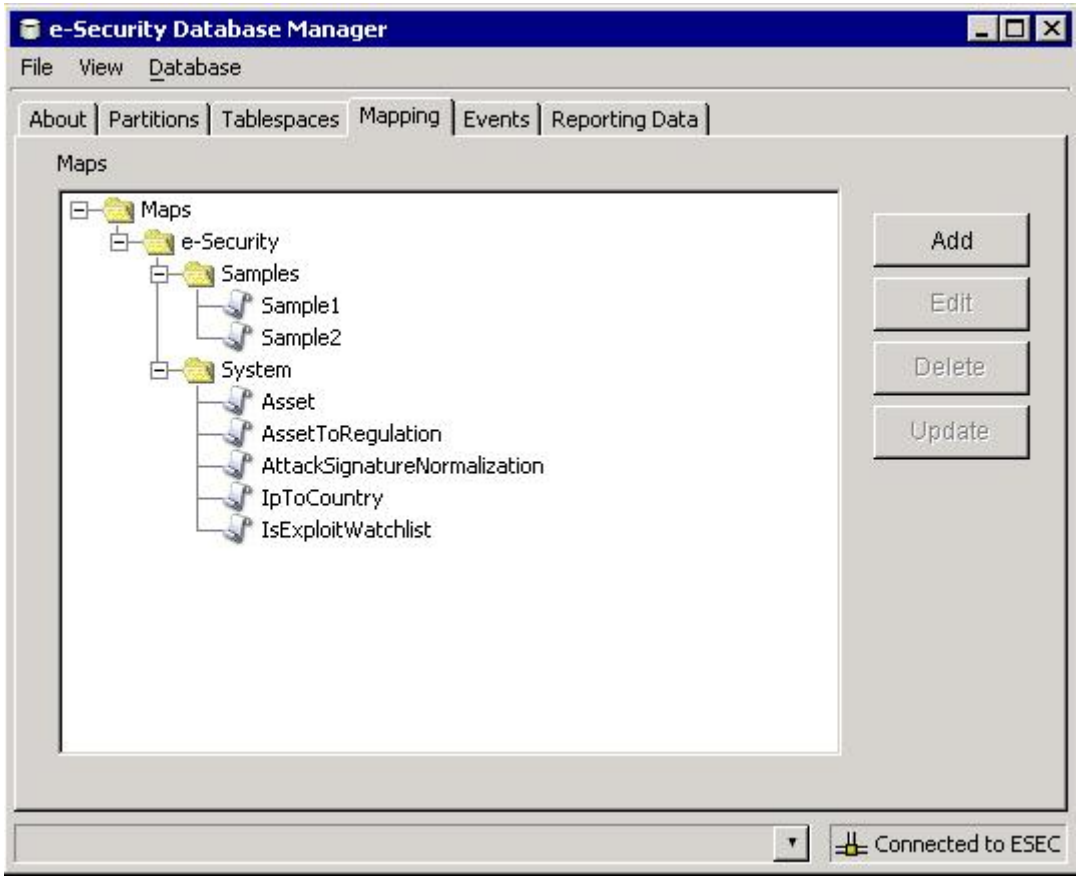


The main Mapping GUI displays a listing of all of the maps that have been defined for the system.

**NOTE**: Maps under the System folder cannot be edited or deleted.

**Adding Map Definitions**

To add a map definition:

1. Click the Mapping tab.
2. Click the Add button.
3. If you are creating a new map folder, click the New... button. Enter a folder name.

   **NOTE**: If this is your first map definition, it is recommended that you create a new map definition folder. Creating a map definition under the System folder will not allow you to edit or delete your map definition.
4. Ensure that the folder you want to enter your map definition into is selected. (i.e. The folder indicates that it is open).
5. Enter your Map Name.

6. Click Next.

**NOTE**: The Map Type field box is disabled.

7. Select either Local File or Remote File.
   - Local File – allows you to browse for your file on your local file system (on the machine were SDM was launched from).
   - Remote File – allows you to choose from existing map source data files on the server where DAS is running. Two files that may already exist on the server (if Advisor is installed and Vulnerability data was uploaded) are attackNormalization.csv and exploitDetection.csv. Remote file points to %ESEC_HOME%\sentinel\bin\map_data (Windows) or $ESEC_HOME/sentinel/bin/map_data (UNIX)

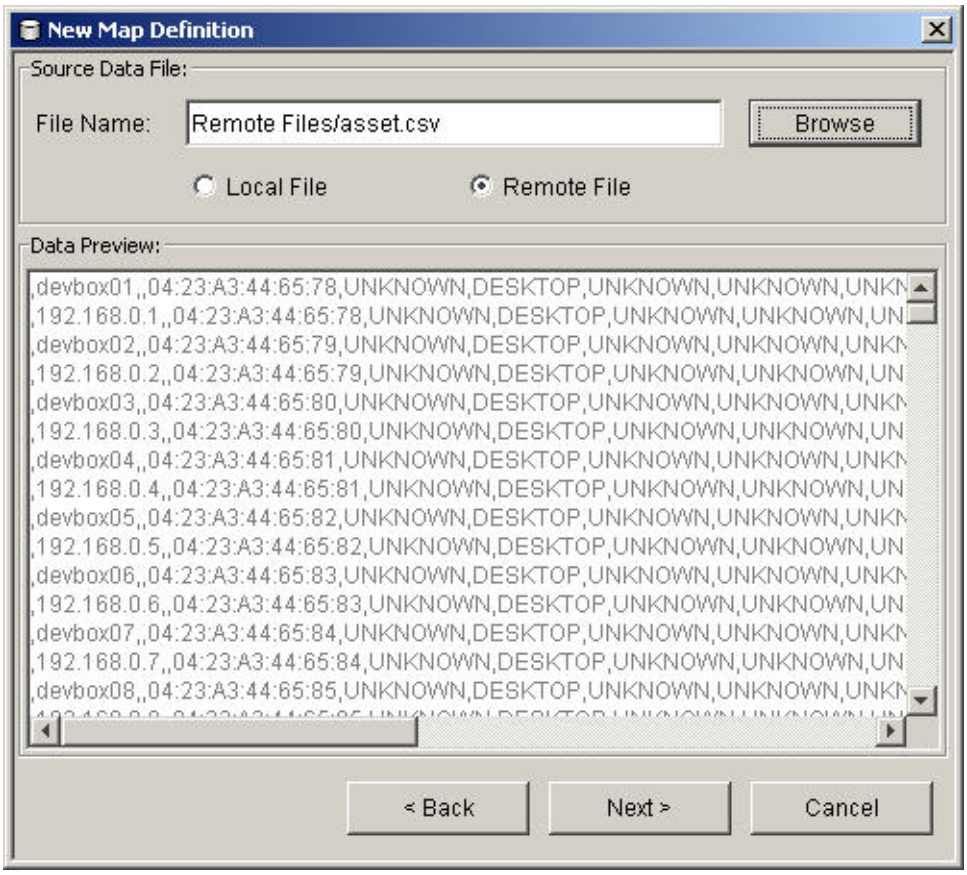


Select your map definition file. Click Next.

**NOTE**: For map files that contain more than 500 lines, you will not see all the lines in the SDM.

8. In the New Map Definition window, set the following:
   - Delimiter (pipe, comma, semicolon, etc...) of data in rows of the map data source file
   - Start at row – The number of rows to skip from the top of the map data source file.
   - Column names
   - Column types (String or NumberRange)

- Active columns – When a column is marked as active, the data in the column will be distributed to processes using maps. All key columns must be active. Only non-key columns that are active can be select as the “Map Column” under the Events tab.
- Key columns - A key is a unique identifier for the row of data in the map data. If more than one column is selected as a key, the overall key of the map will include all of the columns selected as keys.
- Column filtering - A row can be explicitly included or excluded based on matching criteria for a particular column. This can be used to exclude rows from the map source data that are not needed or will interfere with your mapping.

As you configure each setting and filter, the data table will automatically update to allow you to preview your data and ensure your data is being parsed as expected.

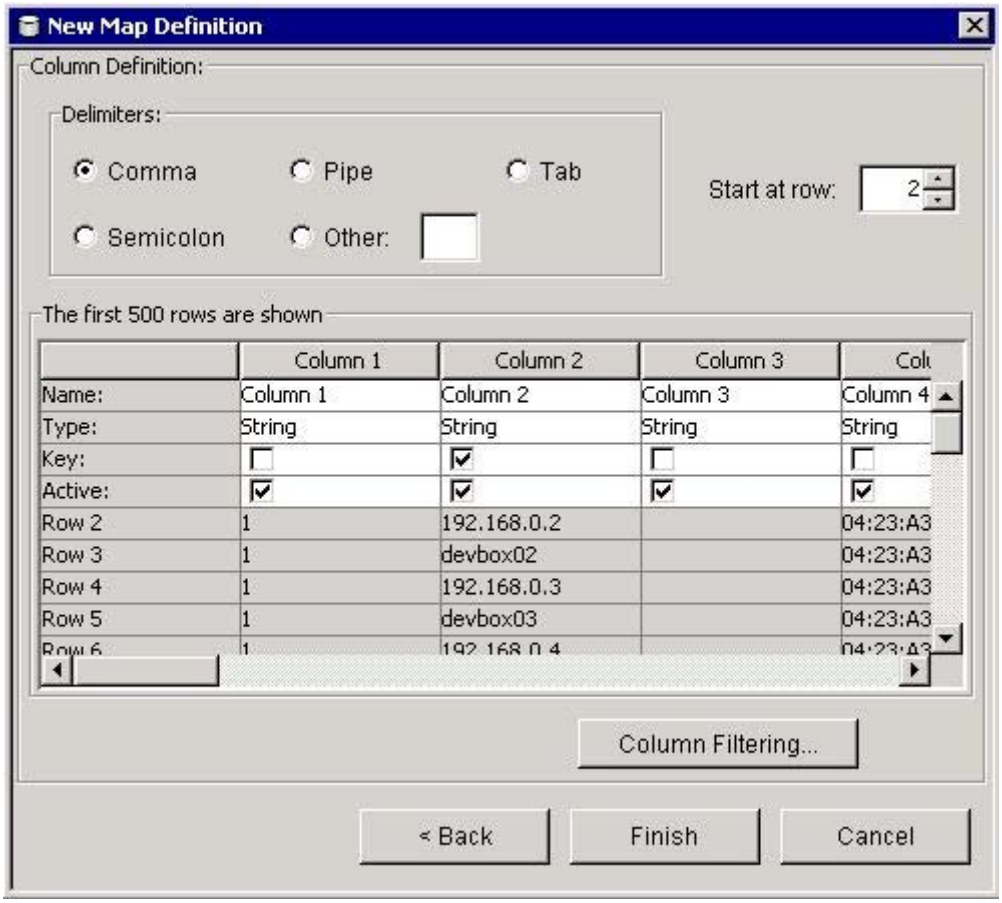


9. Once you finish configuring all parameters and filters for the definition, click Finish.
10. If you chose Local File in step 7 above, you will be prompted to upload your file to the Remote Files virtual folder located: %ESEC_HOME%\sentinel\bin\map_data. Enter a file name and click OK.

### Editing Map Definitions

To edit a map definition:

1. Click the Mapping tab.
2. Expand the folder of interest.
3. Highlight a map definition and click the Edit button.

**NOTE**: The editing function is disabled for map definitions that are under the Systems folder.

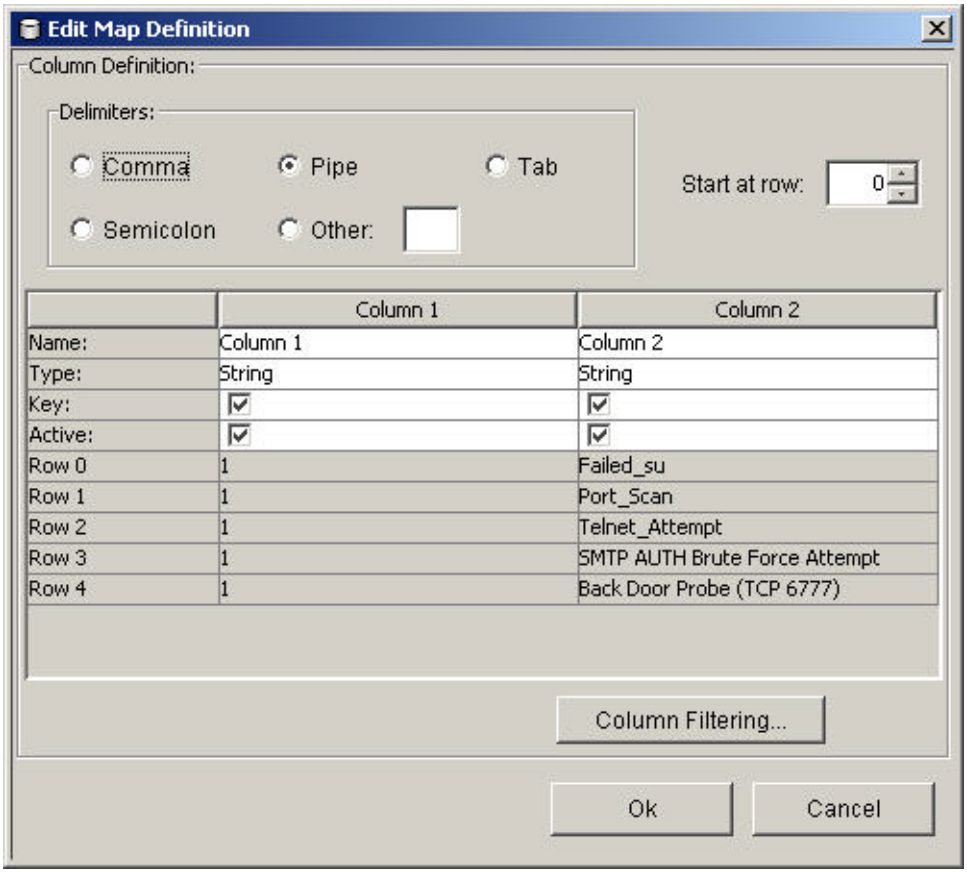


The edit function allows you to:

- set your delimiters
- set which row to start your map
- rename your columns
- activate or deactivate a column
- set your column keys
- column filter

4. After making your changes, Click Ok.

### Deleting Map Definitions

To delete a map definition

1. Click the Mapping tab.
2. Expand the folder of interest.
3. Highlight the map definition to be deleted.
4. Click the Delete button.

**NOTE**: Map definitions under the e-Security folder cannot be deleted.

### Updating Map Data

Updating allows you to replace the map source data file of a map on the server running DAS with another file. Your new map source data file must have the same delimiter, number of columns, and overall structure as the existing map data source file in order for the map to function properly after the update. The new map source data file should only differ from the existing file by the values that appear in the columns. If the new map source data file has a different structure than the existing file, use the Edit feature of the SDM GUI to update the map definition.

To update map data

1. If you haven't already, create a file containing the new map source data on the machine where you run SDM. This file can be generated (e.g.- from a data dump script), created manually from scratch, or be an edited version of the existing map data source file. If needed, you can obtain the existing map data source file from the location:

   For Windows:

   ```
   %ESEC_HOME%\sentinel\bin\map_data
   ```

   For UNIX

   ```
   $ESEC_HOME/sentinel/bin/map_data
   ```

2. Click the Mapping tab.
3. Expand the folder of interest. Highlight the mapping to be updated. Click the Update button.

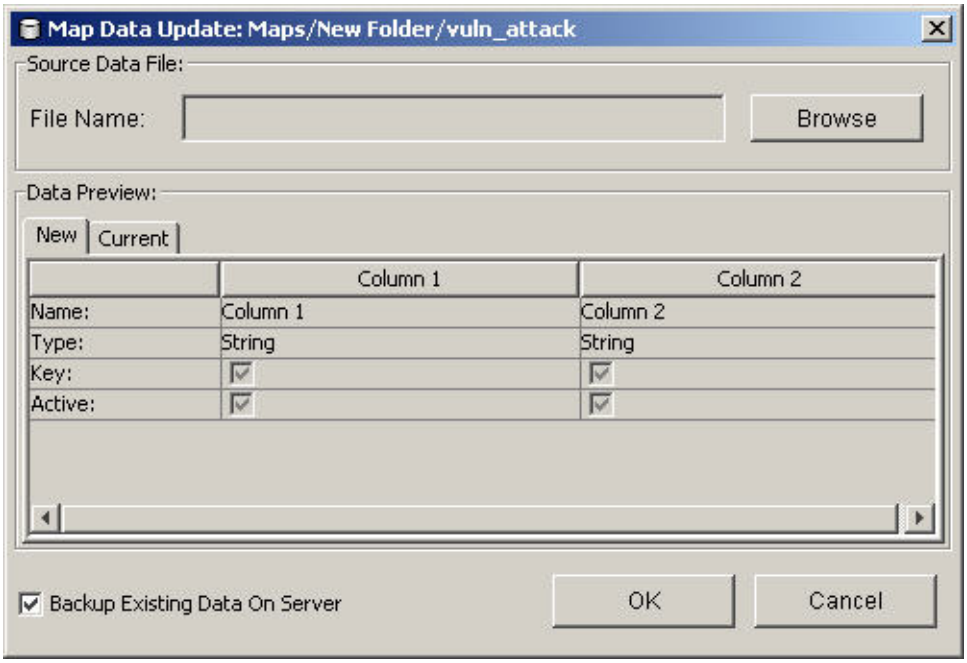


4. Select the new map data source file by clicking the Browse button and selecting the file with the new map data. After selecting the file, the data from the new map data source file will appear under the New tab. The map data you are replacing will be under the Current tab.

5. Uncheck or leave the default setting for 'Backup Existing Data On Server'. Enabling this option results in a backup of the existing map data source file being put in the %ESEC_HOME%\sentinel\bin\map_data (Windows) or $ESEC_HOME/sentinel/bin/map_data (UNIX) folder. The prefix of the name of the backup map data source file will be the name of the existing map data source file. The end of the filename will contain a set of random numbers followed by the .bak suffix. For example: vuln_attacks10197.bak.
6. Click Ok.
7. The data from the new map data source file will be uploaded to the server, replacing the contents of the existing map data source file. After the source data is completely uploaded, the map data will be regenerated and distributed to map clients (e.g. - Agent Manager).

## Events Tab

> **NOTE**: In order to use the Events Tab, your configuration.xml file must be pointing to a Communication Server that also has DAS_Binary and DAS_Query connected to it. This will normally be the case, by default, as long as your Communication Server and DAS processes are running.

### Event Mapping

Event Mapping is a mechanism that allows you to add data to an event by using data already in the event to reference and pull in data from an outside source. The outside data source is a map, which is defined using the Mapping Tab. The data already in the event that should be used as the reference into the map and the data to be pulled from the map into the event are specified using the Events Tab.

Since virtually any data set can be made into a map, Event Mapping is useful for incorporating into the event stream data from elsewhere in your organization. Some opportunities Event Mapping provides are:

- Regulatory Compliance monitoring
- Policy compliance
- Response prioritization
- Enable security data to be analyzed related to business operations
- Enhance accountability

When an Event Mapping is defined, it is applied system-wide to all events from all Agents. Additionally, Sentinel will automatically distribute map data to all processes that perform event mappings as well as keep the map data in these processes up-to-date. For these reasons, Event Mapping provides significant capabilities to support enterprise deployments.

Event Mapping comprises of four main parts:

- Controller - stores all map information
- Distributor - Automatically redistributes modified maps to those processes that registered for the map
- Monitor - a monitor to detect changes in map source data
- Generator - Generates maps from source data

One application of Event Mapping is e-Security's Asset Data functionality. For example, asset information is collected and stored in the Sentinel Database asset schema and is represented by a Physical Asset Entry. Soft assets, such as services and applications, are represented by an entry that is linked to a Physical Asset. The primary automated update mechanism for asset data is through an asset agent reading data from a scanner such as Nmap. The asset agent automates the retrieval of asset information by reading asset data from the scanner and populating the asset schema tables with this data. For Event Mapping, asset information is mapped from the destination IP and source IP.

There are two types of data sources:

- External – An Agent populates that value in the event tag.
- Referenced from Map - Data is retrieved from a map to populate the tag.

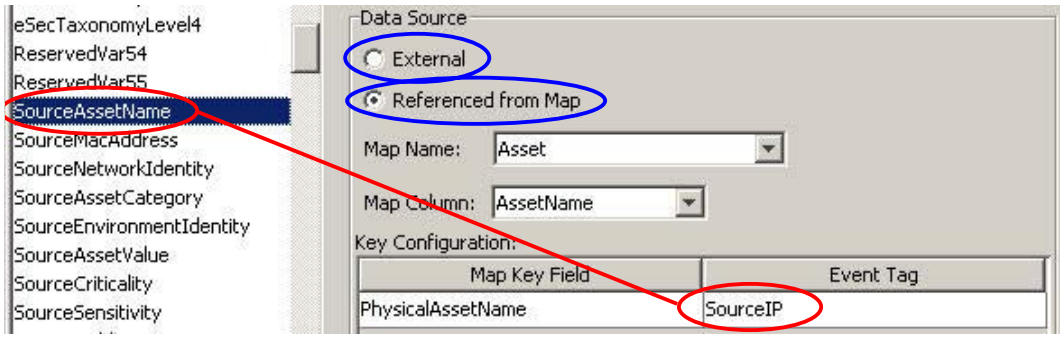


In the above illustration, the SourceAssetName tag is populated from the map called Asset (which has asset.csv as its map data source file). The specific value for SourceAssetName is taken from the AssetName column from the Asset map. The PhysicalAsssetName column is set as the key. When the SourceIP tag of the event matches one of the source IP values in the PhysicalAsssetName column of the map, the row with the matching key is used to intersect the AssetName Column. For instance, in the below example IP 198.168.1.100 corresponds to AssetName Finance35.

> **NOTE**: When a column is set as a key, it will not appear in the Column drop down field.

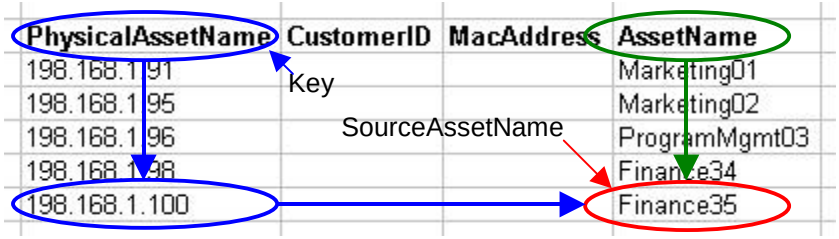


You may have more than one column set as a key as you do not want the map to be a Range Map (Range Maps can only have one key column, with that column type set to NumberRange). For instance (with column type set to String) the AttackId tag has the DeviceName (name of the security device) and DeviceAttackName columns set as keys and uses the NormalizedAttackID column in the AttackNormalization map for its value. In a row where the DeviceName event tag matches the data in Device map column and the DeviceAttackName matches the data in the AttackSignature map column, the

value for AttackId is the value in the NormalizedAttackID column. The configuration for Event Mapping just described is:

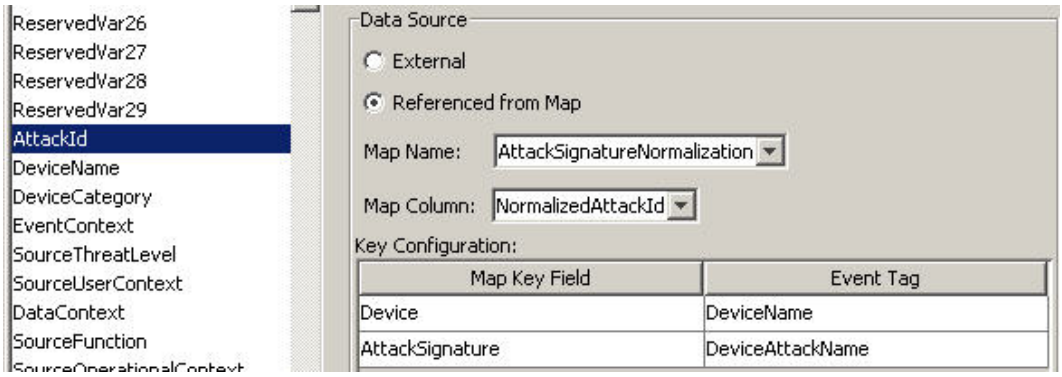


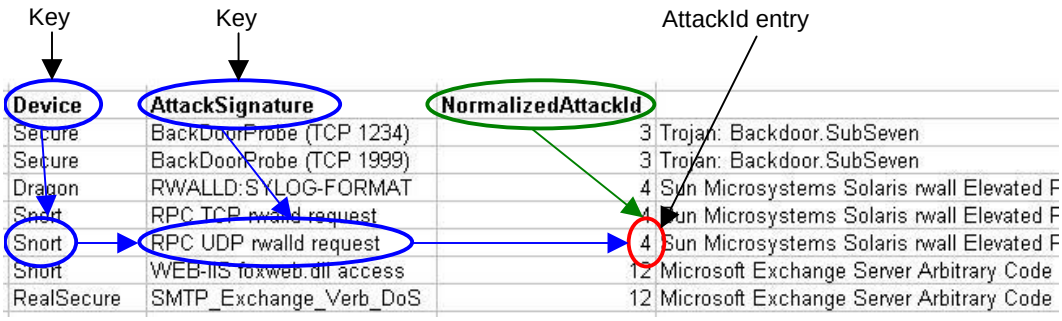


A numeric range can be set as a key, such as 10-200.

> **NOTE**: To use the range map functionality, a map definition must have exactly one key column and the key column must be of type NumberRange. If there are any other key columns, or the key column is of a different type, the mapping service will not consider the map a range map.

To create a range map, select a single column to be the key of the map and select 'NumberRange' as the type of the column. The format of the data in a column of type NumberRange must be 'm-n', where m is the minimum number in the range and n is the maximum number in the range (i.e., 10-200). The maximum number in the range is not included in the range (i.e- [m,n)). This means a range of 10-200 will only key off numbers equal to 10 to 199. An example set of data is with the first column as the key:

```
1-2,AA
2-4,AA
4-12,BB
10-20,BB
30-31,BB
100-200,AA
110-120,CC
```

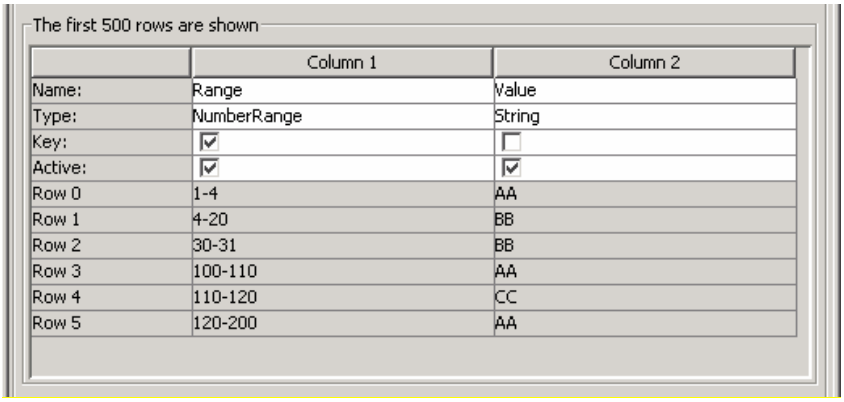
The first 500 rows are shown

| | Column 1 | Column 2 |
|---|---|---|
| Name: | Range | Value |
| Type: | NumberRange | String |
| Key: | ☑ | ☐ |
| Active: | ☑ | ☑ |
| Row 0 | 1-4 | AA |
| Row 1 | 4-20 | BB |
| Row 2 | 30-31 | BB |
| Row 3 | 100-110 | AA |
| Row 4 | 110-120 | CC |
| Row 5 | 120-200 | AA |

Notice how the example table gets transformed.

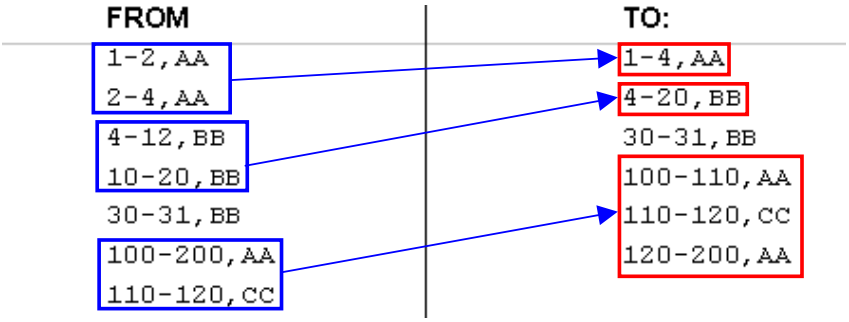


An example event configuration on the above map may look like:

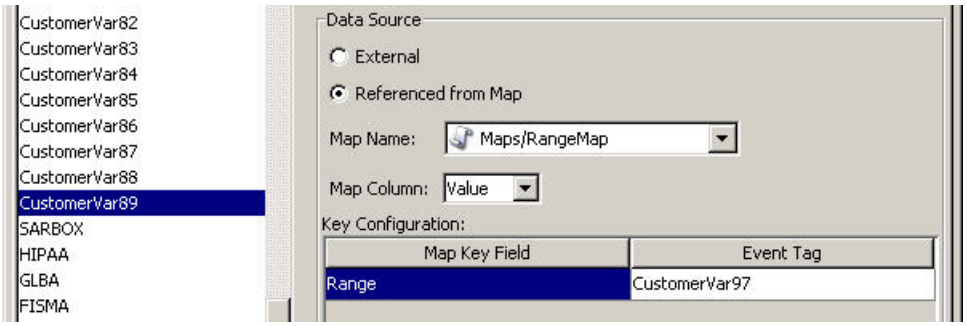


Where CustomerVar97 is expected to contain a numeric value (or is of a type that can be converted to a numeric value, such as an IP or Date).

When performing lookups into the example range map, the value in CustomerVar97 will take the range map and search for the range that the value belongs in (if any). Some examples and their results are:

```
CustomerVar97 = 1; CustomerVar89 will be set to AA

CustomerVar97 = 4; CustomerVar89 will be set to BB

CustomerVar97 = 300; CustomerVar89 will not be set
```

Internally, Sentinel converts IP addresses and dates to an integer for tags of the type IPv4 and date.

IPv4 tags are:

- DestinationIP (dip)
- SourceIP (sip)

Date tags are:

- CustomerVar11 to CustomerVar20 (cv11 to cv20)
- DateTime (dt)
- ReservedVar11 to ReservedVar20 (rv11 to rv20)

For more information on meta-tags, see Sentinel Reference Guide, Chapter 5 – Wizard and Sentinel Meta-tags.

For example, for the table below, column 1 is numerical range equivalent to an IP range of 10.0.0.0 to 10.0.2.255.

```
167772160-167772415,AAA
167772416-167772671,BBB
167772672-167772927,CCC
```

Using the same setup as the previous example, if:

- the Event Tag is set to DestinationIP and key column set to column 1 (range)
- Map Column to column 2 (value). The output values for CustomerVar89.

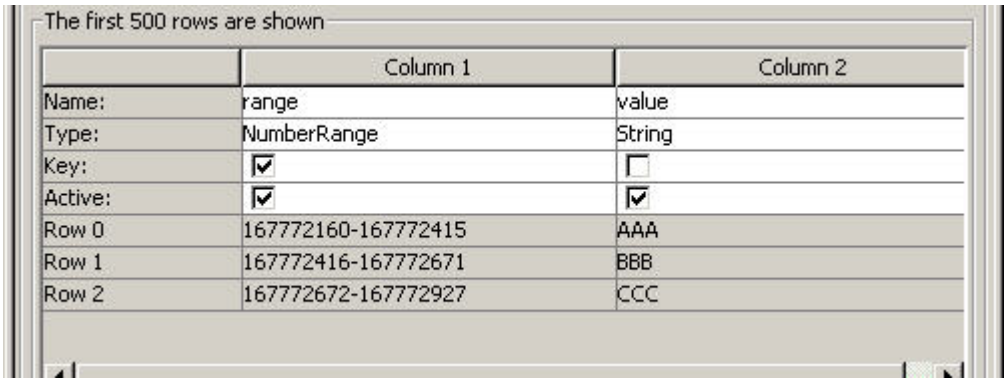


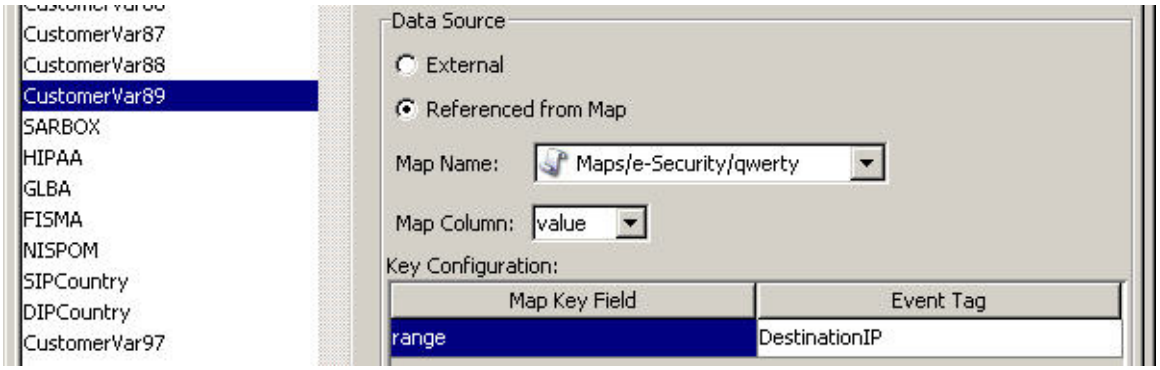


If an event contains a destination IP of 10.0.1.14 (equivalent to numerical value of 167772430), the output for column CustomerVar89 within the event would be BBB.

Additionally, Sentinel Linux supports the following number ranges:

- Range from negative number to negative number (e.g., "-234--34")
- Range from negative number to positive number (e.g., "-234-34")
- Range from positive number to positive number (e.g., "234-236")

- Single number range (negative) (e.g., "-234"). In this case, the min and the max will both be -234.
- Single number range (positive) (e.g., "234"). In this case, the min and the max will both be 234.
- Range from negative number to max number (e.g., "-234-"). In this case, the min will be -234 and the max will be (2^63 - 1).
- Range from positive number to max number (e.g., "234-"). In this case, the min will be 234 and the max will be (2^63 - 1).

**NOTE**: In all cases, the min must be less than or equal to the max (e.g., "-234--235" is NOT valid).

### Configuring Event tags (columns) to use Mapping

1. Click the Events tab.
2. Highlight an event tag entry from the Event Columns list.

**NOTE**: The original Event Tag name appears above the Label field. In addition, the description of the event column is provided.

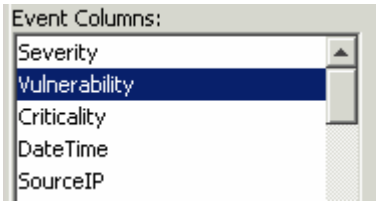


3. Click “Referenced from Map” to configure the event tag to be populated with data from a map. Click External to keep whatever value the Agent put in the event tag (if any).
4. Click the Map Name field down arrow.

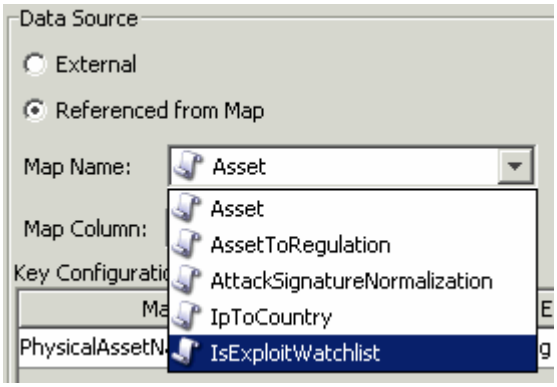


Select one of the following default maps or a map you have created:

- Asset – Contains the data from the map data source file asset.csv. The asset.csv is automatically generated from asset data from Sentinel Database when an asset Agent is run. This file could be populated manually instead, if desired.
- AssetToRegulation – Contains the data from the map data source file AssetToRegulation.csv. This file must be populated manually.
- AttackSignatureNormalization – Contains the data from the map data source file attackNormalization.csv (IDS signatures). The

attackNormalization.csv file is automatically generated from Advisor data from Sentinel Database when an Advisor feed is completed.

- IpToCountry – Contains the data from the map data source file IpToCountry.csv. This file must be populated manually.
- IsExploitWatchlist – Contains the data from the map data source file exploitDetection.csv (vulnerabilities and threats). The exploitDetection.csv file is automatically generated from Advisor and Vulnerability data from Sentinel Database when either an Advisor feed is completed or a vulnerability Agent is run.

5. Click the Map Column field down arrow and select a Map Column name. Depending on your Map Name choice in the previous step, these values will vary.

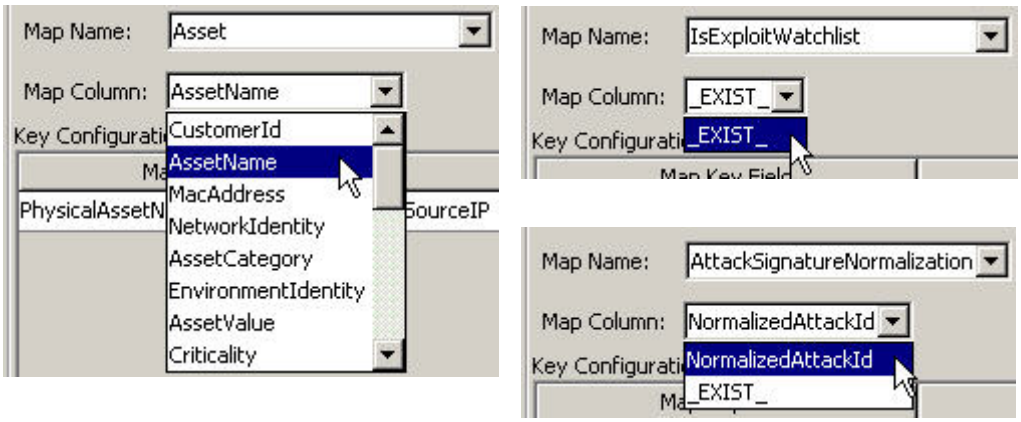


- _EXIST_ - This is a special Map Column that exists in every map. If this Map Column is selected, a “1” will be put in the event tag if the key is in the map data. If the key is not in the map data, a “0” will be put in the event tag.
- All other choices – names of active columns within the map definition that are not set as a key (e.g. - CustomerId column in Asset or NormalizedAttackId column in AttackNormalization)

6. In the Key Configuration, for each row in the table select the event tag in the Event Tag column that will be matched against the map key column specified in the corresponding Map Key Field column. The rows in the Key Configuration table will depend on the Map Name selected.

**NOTE**: A key is a unique identifier for the row of data in the map data.

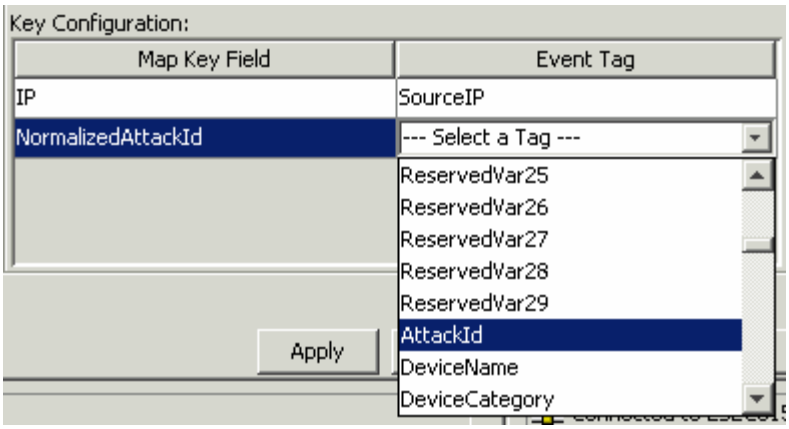


7. Click Apply.

> **NOTE:** Clicking Apply saves the changes you made for the currently selected event column in a temporary buffer. If you don't click Apply, when you select a different event column the changes you made to the previously selected event column are lost. Changes won't be saved to the server until you click Save.

8. If you would like to edit the Event Mapping of another Event Column, repeat the steps above. Remember to click Apply after editing the Event Mapping of each Event Column.
9. Click Save.

> **NOTE**: Clicking Save will save your changes to the server. The save function saves all changes stored in the temporary buffer (when you clicked Apply).

### Renaming Tags

The Events tab also allows you to assign names to existing event tag labels. For example, you can rename the label for event tag Ct2 to City. Doing this will result in the event tag that formally appeared in Sentinel Control Center as “Ct2” to now appear as “City”. Some places where event tags appear in Sentinel Control Center are filters, correlation rules, and Active Views.

Renaming Tags does not change the name of the variable in Agent scripts, however. Therefore, even if the event tag labeled Ct2 is renamed to City, the variable that must be used in an Agent script to reference this meta-tag will still be s_CT2.

Below is a before and after illustration of this feature in an Active View.


Renaming an event column

1. Click the Events tab.

   **NOTE**: The original Event Column name appears above the Label field. In addition, the description of the event column is provided.

2. Highlight an event column entry.
3. Enter a new value for your Event Column in the Label field.

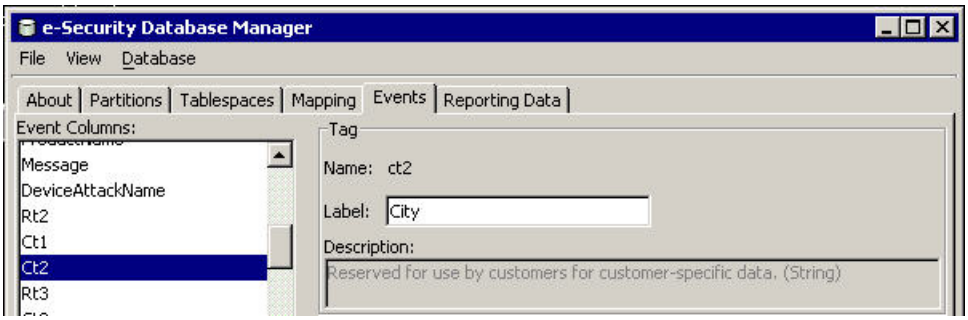


4. Click Apply.

   **NOTE**: Clicking on Apply saves the changes you made for the currently selected event tag in a temporary buffer. If you don't click apply, when you select a different event tag, the changes you made to the previously selected event tag are lost. Changes won’t be saved to the server until you click Save.

5. Click Save.

> **NOTE**: Clicking Save will save your changes to the server. The save function saves all changes stored in the temporary buffer (when you clicked Apply).

6. In order for changes to be visible in Sentinel Control Center, running Sentinel Control Centers must be closed and reopened.

## Reporting Data Tab

> **NOTE**: In order to use the Reporting Data Tab, your configuration.xml file must be pointing to a Communication Server that also has DAS_Binary and DAS_Query connected to it. This will normally be the case, by default, as long as the Communication Server and DAS processes are running.

The Reporting Data tab is a Summary Management Interface for Sentinel. This tab allows you to enable and disable Summaries. Enabling a summary allows Aggregation to start computing the counts for that particular Summary.

A Summary is a defined set of attributes that make up the key for which to compute the number of unique occurrences (event count) by each hour time period (event time). In the case of the EventSevDestPortSummary, when “active”, it saves the count of events for each unique combination of destination port and severity for an hour time frame. These saved computations of the event data allow for quicker summary reporting and querying. These reports are used by Crystal Reports. See the Crystal Reports Install chapters in the Sentinel Install Guide for more information. Certain Summaries will need to be “active” in order for the Summary reports to be accurate.

Aggregation is the process of calculating the running count for all active Summaries as events flow through the system. These running counts are saved to the database in the respective Summary tables.

Summaries Benefits:

- Greatly reduced event data set
- Conformed dimensions that allow the ability to drill-down, roll-up and drill-across on event data
- Summary reports run much faster with pre-computed summaries

Aggregation Benefits:

- Only processes active Summaries
- Does not affect event insertion into the real time database.

Reporting Data tab allows you to:

- enable/disable any predefined Summaries
- view attributes of each summary
- see the validity of a Summary for a timeframe
- query which eventfiles need to be run so that the Summary is complete

The following are all Summaries already defined in the system. It lists the Summary Name, database table name and it's attributes in a brief description about the Summary.

| Summary Name | Table/Description |
|---|---|
| EventSrcSummary | EVT_SRC_SMRY_1<br>This Summary sums the event count by source ip, source asset information, source port, source user, taxonomy, event_name, resource, agent, protocol, severity and event time by hour |
| EventDestSummary | EVT_DEST_SMRY_1<br>This Summary sums the event count by destination ip, destination asset information, destination port, destination user, taxonomy, event_name, resource, agent, protocol, severity and event time by hour. |
| EventSevDestTxnmySummary | EVT_ DEST_TXNMY_SMRY_1<br>This Summary sums the event count by destination ip, destion asset information, taxonomy, severity and event time by hour. |
| EventSevDestEvtSummary | EVT_DEST_EVT_NAME_SMRY_1<br>This Summary sums the event count by destination ip, destination event asset, taxonomy, event name, severity and event time by hour. |
| EventSevDestPortSummary | EVT_PORT_SMRY_1<br>This Summary sums the event count by destination port, severity and event time by hour. |
| EventSevSummary | EVT_SEV_SMRY_1<br>This Summary sums the event count by severity and event time by hour. |

## Disabling/Enabling Summary

1. Click the Reporting Data tab.
2. To disable a Summary, click the "Active" button in the Status column until it changes to say "InActive".
3. To enable a Summary, click the "InActive" button in the Status column until it changes to say "Active".

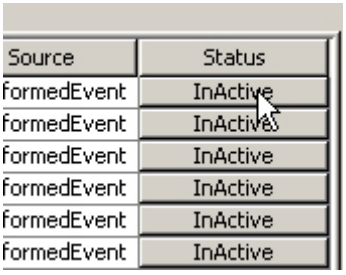


To enable Aggregation, only three summaries need to be enabled. They are:

- EventDestSummary
- EventSevSummary
- EventSrcSummary

## View information for a Summary

1. Click the Reporting Data tab.
2. Click on the "…" button in the Attributes column to see the attributes that makes up a Summary.

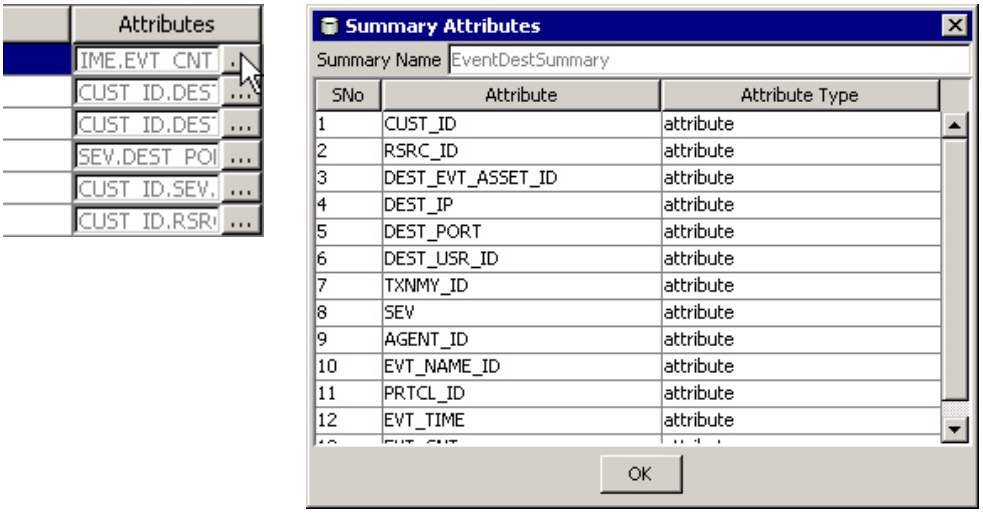


## Check the Validity of a Summary

1. Click the Reporting Data Tab.
2. Select the Status button.
3. Choose the summary or summaries you wish to query.

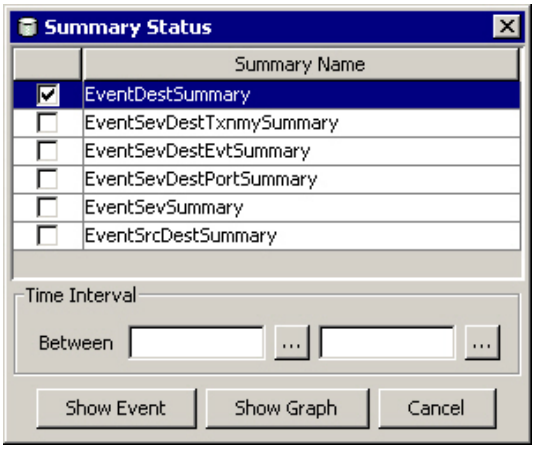


4. Select a time interval.
5. Click the 'Show Graph' button.
6. The green bars signify that the Summary is complete for that time frame. The red sections signify that the Summary is missing data during that time period.

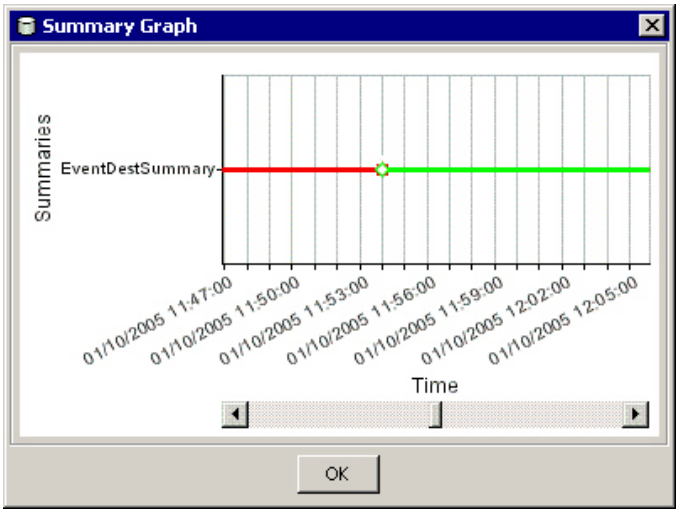


**NOTE**: To complete summaries, see “Run EventFiles for a Summary” section.

## Query the Eventfiles for a Summary

1. Click the Reporting Data Tab.
2. Select the Status button.
3. Choose the Summary or Summaries you wish to query.

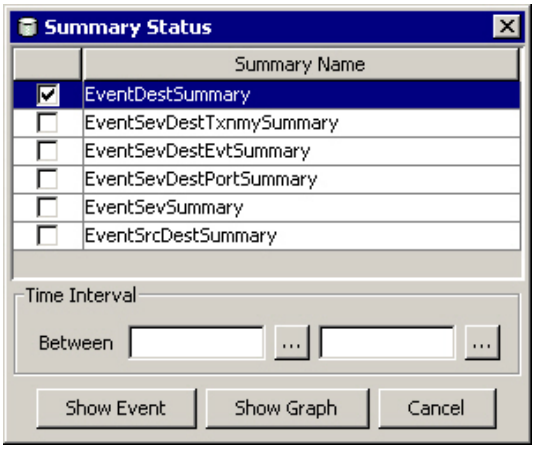


4. Select a time interval.
5. Click the 'Show Event' button.
6. The Eventfiles needed to complete the Summary display in a list format.

**NOTE**: To complete summaries, see “Run EventFile(s) for a Summary” section.

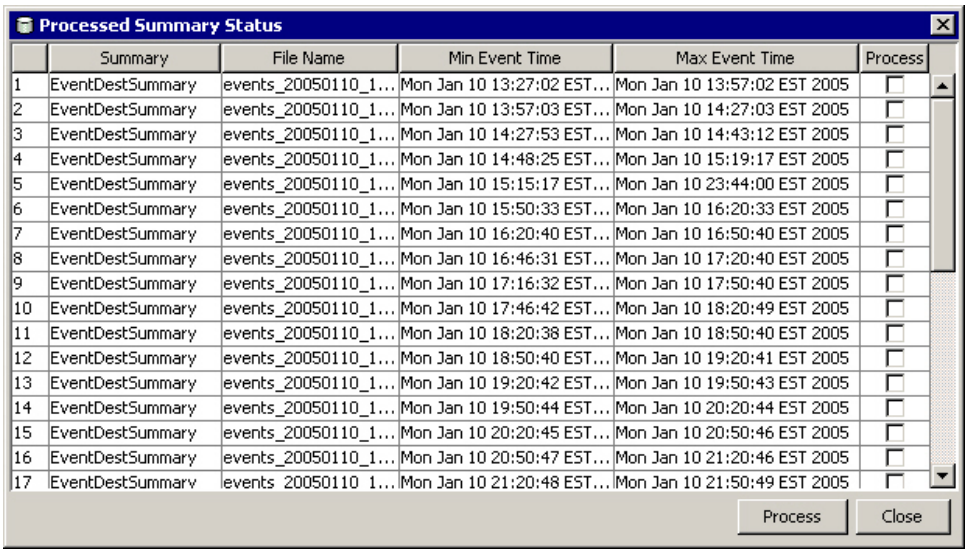


### Running Eventfiles for a Summary

1. Click the Reporting Data Tab.
2. Select the Status button.
3. Choose the Summary or Summaries you wish to query.
4. Select a time interval.
5. Click the 'Show Event' button.
6. The Eventfiles needed to complete the Summary display in a list format.
7. Check the Eventfiles that you would like to run so that the Summary is complete.

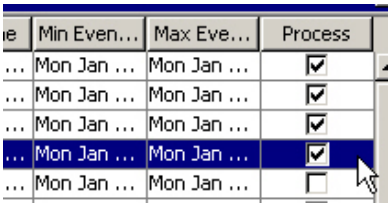


8. Click Process.

## SDM Command Line

**NOTE**: If your machine does not have access to DAS_Binary and DAS_Query, the SDM Command Line can be used in place of the SDM GUI.

## Saving Connection Properties for Sentinel Data Manager

This must be performed prior to using any of the Sentinel Data Manager Command Line actions other than saveConnection.

If you have run the SDM GUI, you can use the sdm.connect file that was created from the GUI. It is located %ESEC_HOME%\sdm for Windows and $ESEC_HOME/sdm for UNIX.

The save connection function saves the following connection details along with the encrypted password (using the keystore specified in configuration.xml) to the file specified.

This command uses the following flags:

| | |
|---|---|
| -action | saveConnection |
| -server | <oracle or mssql> |
| -host | <database host IP Address or host name to connect to> |
| -port | <database port number to connect to [Oracle default: 1521/SQL Server default: 1433]> |
| -database | <database name/SID to connect to> |
| -user | <database username> |
| -password | <database password> |
| -winAuth | Used for Windows authentication. When using this option, do not use -user and –password. |
| -connectFile | <filename to save the connection details [file name of your choosing]> |

The application saves all the above connection details along with the encrypted password to the file specified. The application uses the saved connection details to execute the rest of the commands. This step should be completed first time you start the application and every time you want to change the connection details the application uses.

### Running saveConnection

1. Execute the command as follows:

```
sdm -action saveConnection -server <oracle/mssql> -
   host <hostIp/hostName> -port <portnum> -database
   <databaseName/SID> [-driverProps
   <propertiesFile>] {-user <dbUser> -password
   <dbPass> | -winAuth} -connectFile
   <filenameToSaveConnection>
```

   The following example will save connections for a host with an IP address of 172.16.0.36 at port 1521 (default for Oracle, for SQL Server, default is 1433).

   - Oracle Example:

```
./sdm -action saveConnection -server oracle -host
   172.16.0.36 -port 1521 -database esec -user
   esecdba -password XXXXXX -connectFile sdm.connect
```

   - SQL Server Example:

```
sdm -action saveConnection -server mssql -host
   172.16.0.36 -port 1433 -database esec -user
```

```
esecdba -password XXXXXX -connectFile
sdm.connect
```

The following example will save connections for a host with an IP address of 172.16.0.36, port 1433 with database name of esec_51 for Windows authentication.

- SQL Server Example (Windows Authentication):

```
sdm -action saveConnection -server mssql -host
   172.16.1.3 -port 1433 -database esec_51 –winAuth
   -connectFile  %ESEC_HOME%\sdm\sdm.connect
```

This will save the connection details to the sdm.connect file. All the rest of the commands will take this filename as input to connect to the designated database and to perform their actions.

## Partition Management

### Partition Configuration

This is for Oracle only. This action (partitionConfig) is used to configure your database partitions. This configuration drives how the partitions are added to all the e-Security partitioned tables. This action uses the following flags:

| | |
|---|---|
| -action | partitionConfig |
| -freq | <either “3D” or “2D” or “1D” or 1W><br><br>The following are the only options supported<br>3D - three partitions per day<br>2D - two partitions per day<br>1D - one partition per day<br>1W - one partition per week |
| -days | <Number of days to be added whenever addPartitions is chosen> |
| -connectFile | <path to the filename saved by saveConnection> |

Running partitionConfig

1. Execute this command as follows:

```
./sdm -action partitionConfig –freq <either 3D or
   2D or 1D or 1W> -days <Number Of days to be added
   whenever “addPartitions” is chosen> -connectFile
   <path to the filename saved by “saveConnection”
   (default: $ESEC_HOME/sdm/sdm.connect)>
```

The following example the system will add thirty partitions (3 partitions per 1 DAY = 3 * 10).

```
./sdm -action partitionConfig -freq 3D -days 10 -
   connectFile sdm.connect
```

The following example the system will add ten partitions (1 partitions per 1 DAY = 1 * 10).

```
./sdm -action partitionConfig -freq 1D -days 10 -
   connectFile sdm.connect
```

The following example the system will add one partition (1 partitions per 7 days = 1 * 10/7).

```
./sdm -action partitionConfig -size 1W -days 10 -
   connectFile sdm.connect
```

### Adding Partitions

This action (addPartitions) adds the required number of partitions according to the partition configuration in the following tables:

- Oracle:
  - EVENTS
  - CORRELATED_EVENTS
  - EVT_DEST_EVT_NAME_SMRY_1
  - EVT_DEST_SMRY_1
  - EVT_DEST_TXNMY_SMRY_1
  - EVT_PORT_SMRY_1
  - EVT_SEV_SMRY_1
  - EVT_SRC_SMRY_1
- SQL Server
  - EVENTS
  - CORRELATED_EVENTS
  - EVT_DEST_EVT_NAME_SMRY_1
  - EVT_DEST_SMRY_1
  - EVT_DEST_TXNMY_SMRY_1
  - EVT_PORT_SMRY_1
  - EVT_SEV_SMRY_1
  - EVT_SRC_SMRY_1

If you are configured to have 10 days worth of partitions, every time you run “addPartitions” it checks to see if you have 10 days of partitions ahead. If you have enough partitions for next 10 days it will not do anything. If not, it will add the required number of partitions for 10 days.

This action uses the following flags:

| | |
|---|---|
| -action | addPartitions |
| -connectFile | <path to the filename saved by “saveConnection”> |

#### Running addPartitions

1. Execute this command as follows:

   ```
   sdm -action addPartitions -connectFile <path to the
      filename saved by “saveConnection”>
   ```

   Oracle Example:

   ```
   ./sdm -action addPartitions -connectFile
      sdm.connect
   ```

   SQL Server Example:

   ```
   sdm -action addPartitions -connectFile sdm.connect
   ```

### Dropping Partitions

This action (dropPartition) drops all the partitions older than the flag keepDays from the following tables:

- Oracle:
  - EVENTS
  - CORRELATED_EVENTS
  - EVT_DEST_EVT_NAME_SMRY_1
  - EVT_DEST_SMRY_1
  - EVT_DEST_TXNMY_SMRY_1
  - EVT_PORT_SMRY_1
  - EVT_SEV_SMRY_1
  - EVT_SRC_SMRY_1
- SQL Server
  - EVENTS
  - CORRELATED_EVENTS
  - EVT_DEST_EVT_NAME_SMRY_1
  - EVT_DEST_SMRY_1
  - EVT_DEST_TXNMY_SMRY_1
  - EVT_PORT_SMRY_1
  - EVT_SEV_SMRY_1
  - EVT_SRC_SMRY_1

This action does not drop any partitions that are not archived. If you want to delete unarchived partitions, use the “forceDelete” flag. If forceDelete is used:

| | |
|---|---|
| false or not specified | drops only the partitions older than keepDays and those that are archived |
| true | drops all the partitions older than keepDays including unarchived partitions |

This action uses the following flags:

| | |
|---|---|
| -action | dropPartitions |
| -keepDays | <number of days to keep> |
| [-forceDelete] | <either “true” or “false”> |
| -connectFile | <path to the filename saved by “saveConnection”> |

**NOTE:** If you drop a partition that has not been archived it cannot be imported.

#### Running dropPartition

1. Execute this command as follows:

```
sdm -action dropPartitions [-forceDelete <false>] -
   keepDays <number> -connectFile <path to the
   filename saved by “saveConnection”>
```

   The following examples drops all the partitions older than 30 days making sure all the partitions are archived. All partitions that were skipped (not removed) because they have not been archived are listed when the operation completes.

   Oracle Example:

```
./sdm -action dropPartitions -keepDays 30 -
   connectFile sdm.connect


./sdm -action dropPartitions -forceDelete false -
   keepDays 30 -connectFile sdm.connect
```

SQL Server Example:

```
sdm -action dropPartitions -keepDays 30 -
   connectFile sdm.connect


sdm -action dropPartitions -forceDelete false -
   keepDays 30 -connectFile sdm.connect
```

**Viewing Partition Summaries**

This action (ViewPartitions) displays the partition summary of the following supported tables:

- Oracle:
  - EVENTS
  - CORRELATED_EVENTS
  - EVT_DEST_EVT_NAME_SMRY_1
  - EVT_DEST_SMRY_1
  - EVT_DEST_TXNMY_SMRY_1
  - EVT_PORT_SMRY_1
  - EVT_SEV_SMRY_1
  - EVT_SRC_SMRY_1
- SQL Server
  - EVENTS
  - CORRELATED_EVENTS
  - EVT_DEST_EVT_NAME_SMRY_1
  - EVT_DEST_SMRY_1
  - EVT_DEST_TXNMY_SMRY_1
  - EVT_PORT_SMRY_1
  - EVT_SEV_SMRY_1
  - EVT_SRC_SMRY_1

This command uses the following flags:

| | |
|---|---|
| -action | startGui |
| -tableName | <name of one the above named tables> |
| -connectFile | <path to the filename saved by “saveConnection”> |

To View Partition Summaries

1. Execute this command as follows:

```
sdm -action viewPartitions -tableName <table name>
   –connectFile <path to the filename saved by
   “saveConnection”>
```

   The following example, displays the list of partitions of the EVENTS table and status of each partition.

- Oracle Example:

```
./sdm -action viewPartitions -tableName EVENTS -
   connectFile sdm.connect
```

- SQL Server Example:

```
sdm -action viewPartitions –tableName EVENTS –
   connectFile sdm.connect
```

## Archive Management

### Archive Configuration

This action (archiveConfig) is used to configure archiving. This configuration drives how the data is archived from the e-Security tables.

This action uses the following flags:

| | |
|---|---|
| -action | archiveConfig |
| -dirPath | <valid directory path to write the archived files to> |
| -keepDays | <number of days to keep> |
| -fileSize | (Oracle only) <maximum size of each archived file. Specify either KB, MB or GB> |
| -connectFile | <path to the filename saved by "saveConnection"> |

For Oracle, the dirPath directory path should be specified as UTL_FILE_DIR parameter in init.ora file according to Oracle requirements. You should have one of the following:

- UTL_FILE_DIR = *
- UTL_FILE_DIR = specific directory where you want to write files to in your init.ora file

Running archiveConfig

1. Execute this command as follows:

```
sdm -action archiveConfig  -dirPath <directory path
   to write the archived files to> -keepDays <number
   of days to keep> -fileSize <maximum size of each
   archived file, specified in KB, MB or GB> -
   connectFile <path to the filename saved by
   "saveConnection">
```

- Oracle Example:

The following example archives all data older than 13 days to /tmp directory in chunks greater than 1GB.

```
./sdm -action archiveConfig  -dirPath
   /tmp -keepDays 13 -fileSize 1GB -connectFile
   sdm.connect
```

The following example archives all data older than 13 days to /tmp directory in chunks greater than 40MB.

```
./sdm -action archiveConfig -dirPath /tmp -keepDays
   13 -fileSize 40MB -connectFile sdm.connect
```

### Archiving Data

Run this action (archiveData) after you set your archive configuration (archiveConfig). This action archives the data from the given table name according to the archive configuration. It archives data from:

- Oracle:
  - EVENTS
  - CORRELATED_EVENTS
- SQL Server
  - EVENTS
  - CORRELATED_EVENTS

**NOTE**: Aggregation tables are not archived.

This command uses the following flags:

-action archiveData
-connectFile <path to the filename saved by "saveConnection">

Running archiveData

1. Execute this command as follows:

```
sdm -action archiveData -connectFile <path to the
   filename saved by "saveConnection">
```

   - Oracle Example:

   The following example archives events, their custom and reserved values and correlated events from the EVENTS, EVT_RESERVED_VALUES, EVT_CUSTOM_VALUES and ASSOCIATIONS table according to the value set in your archive configuration (archiveConfig). Using the value set in the example provided under the section on Archive Management, this will archives data older than 13 days.

```
./sdm -action archiveData -connectFile sdm.connect
```

   - SQL Server Example:

   The following example archives events and correlated events according to the value set in your archive configuration archiveConfig). Using the value set in the example provided under the section on Archive Management, this will archives data older than 13 days.

```
sdm -action archiveData -connectFile sdm.connect
```

### Deleting Data

This action (deleteData) deletes the data older than keep days from the given table name. It deletes data from:

- Oracle:
  - EVENTS
  - CORRELATED_EVENTS
  - EVT_DEST_EVT_NAME_SMRY_1
  - EVT_DEST_SMRY_1
  - EVT_DEST_TXNMY_SMRY_1
  - EVT_PORT_SMRY_1
  - EVT_SEV_SMRY_1

- EVT_SRC_SMRY_1

- SQL Server
  - EVENTS
  - CORRELATED_EVENTS
  - EVT_DEST_EVT_NAME_SMRY_1
  - EVT_DEST_SMRY_1
  - EVT_DEST_TXNMY_SMRY_1
  - EVT_PORT_SMRY_1
  - EVT_SEV_SMRY_1
  - EVT_SRC_SMRY_1

This action does not drop any partitions that are not archived. If you want to delete unarchived partitions, the optional flag “forceDelete” has to be specified with a value of true. If forceDelete is used:

| | |
|---|---|
| false or not specified | drops only the partitions older than keepDays and those that are archived |
| true | drops all the partitions older than keepDays including unarchived partitions |

This command uses the following flags:

| | |
|---|---|
| -action | deleteData |
| -keepDays | <number of days to keep> |
| [-forceDelete] | <either true or false> |
| -connectFile | <path to the filename saved by “saveConnection”> |

### Running deleteData

1. Execute this command as follows:

   ```
   sdm -action deleteData -keepDays <number of days to
      keep> -connectFile <path to the filename saved by
      “saveConnection”>
   ```

   - Oracle Example:

   The following example drops partitions from all tables older than 13 days making sure all dropped partitions are archived. In the end, a list is generated of any partitions that were not deleted if they have not been archived.

   ```
   ./sdm -action deleteData -keepDays 13 -connectFile
      sdm.connect
   ```

   - SQL Server Example:

   The following example drops the partitions from all tables older than 13 days making sure all dropped partitions are archived. In the end, it lists any partitions that were not deleted if they have not been archived.

   ```
   sdm -action deleteData -keepDays 13 -connectFile
      sdm.connect
   ```

## Import Management

### Listing Files to Import

This action (filesToImport) is used to list the files needed to import the data between the given dates of the following supported tables:

- Oracle:
  - HIST_EVENTS
  - HIST_CORRELATED_EVENTS
- SQL Server
  - HIST_EVENTS
  - HIST_CORRELATED_EVENTS

This command uses the following flags:

| | |
|---|---|
| -action | filesToImport |
| -startDate | <mm/dd/yyyy hh24:mi:ss> |
| -endDate | <mm/dd/yyyy hh24:mi:ss> |
| -connectFile | <path to the filename saved by “saveConnection”> |

**NOTE:** hh24 is hours represented in 24 hour format. For example, 1:15:00 p.m. is 13:15:00 and 3:00:00 a.m. is 03:00:00.

Running filesToImport

1. Execute this command as follows:

```
sdm -action filesToImport -startDate <mm/dd/yyyy
   hh24:mi:ss> -endDate <mm/dd/yyyy hh24:mi:ss> -
   connectFile <path to the filename saved by
   “saveConnection”>
```

   The following example lists all files containing data between dates "09/25/2003 00:00:00" (Sep 25th midnight) and "09/26/2003 00:00:00" (Sep 26th midnight) that has been archived earlier and can be imported back.

   - Oracle Example:

```
./sdm -action filesToImport -startDate 09/25/2003
   00:00:00 -endDate 09/26/2003 00:00:00 -
   connectFile sdm.connect
```

   - SQL Server Example:

```
sdm -action filesToImport -startDate 09/25/2003
   00:00:00 -endDate 09/26/2003 00:00:00 -
   connectFile sdm.connect
```

   The following example lists all the files containing the data between dates "09/25/2003 16:00:00" (Sep 25th 4 PM) and "09/26/2003 18:00:00" (SEP 26, 6 PM) that has been archived earlier and can be imported back.

   - Oracle Example:

```
./sdm -action filesToImport -startDate 09/25/2003
   16:00:00 -endDate 09/26/2003 18:00:00 -
   connectFile sdm.connect
```

   - SQL Server Example:

```
sdm -action filesToImport -startDate 09/25/2003
   16:00:00 -endDate 09/26/2003 18:00:00 -
   connectFile sdm.connect
```

**Importing Data**

This action (importData) imports data between the given dates into the following supported tables:

- Oracle:
  - HIST_EVENTS
  - HIST_CORRELATED_EVENTS
- SQL Server
  - HIST_EVENTS
  - HIST_CORRELATED_EVENTS

If the data has already been imported or there is no archived data is found between the specified dates, it returns a message.

The application imports each file into a table and builds the historical view on all the historical tables. The report view joins on the original table and historical view. All reports use the report view and thus will see any imported data.

This command uses the following flags:

| | |
|---|---|
| -action | importData |
| -startDate | <mm/dd/yyyy hh24:mi:ss> |
| -endDate | <mm/dd/yyyy hh24:mi:ss> |
| -dirPath | <directory to import files from> |
| -connectFile | <path to the filename saved by “saveConnection”> |

**NOTE:** hh24 is hours represented in 24 hour format. For example, 1:15:00 p.m. is 13:15:00 and 3:00:00 a.m. is 03:00:00.

Running importData

1. Place all the files you wish to import in a specific directory (i.e., dirPath - <directory to import files from>).
2. Execute this command as follows:

   ```
   sdm -action importData -dirPath <directory to
      import files from> -startDate <mm/dd/yyyy
      hh24:mi:ss> -endDate <mm/dd/yyyy hh24:mi:ss> -
      connectFile <path to the filename saved by
      “saveConnection”>
   ```

   The following example imports the archived files from the tmp directory containing the data between dates "09/25/2003 00:00:00" (Sep 25 midnight) and "09/26/2003 00:00:00" (Sep 26 midnight) into the above mentioned tables.

   - Oracle Example:

     ```
     ./sdm -action importData -dirPath /tmp -startDate
        09/25/2003 00:00:00 -endDate 09/26/2003
        00:00:00 -connectFile sdm.connect
     ```

   - SQL Server Example:

```
sdm -action importData -dirPath c:\tmp -startDate
   09/25/2003 00:00:00 -endDate 09/26/2003
   00:00:00 -connectFile sdm.connect
```

The following example imports the archived files from the tmp directory containing the data between dates "09/25/2003 08:30:00" (Sep 25 8:30 AM) and "09/26/2003 20:00:00" (Sep 26 8:00 PM) into the above mentioned tables.

- Oracle Example:

```
./sdm -action importData -dirPath /tmp -startDate
   09/25/2003 08:00:00 -endDate 09/26/2003
   20:00:00 -connectFile sdm.connect
```

- SQL Server Example:

```
sdm -action importData -dirPath c:\tmp -startDate
   09/25/2003 08:00:00 -endDate 09/26/2003
   20:00:00 -connectFile sdm.connect
```

**Deleting Imported Data**

This action (dropImported) deletes the imported data between the given dates from the following supported tables:

- Oracle:
  - HIST_EVENTS
  - HIST_CORRELATED_EVENTS
- SQL Server
  - HIST_EVENTS
  - HIST_CORRELATED_EVENTS

If there is no data imported between two specified dates, it returns a message.

This command uses the following flags:

| | |
|---|---|
| -action | dropImported |
| -startDate | <mm/dd/yyyy hh24:mi:ss> |
| -endDate | <mm/dd/yy hh24:mi:ss> |
| -connectFile | <path to the filename saved by “saveConnection”> |

**NOTE:** hh24 is hours represented in 24 hour format. For example, 1:15:00 p.m. is 13:15:00 and 3:00:00 a.m. is 03:00:00.

Running dropImported

1. Execute this command as follows:

```
sdm -action dropImported -startDate <mm/dd/yyyy
   hh24:mi:ss> -endDate <mm/dd/yyyy hh24:mi:ss> -
   connectFile <path to the filename saved by
   “saveConnection”>
```

The following example deletes the imported data between the given dates from the above mentioned tables.

- Oracle Example:

```
./sdm -action dropImported -startDate 09/25/2003
   00:00:00 -endDate 09/26/2003 00:00:00 -
   connectFile sdm.connect
```

- SQL Server Example:

```
sdm -action dropImported -startDate 09/25/2003
   00:00:00 -endDate 09/26/2003 00:00:00 -
   connectFile sdm.connect
```

## Tablespace Management

In Tablespace Management, you have a command line option and GUI option. The command line allows you to:

- View Sentinel database space usage

The GUI allows you to:

- View partitions
- View archived partitions
- View import partitions
- View space usage

### Viewing Sentinel Database Space Usage (Command Line)

This action (dbstats) displays the Sentinel database usage for all Sentinel tablespaces in Oracle and Sentinel filegroups in MS SQL.

This command uses the following flags:

| | |
|---|---|
| -action | dbstats |
| -connectFile | <path to the filename saved by “saveConnection”> |

Viewing Sentinel Database Space Usage (Command Line)

1. Execute the following command:

```
sdm -action dbStats -connectFile <path to the
   filename saved by ”saveConnection”>
```

- Oracle Example:

The following example displays the tablespaces of Sentinel database with their total space, used space and free space available.

```
./sdm -action dbStats -connectFile sdm.connect
```

- SQL Server Example:

The following example displays the file groups of Sentinel database with their total space, used space and free space available.

```
sdm -action dbStats -connectFile sdm.connect
```

## Updating Mappings (Command Line)

This action (updateMapData) allows you to replace a map source data file with another. Your new source data file should have the same delimiters, key columns and activated column of your previous mapping. If not, use the Edit feature of the SDM GUI.

This command uses the following flags:

| | |
|---|---|
| -action | updateMapData |
| -map | <map name> |
| -file | <filename> |
| -backup | <true/false> (default: true) |
| -connectFile | <path to the filename saved by "saveConnection"> |

The –backup flag allows you to backup the original mapping file in the map_data folder. The backed-up data map file will be saved as .bak file with a set of random numbers at the end of file. For example: threat10197.bak.

Updating (replace) a Mapping

1. Execute the following command:

```
sdm -action updateMapData -map <mapName> -file
   <fileName> [-backup <true/false> (DEFAULT: true)]
   -connectFile <path to the filename saved by
   "saveConnection">
```

   The following example replaces the mappings in the map 'threat' with the mappings from the map file 'vuln_attacks.txt'.

```
sdm -action updateMapData -map threat -file
   vuln_attacks.txt -connectFile sdm.connect
```

Since the flag –backup was not used, the default operation will create a backup of the original mapping prior to updating it the map file 'vuln_attack.txt'.

## Using eSecurity Supplied Auto Manage Script (Windows Only)

eSecurity has developed a batch file that can be scheduled so that many of the management actions of SDM can be preformed automatically.

> **NOTE**: If your machine does not have access to DAS_Binary and DAS_Query, the SDM Command Line can be used in place of the SDM GUI.

This procedure is only applicable to Windows. Ensure that while performing your pre-configuration and configuration that the following is done:

- Make sure sdm.connect is initialized either by using SDM GUI or command line.
- Make sure the archive directory exists.
- Make sure the archiveConfig & dropPartitions days are equal.
- Make sure the batch file runs correctly from command prompt at least once before scheduling it to run automatically.

> **NOTE**: If the scheduled task fails, it will not send a notification. It will log it in SDM_*.log

## Setting up Manage_data.bat file to Archive Data and Add Partitions

### Pre-Configuration

Prior to automatically setting Archive Data and Add Partitions, you must:

- Save connection properties
- Establish archival parameters

> **NOTE**: If you saved a connect file to a different location or filename than the default (%ESEC_HOME%\sdm\sdm.connect), you will have edit the manage_data.bat file to update the path to your connect file.

### Establishing Archival Parameters

This can be done using the Command Line.

This action (archiveConfig) is used to configure archiving. This configuration drives how the data is archived from the e-Security tables.

This action uses the following flags:

| | |
|---|---|
| -action | archiveConfig |
| -dirPath | <valid directory path to write the archived files to> |
| -keepDays | <number of days to keep> |
| -connectFile | <path to the filename saved by "saveConnection"> |

#### Establishing Archival Parameters via the Command Line

1. Create an archive output directory at the root called SDM_archive (c:\SDM_archive).

   > **NOTE**: If you create a different output directory or location, you will have to edit the manage_data.bat file.

2. Execute this command as follows:

   ```
   sdm -action archiveConfig  -dirPath <directory path
      to write the archived files to> -keepDays <number
      of days to keep> -connectFile <path to the
      filename saved by "saveConnection">
   ```

   The following example archives all data older than 30 days to c:\SDM_archive directory.

   ```
   sdm -action archiveConfig –dirpath c:\SDM_archive -
      keepDays 30 -connectFile sdm.connect
   ```

#### Establishing Archival Parameters via the GUI

1. Create an archive output directory at the root called SDM_archive (c:\SDM_archive).

   > **NOTE**: If you create a different output directory or location, you will have to edit the manage_data.bat file.

2. The SDM GUI does not require archival parameters. The GUI can directly archive data without having to establish archival parameters.

**Delete Data (Drop Partitions)**

This action (deleteData) deletes the data older than keep days from the given table name. It deletes data from:

- EVENTS
- CORRELATED_EVENTS
- EVT_DEST_EVT_NAME_SMRY_1
- EVT_DEST_SMRY_1
- EVT_DEST_TXNMY_SMRY_1
- EVT_PORT_SMRY_1
- EVT_SEV_SMRY_1
- EVT_SRC_SMRY_1

This action does not drop any partitions that are not archived. If you want to delete unarchived partitions, the optional flag “forceDelete” has to be specified with a value of true. If forceDelete is used:

| | |
|---|---|
| false or not specified | drops only the partitions older than keepDays and those that are archived |
| true | drops all the partitions older than keepDays including unarchived partitions |

This command uses the following flags:

| | |
|---|---|
| -action | deleteData |
| -keepDays | <number of days to keep> |
| [-forceDelete] | <either true or false> |
| -connectFile | <path to the filename saved by “saveConnection”> |

Running deleteData

1. Execute this command as follows:

```
sdm -action deleteData -keepDays <number of days to
   keep> -connectFile <path to the filename saved by
   “saveConnection”>
```

   The following example drops the partitions from tables older than 30 days making sure all dropped partitions are archived. In the end, it lists any partitions that were not deleted if they have not been archived.

```
sdm -action deleteData -keepDays 30 -connectFile
   sdm.connect
```

## Scheduling Manage_data.bat to Archive Data and Add Partitions

**NOTE**: The manage_data.bat file is set to a keep day value of 30, archive output to c:\SDM_archive and connect file to %ESEC_HOME%\sdm\sdm.connect. If your values are different, you will need to edit the manage_data.bat file.

If you have set your connection properties and archival parameters, run the manage_data.bat from the command prompt to ensure that it works.

### To Automatically Archive Data and Add Partitions

> **NOTE**: The following steps are for Windows 2000 Professional. Steps for Windows 2000 Server and XP may be different, but similar.

1. In Windows, click Start > Setting > Control Panel.
2. Double-click 'Scheduled Tasks'.
3. Double-click 'Add Scheduled Task'. Click Next.
4. Click the Browse button and navigate to the manage_data.bat file.
5. Enter a name for the scheduled task such as SDM_Archive. Select Daily under 'Perform this task:'. Click Next.
6. Select a time a day to run this task. Click Next.
7. Enter a time and date of choice. Click Next.

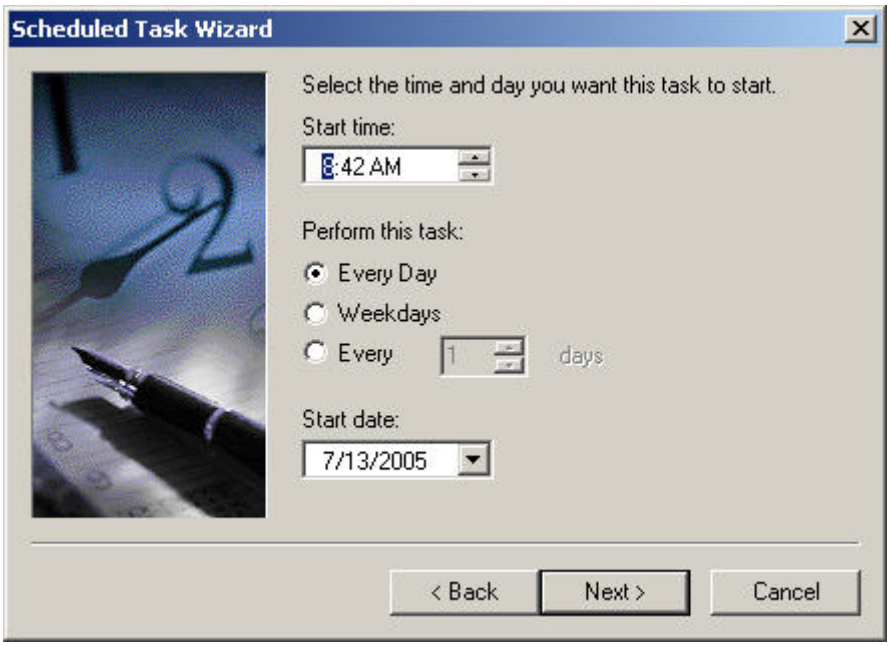


8. Enter a user that this task will run under. The user cannot be the local system account. It must be run as a specific user. Click Next.
9. Click Finish to complete as scheduled task.

# Chapter 11 – Utilities

## Starting and Stopping the Sentinel Server and Agent Manager - UNIX

### Starting the UNIX Sentinel Server

On UNIX, starting Sentinel Server also starts the Communication Server.

Starting the UNIX Sentinel Server

1. As user esecadm, cd to the $ESEC_HOME/sentinel/scripts directory.
2. Run the following command:

```
./sentinel.sh start
```

### Stopping the UNIX Sentinel Server

On UNIX, stopping Sentinel Server also stops the Communication Server.

Stopping the UNIX Sentinel Server

1. As user esecadm, cd to $ESEC_HOME/sentinel/scripts directory.
2. Run the following command:

```
./sentinel.sh stop
```

### Starting the UNIX Agent Manager

Starting the UNIX Agent Manager

1. As user esecadm, cd to the $WORKBENCH_HOME.
2. Run the following command:

```
./agent-manager.sh start
```

### Stopping the UNIX Agent Manager

Stopping the UNIX Agent Manager

1. As user esecadm, cd to the $WORKBENCH_HOME.
2. Run the following command:

```
./agent-manager.sh stop
```

## Starting and Stopping the Sentinel Server and Agent Manager - Windows

Depending upon your installation configuration, you can have up to three e-Security services running on your machine. They are:

- eSecurity – Watchdog, this service starts all other sentinel server processes.
- eSecurity Communication – This service is your encrypted Communication Server.
- Agent Manager – This service is your Wizard.

Under Windows Services, you can manually start, restart and stop any of these services.

## Starting the Windows Agent Manager

Starting the Windows Agent Manager

1. Click Start > Settings > Control Panel.
2. Double-click on Administrative Tools.
3. Double-click on Services.
4. Right click on Agent Manager > Start.

## Stopping the Windows Agent Manager

Stopping the Windows Agent Manager

1. Click Start > Settings > Control Panel.
2. Double-click on Administrative Tools.
3. Double-click on Services.
4. Right click on Agent Manager > Stop.

## Starting the Sentinel Server for Windows

Starting the Windows Sentinel Server

1. Click Start > Settings > Control Panel.
2. Double-click on Administrative Tools.
3. Double-click on Services.
4. In the Services window, highlight eSecurity.
5. Right-click > Start or click the start button in the tool bar.

## Stopping the Sentinel Server for Windows

Stopping the Windows Sentinel Server

1. Click Start > Settings > Control Panel.
2. Double-click on Administrative Tools.
3. Double-click on Services.
4. In the Services window, highlight eSecurity.
5. Right-click > Stop or click the stop button in the tool bar.

## Starting the Sentinel Communication Server for Windows

Starting the Windows Sentinel Communication Server

1. Click Start > Settings > Control Panel.
2. Double-click on Administrative Tools.
3. Double-click on Services.
4. In the Services window, highlight eSecurity Communication.
5. Right-click > Start or click the start button in the tool bar.

### Stopping the Sentinel Communication Server for Windows

Stopping the Windows Sentinel Communication Server

1. Click Start > Settings > Control Panel.
2. Double-click on Administrative Tools.
3. Double-click on Services.
4. In the Services window, highlight eSecurity Communication.
5. Right-click > Stop or click the stop button in the tool bar.

## Sentinel Script Files

Depending upon your installation configuration, the ESEC_HOME/sentinel/scripts directory may contain some or all of the following script files:

| Script File: | Description: |
|---|---|
| ▪ remove_sonic_lock.bat | This script removes the communication server lock file(s). |
| ▪ start_broker.bat<br>▪ start_broker.sh | These scripts start the communication server on the command line in console mode. |
| ▪ stop_broker.bat<br>▪ stop_broker.sh | These scripts stop the communication server on the command line in console mode. |
| ▪ stop_container.bat<br>▪ stop_container.sh | This script restart the following containers:<br>▪ DAS_Aggregation<br>▪ DAS_RT<br>▪ DAS_iTRAC<br>▪ DAS_Binary<br>▪ DAS_Query |
| ▪ sentinel.sh | This script stops or starts the Sentinel Server. See Starting the UNIX Sentinel Server or Stopping the UNIX Sentinel Server. |

### Removing the Communication Server Lock Files

In the event of an improper shutdown, the communication server may be locked. After removing the lock files, you will have to restart the communication server. These files are located:

For Windows:

```
%ESEC_HOME%\3rdparty\SonicMQ\MQ6.1\esecDomain\data\_MF
   System\lock
%ESEC_HOME%\3rdparty\SonicMQ\MQ6.1\SonicMQStore\db.lck
```

For UNIX:

```
$ESEC_HOME/3rdparty/SonicMQ/MQ6.1/esecDomain/data/_MFS
   ystem/lock
$ESEC_HOME /3rdparty/SonicMQ/MQ6.1/SonicMQStore/db.lck
```

### Removing the Communication Server Lock File (Windows)

1. Either cd or navigate via Windows Explorer to:

   ```
   %ESEC_HOME%\sentinel\scripts
   ```

2. Either double-click on (via Windows Explorer) or execute the following file:

   ```
   remove_sonic_lock.bat
   ```

### Removing the Communication Server Lock File (UNIX)

Removing the lock file on UNIX is typically not required because the lock file is usually removed automatically when Sentinel Server is started. If these files need to be removed manually, you must remove them using typical UNIX file system commands (such as rm).

## Starting the Communication Server in Console Mode

These scripts start the communication server on the command line in console mode. These scripts are useful for debugging the communication server without forcing you to run the rest of Sentinel Server. During normal operations, you should not use these script (use the instructions at Starting the UNIX Sentinel Server or Starting the Sentinel Server for Windows instead).

### Starting the Communication Server (Windows)

> **NOTE**: When starting this script in Windows, it will not indicate as started in the Services window and will only run if the Command Prompt window remains open.

1. Either cd or navigate via Windows Explorer to:

   ```
   %ESEC_HOME%\sentinel\scripts
   ```

2. Either double-click on (via Windows Explorer) or execute the following file:

   ```
   start_broker.bat
   ```

### Starting the Communication Server (UNIX)

1. Login as user esecadm.
2. cd to:

   ```
   $ESEC_HOME/sentinel/scripts
   ```

3. Enter:

   ```
   ./start_broker.sh
   ```

## Stopping the Communication Server in Console Mode

These scripts stop the communication server on the command line in console mode. These scripts are useful for debugging the communication server without forcing you to stop the rest of Sentinel Server. During normal operations, you should not use these scripts (use the instructions at Stopping the UNIX Sentinel Server or Stopping the Sentinel Server for Windows instead).

### Stopping the Communication Server (Windows)

1. Either cd or navigate via Windows Explorer to:

   ```
   %ESEC_HOME%\sentinel\scripts
   ```

2. Either double-click on (via Windows Explorer) or execute the following file:

   ```
   stop_broker.bat
   ```

### Stopping the Communication Server (UNIX)

1. Login as user esecadm.
2. cd to:

   ```
   $ESEC_HOME/sentinel/scripts
   ```

3. Enter:

   ```
   ./stop_broker.sh
   ```

## Restarting Sentinel Containers

These scripts restarts the containers listed below. The script sends a message to the specified service to shut itself down. The Sentinel Watchdog then restarts the service.

When running Sentinel Server on Linux, the preferred method of stopping, starting, or restarting these container services is to use the Server Views in the Admin tab of Sentinel Control Center.

| Name | Description |
|---|---|
| ▪ DAS_Aggregation | (das_aggregation.xml) used for executing and configuring the aggregation service. |
| ▪ DAS_RT | (das_rt.xml) used for executing and configuring real time views service. |
| ▪ DAS_iTRAC | (das_itrac.xml) used for configuring the iTRAC service. |
| ▪ DAS_Binary | (das_binary.xml) used for event and correlated event insertion operation. |
| ▪ DAS_Query | (das_query.xml) all other database operations. |

### Restarting a Sentinel Container (Windows)

1. cd to:

   ```
   %ESEC_HOME%\sentinel\scripts
   ```

2. Enter:

   ```
   stop_container.bat <host machine> <container name>
   ```

   For example:

   ```
   stop_container.bat localhost DAS_RT
   ```

### Restarting a Sentinel Container (UNIX)

1. Login as user esecadm.

2. cd to:

   ```
   $ESEC_HOME/sentinel/scripts
   ```

3. Enter:

   ```
   ./stop_container <host machine> <container name>
   ```

   For example:

   ```
   ./stop_container localhost DAS_RT
   ```

## Version Information

### Sentinel Server Version Information

Sentinel Server has a command line option to display the version information of the following processes:

- watchdog
- rulelg_checker
- correlation_engine
- data_synchronizer
- query_manager
- DAS

How to get Sentinel version information (UNIX)

1. cd to:

   ```
   $ESEC_HOME/sentinel/bin
   ```

2. Enter:

   ```
   ./<process> -version
   ```

   For example:

   ```
   ./correlation_engine –version
   ```

How to get Sentinel version information (Windows)

1. cd to:

   ```
   %ESEC_HOME%\sentinel\bin
   ```

2. Enter:

   ```
   <process> -version
   ```

   For example:

   ```
   correlation_engine –version
   ```

### e-Security .dll and .exe File Version Information

How to get e-Security .dll and .exe file version information

1. cd to %ESEC_HOME%.
2. Within the various different sub-directories, right-click on either a .dll or .exe file and select properties.
3. Click the version tab.

4. In the Item Name pane, select Product Version. The version number of the file will appear in the Value pane.

### e-Security .jar Version Information

How to get e-Security .jar file version information

1. At the Sentinel Server, login as user:
   For UNIX:
   ```
   esecadm
   ```
   For Windows, login as a user with rights to Sentinel Server.
2. cd to:
   For UNIX:
   ```
   $ESEC_HOME/utilities
   ```
   For Windows:
   ```
   %ESEC_HOME%\utilities
   ```
3. At the command line, enter:
   For UNIX:
   ```
   ./versionreader.sh <path/jar file name>
   ```
   For Windows
   ```
   versionreader <path/jar file name>
   ```

## Configuring Sentinel email

Sentinel email configuration settings are stored in the execution.properties file during installation. This file can be edited after installation. This file is on the machine where DAS is installed and is located:

For Windows:

```
%ESEC_HOME%\sentinel\config
```

For UNIX:

```
$ESEC_HOME/sentinel/config
```

There are two scripts (mailconfig.sh and mailconfigtest.sh for UNIX and mailconfig.bat and mailconfigtest.bat for Windows) that change and test the email settings within the execution.properties file. The mailconfig.* script changes the email settings and the mailconfigtest.* script tests the email settings. The bolded areas are the email settings that can be changed.

The properties within execution.properties are:

**mail.authentication.user=<domain\\user>**

correlated events retry wait=5000

**mail.smtp.host=<SMTP_HOST>** The SMTP host that will be used to send email.

mail.events.max=1000 Maximum number of events that

| | |
|---|---|
| | will be sent in an email that is automatically triggered by the correlation engine. Its purpose is to limit the size of emails for correlated events that have a very large set of trigger events. |
| correlated events retry count=10 | |
| **mail.address.from=<SMTP_FROM_ADDR>** | The email address that appears in the From field of the email sent from DAS. |
| **mail.authentication.password=<password>** | password for mail.authentication.user. |

The mailconfig.sh and mailconfig.bat scripts use the following arguments:

| | |
|---|---|
| -host | SMTP host name or IP address |
| -from | From field of the email |
| -user | The mail authentication user |
| -password | Password for the mail authentication user |

> **NOTE**: Do not enter your password after the –password argument. You will be prompted for a new password after you enter the command. The console output will be masked by asterisks (*).

The mailconfigtest.sh and mailconfig.bat file uses the following arguments:

| | |
|---|---|
| -to | Destination email address |

To set email properties in the execution.properties file

1. On the machine where you have DAS installed, cd to:
   For UNIX:
   ```
   $ESEC_HOME/sentinel/config
   ```
   For Windows
   ```
   %ESEC_HOME%\sentinel\config
   ```
2. Execute mailconfig as follows:
   For UNIX:
   ```
   ./mailconfig.sh -host <SMTP Server> -from <source
      email address> -user <mail authentication user> –
      password
   ```
   For Windows:
   ```
   mailconfig.bat -host <SMTP Server> -from <source
      email address> -user <mail authentication user> -
      password
   ```
   UNIX example:

```
./mailconfig.sh -host 10.0.1.14 -from
   my_name@domain.com -user my_user_name –password
```

Windows example:

```
mailconfig.bat -host 10.0.1.14 -from
   my_name@domain.com -user my_user_name –password
```

After entering this command you will be prompted for a new password.

```
Enter your password:*********
Confirm your password:*********
```

**NOTE**: When using the password option, it must be the last argument.

To test your email settings in the execution.properties file

1. On the machine where you have DAS installed, cd to:

   For UNIX:

   ```
   $ESEC_HOME/sentinel/config
   ```

   For Windows

   ```
   %ESEC_HOME%\sentinel\config
   ```

2. Execute mailconfigtest as follows:

   For UNIX:

   ```
   ./mailconfigtest.sh -to <destination email address>
   ```

   For Windows:

   ```
   mailconfigtest.bat -to <destination email address>
   ```

   If your mail is sent successfully, you will get the following on screen output and e-mail received at the destination address.

   ```
   Email has been sent successfully!
   ```

   Check the destination e-mail mailbox to confirm receipt of email. The subject line and content should be:

   ```
   Subject: Testing e-Security mail property
   This is a test for e-Security mail property set up.
      If you see this message, your e-Security mail
      property has been configured correctly to send
      emails
   ```

## Updating Your License Key

If your e-Security license key has expired and e-Security has issued you a new one, run the software key program to update your license key.

How to update your license key (UNIX)

1. Login as user esecadm.
2. Go to $ESEC_HOME/utilities.
3. Enter the following command:

```
./softwarekey
```

4. Enter the number 1 to set your primary key. Press enter.

### How to update your license key (Windows)

1. Login as a use with administrative rights.
2. Go to %ESEC_HOME%\utilities.
3. Enter the following command:

```
softwarekey.exe
```

4. Enter the number 1 to set your primary key. Press enter.

# Chapter 12 – Quick Start

This chapter discusses quick start procedures for:

- Security Analysts
- Report Analysts
- Administrators

The following topics are discussed:

- Active Views™
- Exploit Detection
- Asset Data
- Event Query
- Analysis Reporting via Crystal Reports
- Basic Correlation

Below are links to video demos to some of the functionality with Sentinel 5.

- Active Views™ - video demo with Asset data and correlation
- Active Views™ - video demo with Quick Query
- Active Views™ - video demo with refining the event table (real-time filtering)
- Vulnerability – video demo
- Analysis Reporting – video demo (Crystal 9)
- iTRAC™ – video demo, assigning an iTRAC role
- iTRAC™ – video demo, iTRAC Process Monitor

## Security Analysts

> **NOTE**: Assumption, your Security Administrator or you have built the necessary filters and configured the necessary agents for your system.

### Active Views Tab

In the Active Views tab, you can monitor events as they happen, performing queries on these events. You can monitor them in a table form or through a 3-D graphical representation.

To get a Real-Time events started

1. Click Active Views > Create an Active View, click the Filter down arrow, select a filter and click Select.

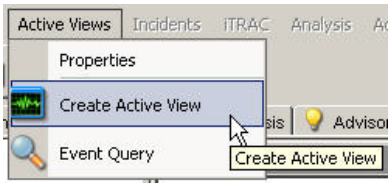

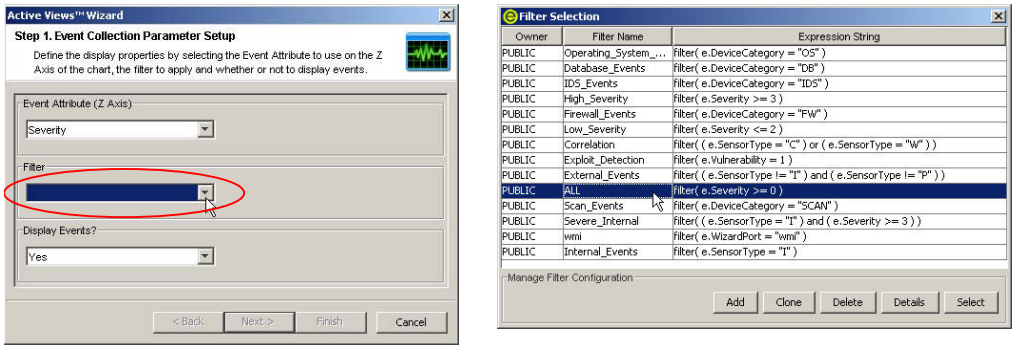


2. Click Finish. If you have an active network, you may see something similar to:

**NOTE**: To display a 3-D graph without real time events, click the Display Events down arrow and select No.

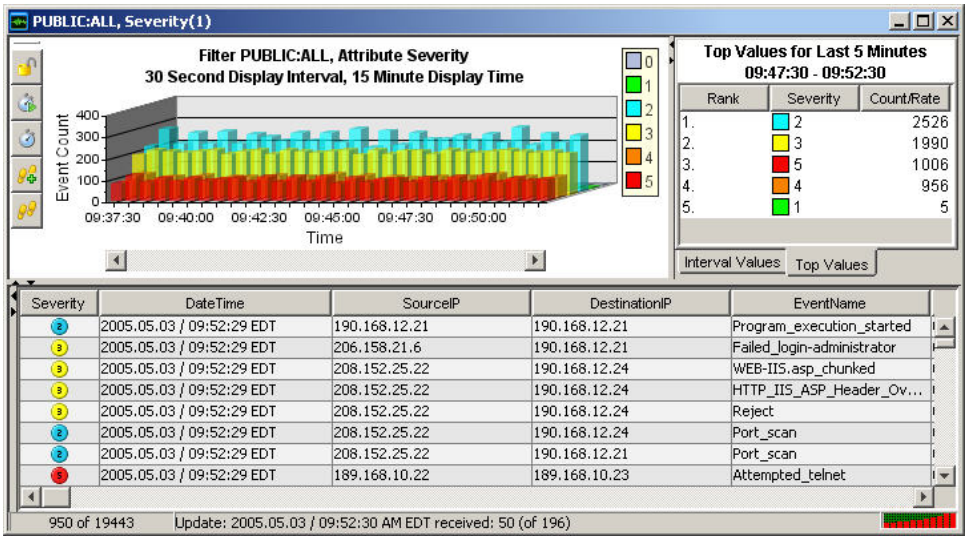


## Exploit Detection

To view any events indicating a possible exploitation, you must have the following:

- Advisor Feed
- Intrusion detection
- Vulnerability scanning

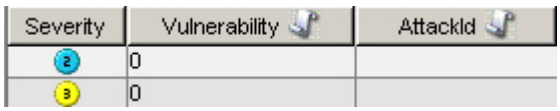


Within an event, when the Vulnerability field (*vul*) equals 1, the asset or destination device is exploited. If the vulnerability field equals 0, the asset or destination device is not exploited. If the Vulnerability field is blank, the exploit detection feature of Sentinel is not active.

To view events that indicate a possible exploitation, create an Active View with a filter where Vulnerability equals 1. If you have Nmap and have run the Nmap agent, you can view asset information on the exploited asset or any asset.

For more information on how exploit detection works and which Intrusion Detection Systems and Vulnerability Scanners are supported, see Chapter 1 – Introduction or Chapter 10 – Sentinel Data Manager.

## Asset Data

To view Asset information for any event, right-click on an event or events > Analysis > Asset Data, a window similar to the one below will appear.

**Asset Report**

**desk.acmeinc.net**

**Hardware**

| | | | |
|---|---|---|---|
| **MAC Address** | A0:12:56:78:90:00 | | |
| **Name** | Build Machine | **Value** | 500 |
| **Type** | Server | **Criticality** | High |
| **Vendor** | Dell | **Sensitivity** | Low |
| **Product** | Precision | **Environment** | Production |
| **Version** | 360 | **Location** | Internal |

**Network**

| IP | Hostname |
|---|---|
| 199.16.2.23 | desk.acmeinc.net |

**Software**

| Name | Type | Vendor | Product | Version |
|---|---|---|---|---|
| ClearCase | APPLICATION | IBM | ClearCase | 5.0 |
| C++ | APPLICATION | Microsoft | Visual C++ | 6.0 |

**Contacts**

| Order | Name | Role | Email | Phone Number |
|---|---|---|---|---|
| 1 | Erickson, Stein | USER | serickson@acmedomain.net | (703) 555-8865 |
| 2 | IT | Administrator | LAN_FOLKS@acemedomain.net | (703) 555-9876 |

**Location**

| | |
|---|---|
| **Room** | server room |
| **Rack** | #17 |
| **Address** | HQ<br>Agent 86 Security Circle<br>Suite 86<br>Washington DC 12345 USA |

## Event Query

Example scenario - During monitoring, you see numerous telnet attempts from source IP 189.168.10.22. Telnet attempts could be an attack. Telnet potentially allows an attacker to remotely connect to a remote computer as if they were locally connected. This can lead to unauthorized configuration changes, installation of programs, viruses, etc.

You can Event Query to determine how often this possible attacker has attempted a telnet, you can setup a filter to query for this particular attacker. For example, you know the following:

- Source IP: 189.168.10.22
- Destination IP: 189.168.10.23
- Severity: 5
- Event Name: Attempted_telnet
- Sensor Type: H (Host Intrusion Detection)

### To Perform an Event Query

1. Click the Event Query button (magnifying glass icon) and click the Filter field down arrow.
2. Click Add, enter a filter name of "telnet SIP 189_168_10_22". In the field below the Filter, enter:
   - SourceIP = 189.168.10.22
   - EventName = Attempted_telnet
   - Severity = 5
   - SensorType = H

- Match if, select (and)
- DestinationIP = 189.168.10.23

3. Click Save. Highlight your filter and click Select.
4. Enter your time period of interest, click the search button (magnifying glass icon). The results of your query will appear.

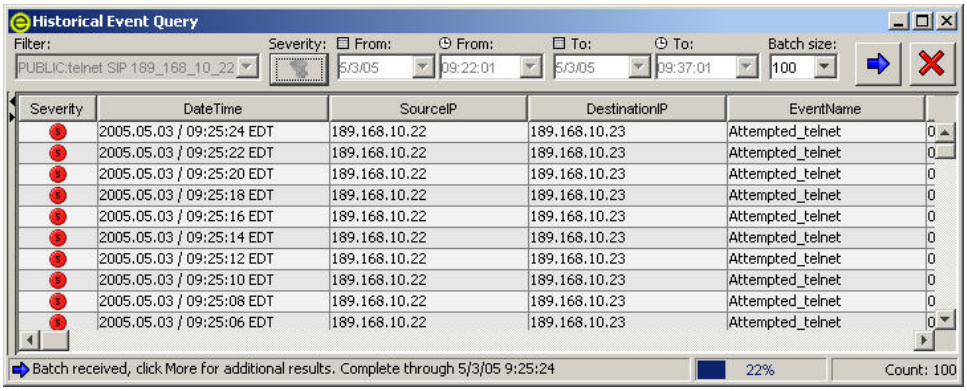


If you want to see how often in general this user is attempting telneting, remove DestinationIP, SensorType and Severity from your filter or create a new filter. The results will show all the destinationIPs this user is attempting to telnet to.

If any of your events are correlated events (SensorType = C or W), you can right-click > View Trigger Events to find what events triggered that correlated event.

Another event of interest could be excessive FTP events. This can also be a remote connection, allowing for transferring, copying and deleting of files.

Below is a short list of attacks of interest. Types of attacks is an extensive list. For more information about network/host attacks, there are many resources available (i.e., books and the internet) that explain different types of attacks in detail.

- SYN Flood
- ICMP and UDP Flood
- Packet Sniffing
- Denial of Service
- Smurf and Fraggle
- Dictionary Attack

## Report Analyst

> **NOTE**: Assumption, your Security Administrator has configured your Crystal Enterprise web server and published a list of available reports.

### Analysis Tab

The Analysis tab allows for historical reporting. Historical and vulnerability reports are published on a Crystal web server, these run directly against the e-Security database. These reports can be useful to track and investigate activity over a large time frame, for instance a week or a month. These reports can also be used as a high level reporting method to your supervisors. If your reporting web server is installed, look in the navigator bar to see what reports are available.

> **NOTE**: The following is a Crystal 9 example. Crystal 11 procedures are the same with different report names.

For example, if you are responsible for generating reports to upper management within your organization. Chances are you will run SourceDestinationReports.

These are Top 10 Source to Destination IP Pairs on hosts names, ports, IPs and users. To run this report, do the following:

### Running a Crystal Report

1. Expand Top 10 and highlight Top 10 Source to Destination IP Pairs Summary and click Create Reports button (magnifying glass).
2. Enter esecrpt (for SQL authentication and Oracle) as the username or your Windows Authentication username and enter your password.
3. Under Report Type, select Weekly Report (select Specific Date Range if you want a specific date range).

> **NOTE**: Other reports may have additional parameters such as resource name and severity range.

4. Click View Report.

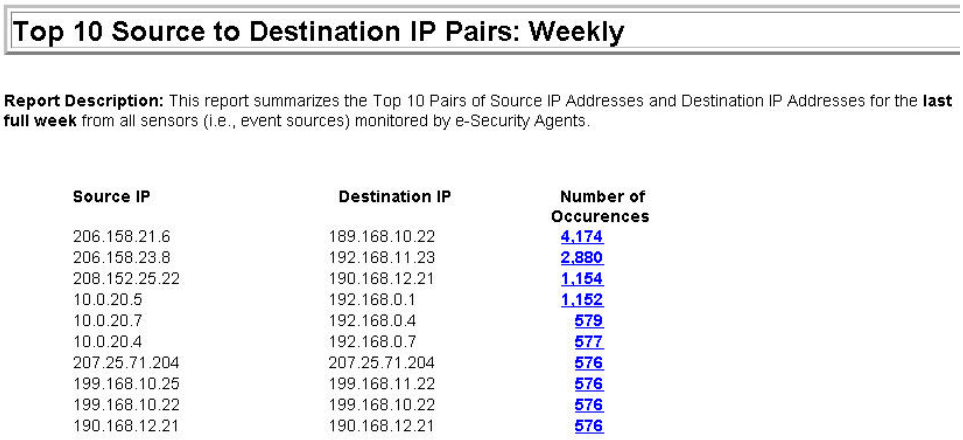

**Top 10 Source to Destination IP Pairs: Weekly**

**Report Description:** This report summarizes the Top 10 Pairs of Source IP Addresses and Destination IP Addresses for the **last full week** from all sensors (i.e., event sources) monitored by e-Security Agents.

| Source IP | Destination IP | Number of Occurences |
|---|---|---|
| 206.158.21.6 | 189.168.10.22 | 4,174 |
| 206.158.23.8 | 192.168.11.23 | 2,880 |
| 208.152.25.22 | 190.168.12.21 | 1,154 |
| 10.0.20.5 | 192.168.0.1 | 1,152 |
| 10.0.20.7 | 192.168.0.4 | 579 |
| 10.0.20.4 | 192.168.0.7 | 577 |
| 207.25.71.204 | 207.25.71.204 | 576 |
| 199.168.10.25 | 199.168.11.22 | 576 |
| 199.168.10.22 | 199.168.10.22 | 576 |
| 190.168.12.21 | 190.168.12.21 | 576 |

5. You can export this file as a Word, PDF, rtf, Excel or as a Crystal Report by clicking the Export Button (envelope).

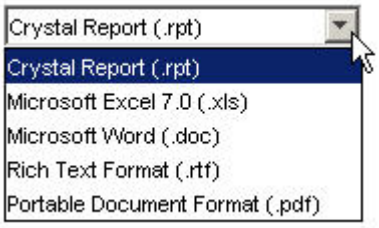


## Event Query

Similar to the Security Analyst, if you have an event or events of interest within your reports, you can run an Event Query under the Analysis tab. To run a query, highlight Historical Events > Historical Event Queries and click the Create Reports button (magnifying glass). For more information, see Security Analyst - Event Query Sample Scenario.

# Administrators

## Basic Correlation

Correlation is the process of analyzing security events to identify potential relationships between two or more events. Correlation allows quick association of priority attacks based on common elements of event data.

In reference to the telnet scenario under Security Analyst - Event Query Sample Scenario, a Basic Correlation Rule can be created that will trigger a correlated event when 4 telnet attempts are done in a 10 second period.

### To Create a Correlation Rule

1. Go to the Admin tab and highlight Correlation Rules in the navigation bar.
2. Create a new folder and place your rule in it. This done through a right-click option.
3. Highlight Basic Correlation, enter a name and click Next. In the next pane, click the down arrow and select Filter Manager. Click the Selected Filter down arrow and in the Filter Selection pane, click Add.
4. Enter the following:
   - Name: telnet_attempt_189_168_10_22
   - Filter Name: telnet attempt 189_168_10_22
   - SourceIP = 189.168.10.22
   - EventName = Attempted_telnet
   - select And
   - Severity = 5
   - SensorType = H
   - DestinationIP = 189.168.10.23
5. Click Save. Highlight your filter and click Select.
6. Click Next, enter the value of 4 for when condition is met and 10 seconds in the Threshold Grouping Criteria pane. Click Next.
7. In the Correlated Events and Actions pane, change the severity level to 2 (click the down arrow). Click Finish.
8. To deploy this rule, highlight Correlation Engine Manager in the Navigation pane, highlight a correlation engine, right-click > Deploy Rules. In the Deploy rules pane, find your rule and check mark it. Click OK. Ensure that your Correlation Engine and Correlation Rule have a green check marks indicating that they are enabled. This is done by right-clicking.
9. There are several different methods to view if you have correlated events. Some methods are:
   - Create an Active View Events window using the correlation filter you created
   - Create an Active View Events window using the provided correlation filter

- Create an Active View Events window using the provided All filter, take a snap shot and sort by SensorType and view all events with SensorType equal to C.
- Quick Query using the filter you created or using the correlation filter.

Right-click on the correlated event and select View Trigger Events to see how many telnet events (could be more than 4) triggered this correlation rule.

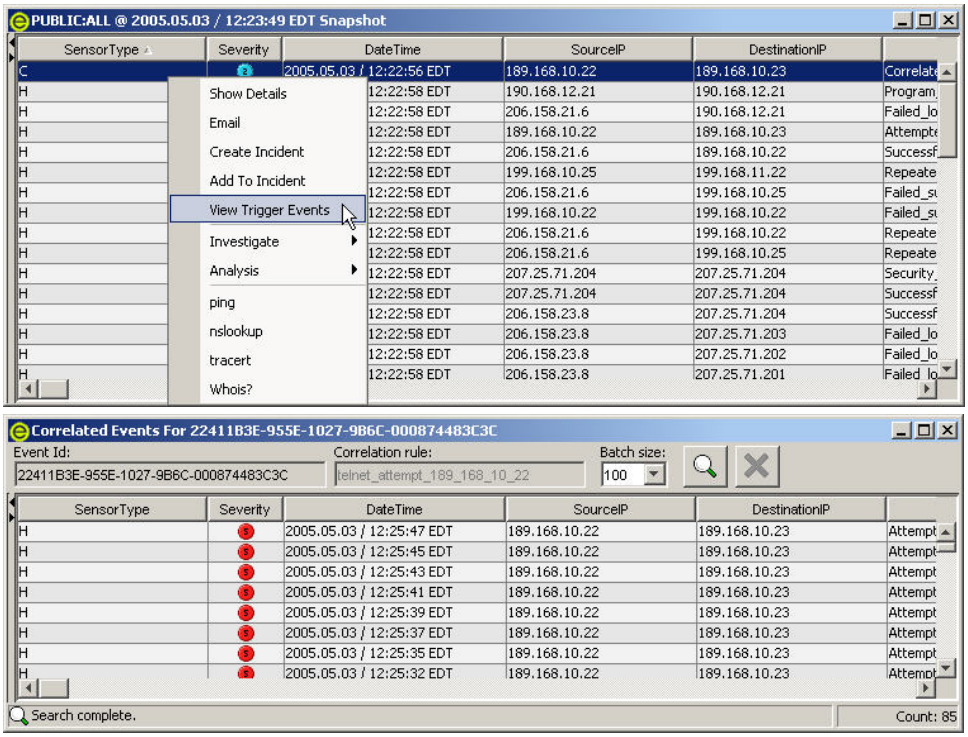

# Appendix A – System Events for Sentinel 5

In the description tables below, words in italics surrounded by <…> are replaced by relevant values in the real messages.

## Authentication Events

### Failed Authentication

When a user authentication fails, the following event is generated.

| Tag | Value |
|---|---|
| Severity | 4 |
| Event Name | AuthenticationFailed |
| Resource | UserAuthentication |
| SubResource | Authenticate |
| Message | Authentication of user <*name*> with OS name <*domUser*> from <*IP*> failed |

### No Such User Event

When a user attempts to login into the application and authentication succeeds but the user is not an e-Security user, the following event is generated.

| Tag | Value |
|---|---|
| Severity | 4 |
| Event Name | NoSuchUser |
| Resource | UserAuthentication |
| SubResource | Authenticate |
| Message | No existing user with name <*name*> found |

### Duplicate User Objects

When there is an unexpected second active user object, this should not happen, the following event is generated. This is an internal error.

| Tag | Value |
|---|---|
| Severity | 4 |
| Event Name | TooManyActiveUsers |
| Resource | UserAuthentication |
| SubResource | Authenticate |
| Message | Error in user table : Multiple users with the name <*name*> found |

### Locked Account

When a locked user account is attempting to login, the following event is generated.

| Tag | Value |
|---|---|
| Severity | 4 |
| Event Name | LockedUser |
| Resource | UserAuthentication |
| SubResource | Authentication |

| Tag | Value |
|---|---|
| Message | Attempt to login using locked account <*acct*> |

## User Sessions

### User Logged Out

When a user logs out, the following internal event is generated.

| Tag | Value |
|---|---|
| Severity | 1 |
| Event Name | UserLoggedOut |
| Resource | UserSessionManager |
| SubResource | User |
| Message | Closing session for <*user*> OS name <*osName*> from <*IP*> was on since <*date*>; currently <*num*> active users |

### User Logged In

When a user logs in, the following internal event is generated.

| Tag | Value |
|---|---|
| Severity | 1 |
| Event Name | UserLoggedIn |
| Resource | UserSessionManager |
| SubResource | User |
| Message | User <*user*> with OS name <*osName*> at <*IP*> logged in; currently <*num*> active users |

### User Discovered

If the server restarts, it loses the session information. It will then reconstruct the session when it receives messages from active users. When it discovers a connected user, the following internal event is generated.

| Tag | Value |
|---|---|
| Severity | 1 |
| Event Name | UserLoggedIn |
| Resource | UserSessionManager |
| SubResource | User |
| Message | Discovered active user <*user*> with OS name <*osName*> at <*IP*> logged in; currently <*num*> active users |

## Event

### Error Moving Completed File

When an event file is completed it is moved to the output directory. If that move fails the following internal event is generated.

| Tag | Value |
|---|---|
| Severity | 3 |
| Event Name | MoveArchiveFileFailed |
| Resource | <*DAS name*> |
| SubResource | ArchiveFile |
| Message | Error moving completed archive file <*fname*> to <*dir*> |

### Error inserting events

When inserting events into the database fails the following internal event is generated.

| Tag | Value |
|---|---|
| Severity | 5 |
| Event Name | InsertEventsFailed |
| Resource | EventSubsystem |
| SubResource | Events |
| Message | Error inserting events into the Database—the events may be permanently lost. Please check the Database and backend server logs <*Exception*> |

### Opening Archive File failed

When opening an archive file for storing the events for aggregation fails, the following internal event is generated.

| Tag | Value |
|---|---|
| Severity | 3 |
| Event Name | OpenArchiveFileFailed |
| Resource | <*Das name*> |
| SubResource | ArchiveFile |
| Message | Error opening archive file <*name*> in <*dir*> |

### Writing to Archive File failed

When opening an archive file for storing the events for aggregation fails, the following internal event is generated.

| Tag | Value |
|---|---|
| Severity | 3 |
| Event Name | WriteArchiveFileFailed |
| Resource | <*Das name*> |
| SubResource | ArchiveFile |
| Message | Error writing newly received events to aggregation archive file <*fname*> |

### Writing to the overflow partition (P_MAX)

An event is sent approximately every 5 minutes notifying the user when events are being written to the overflow partition (P_MAX). When this occurs, the administrator needs to use SDM and add more partitions otherwise performance will start degrading.

| Tag | Value |
|---|---|
| Severity | 5 |
| Event Name | InsertIntoOverflowPartition |
| Resource | EventSubSystem |
| SubResource | Events |
| Message | Error: currently inserting into the overflow partitions (P_MAX), add more partitions |

### Event Insertion is blocked

If DAS is writing into the overflow partition and the user attempts to add partitions SDM will send a request to DAS to temporarily stop inserting events into the database. When this happens DAS will send internal events every time it attempts to insert events into the database.

| Tag | Value |
|---|---|
| Severity | 4 |
| Event Name | EventInsertionIsBlocked |
| Resource | EventSubSystem |
| SubResource | Events |
| Message | Event insertion is blocked, waiting <*num*> sec |

### Event Insertion is resumed

When event insertion is resumed after being blocked, the following event is sent.

| Tag | Value |
|---|---|
| Severity | 2 |
| Event Name | EventInsertionResumed |
| Resource | EventSubSystem |
| SubResource | Events |
| Message | Event insertion has resumed after being blocked |

### Database Space Reached Specified Time Threshold

When event insertion is resumed after being blocked, the following event is sent.

| Tag | Value |
|---|---|
| Severity | 0 |
| Event Name | DbSpaceReachedTimeThrshld |
| Resource | Database |
| SubResource | Database |
| Message | Tablespace <*string*> has <*num*> MB left and growing <*num*> bytes per second and will run out space within the time threshold specified <*num*> seconds |

### Database Space Reached Specified Percent Threshold

When event insertion is resumed after being blocked, the following event is sent.

| Tag | Value |
|---|---|
| Severity | 0 |
| Event Name | DbSpaceReachedPercentThrshld |
| Resource | Database |
| SubResource | Database |
| Message | Tablespace <*string*> has current size of <*num*> MB with a max size of <*num*> MB and has reached the percentage threshold of <*num*> % |

### Database Space Very Low

When event insertion is resumed after being blocked, the following event is sent.

| Tag | Value |
|---|---|
| Severity | 5 |
| Event Name | DbSpaceVeryLow |
| Resource | Database |
| SubResource | Database |
| Message | Tablespace <*string*> has current size of <*num*> MB and has reached the physical threshold of <*num*> MB |

## Aggregation

### Error inserting summary data into the database

If an error is encountered while writing aggregation data into the database, the following internal event is generated.

| Tag | Value |
|---|---|
| Severity | 4 |
| Event Name | SummaryUpdateFailure |
| Resource | Aggregation |
| SubResource | Summary |
| Message | Error saving summary batch to the database for summary <*summaryName*> |

## Mapping Service

### Error initializing map with ID

This internal event is generated from the client side of the mapping service (the one that is part of the agent manager). This error is generated when the agent manager attempts to retrieve a map that does not exist. This should not happen but may happen if maps are created and deleted.

| Tag | Value |
|---|---|
| Severity | 4 |
| Event Name | ErrorNoSuchMap |
| Resource | MappingService |
| SubResource | ReferentialDataObjectMap |

| Tag | Value |
|---|---|
| Message | Error initializing map with id <ID>: no such map |

### Refreshing Map from Cache

This internal event is generated from the client side of the mapping service (the one that is part of the agent manager). When the agent manager is told to refresh the map because it has been modified or its definition has changed it sends an internal event. This means that its cache is up to date and is refreshing the map from cache.

| Tag | Value |
|---|---|
| Severity | 1 |
| Event Name | LoadingMapFromCache |
| Resource | MappingService |
| SubResource | ReferentialDataObjectMap |
| Message | Loading from cache v<*version*> of map <*mapName*> (ID <*id*>) |

### Refreshing Map from Server

This internal event is generated from the client side of the mapping service (the one that is part of the agent manager). When the agent manager is told to refresh the map because it has been modified or its definition has changed it sends an internal event. This means that the map was either not in the cache or the version in the cache was not up to date and the agent manager is retrieving the map from the server.

| Tag | Value |
|---|---|
| Severity | 1 |
| Event Name | RefreshingMapFromServer |
| Resource | MappingService |
| SubResource | ReferentialDataObjectMap |
| Message | Refreshing from server map <name> with id <ID> |

### Timeout Refreshing Map

This internal event is generated from the client side of the mapping service (the one that is part of the agent manager). When the agent manager is told to refresh the map because it has been modified or its definition has changed it sends an internal. This means that the agent manager attempted to retrieve the map from the server and the server never acknowledged the request and timed out. This error is considered transient and the agent manager will retry.

| Tag | Value |
|---|---|
| Severity | 4 |
| Event Name | TimeoutRefreshingMap |
| Resource | MappingService |
| SubResource | ReferentialDataObjectMap |
| Message | Request timed out while refreshing map <name>: <exception> |

### Error Refreshing Map

This internal event is generated from the client side of the mapping service (the one that is part of the agent manager). When the agent manager is told to refresh the map because it has been modified or its definition has changed it sends an internal event. This means that there was some unexpected non-transient error while trying to refresh a map. The agent manager will wait 15 minutes and will try again. If this happens during initialization the initialization will proceed and this map will be ignored until it can be successfully loaded.

| Tag | Value |
|---|---|
| Severity | 4 |
| Event Name | ErrorRefreshingMapData |
| Resource | MappingService |
| SubResource | ReferentialDataObjectMap |
| Message | Error refreshimg map <*mapName*>: <*exc*> |

### Loaded Large Map

This internal event is an information event sent by the mapping service informing that a large map was loaded to the agent manager. A map is considered large if the number of rows exceeds 100,000.

| Tag | Value |
|---|---|
| Severity | 0 |
| Event Name | LoadedLargeMap |
| Resource | MappingService |
| SubResource | ReferentialDataObjectMap |
| Message | Finished loading map <name> with id <ID> and <num> entries and total size <#>Kb in <##>sec |

### Long time to load Map

This internal event is an information event sent by the mapping service informing that loading a map took an unusually long time (greater than one minute).

| Tag | Value |
|---|---|
| Severity | 0 |
| Event Name | LongTimeToLoadMap |
| Resource | MappingService |
| SubResource | ReferentialDataObjectMap |
| Message | It took <##>sec to load map <name> with id <ID> and <num> entries and total size <##>Kb |

### TimedoutWaitingForCallback

When the agent manager needs to refresh a map it sends a request to the backend. This request contains a callback. The backend generates the map and when it is ready it sends the map to the agent manager using the callback. If it takes too long for the response to arrive (more than ten minutes) the agent manager will submit a second request assuming the first was lost. When this occurs, the following internal event is genrerated.

| Tag | Value |
|---|---|
| Severity | 2 |
| Event Name | TimedoutWaitingForCallback |
| Resource | MappingService |
| SubResource | ReferentialDataObjectMap |
| Message | Map <name> timed out waiting for callback with new map data--retrying |

### ErrorApplyingIncrementalUpdate

This event is sent when the mapping service fails to apply an update to an existing client map.

| Tag | Value |
|---|---|
| Severity | 4 |
| Event Name | ErrorApplyingIncrementalUpdate |
| Resource | MappingService |
| SubResource | ReferentialDataObjectMap |
| Message | The error <error> occurred while applying updates to map <*mapName*> (ID <*mapId*>) v.<*version*>. Rescheduling a refresh to complete map update. |

### OutOfSyncDetected

This event is sent when the mapping service detects that a map is out of date. The mapping service will automatically schedule a refresh.

| Tag | Value |
|---|---|
| Severity | 2 |
| Event Name | OutOfsyncDetected |
| Resource | MappingService |
| SubResource | ReferentialDataObjectMap |
| Message | Map <*mapName*> detected the map data is out-of-sync, probably due to a missed update notification--scheduling a refresh |

## Event Router

### Event Router is Running

Event router is the main component of the agent manager (the one that performs the maps, applies global filters and publishes the events). This internal event is sent when the event router is ready during initialization. When the agent manager is restarted, another event will be sent when it is ready.

This event is not sent until the event router successfully loaded all the global filters and map information.

| Tag | Value |
|---|---|
| Severity | 1 |
| Event Name | EventRouterIsRunning |
| Resource | AgentManager |
| SubResource | EventRouter |

| Tag | Value |
|---|---|
| Message | Event router completed its initialization in <mode> mode |

### Event Router is Initializing

This event is sent when an event router starts its initialization. The event router starts initializing when it has established a connection with the backend (DAS Query).

| Tag | Value |
|---|---|
| Severity | 1 |
| Event Name | EventRouterInitializing |
| Resource | AgentManager |
| SubResource | EventRouter |
| Message | Event router is initializing in <mode> mode |

### Event Router is Stopping

This event is sent when a request is received by the event router to stop during shutdown.

| Tag | Value |
|---|---|
| Severity | 2 |
| Event Name | EventRouterStopping |
| Resource | AgentManager |
| SubResource | EventRouter |
| Message | Event router is stopping |

### Event Router is Terminating

This event is sent when a request is received by the event router to stop during shutdown.

| Tag | Value |
|---|---|
| Severity | 2 |
| Event Name | EventRouterTerminating |
| Resource | AgentManager |
| SubResource | EventRouter |
| Message | Event router is terminating |

## Correlation Engine

### Correlation Engine is Running

The correlation engine process can be idled by the user. Its running state determines whether the active process is processing events or not. The process starts in the idle (stopped) state and waits to retrieve its configuration from the database. This event is sent when the engine changes state from stopped to running.

| Tag | Value |
|---|---|
| Severity | 1 |
| Event Name | EngineRunning |

| Tag | Value |
|---|---|
| Resource | CorrelationEngine |
| SubResource | CorrelationEngine |
| Message | Correlation Engine is processing events. |

### Correlation Engine is Stopped

This event is sent out when the engine changes state from running to stopped.

| Tag | Value |
|---|---|
| Severity | 1 |
| Event Name | EngineStopped |
| Resource | CorrelationEngine |
| SubResource | CorrelationEngine |
| Message | Correlation Engine has stopped processing events. |

### Rule Deployment is Started

This event is sent out when an engine successfully loads a rule deployment. This message is sent out regardless of the engine running state.

| Tag | Value |
|---|---|
| Severity | 1 |
| Event Name | DeploymentStarted |
| Resource | CorrelationEngine |
| SubResource | Deployment |
| Message | deployment <name> started |

### Rule Deployment is Stopped

This event is sent out when an engine successfully unloads a rule deployment. This message is sent out regardless of the engine running state.

| Tag | Value |
|---|---|
| Severity | 1 |
| Event Name | DeploymentStopped |
| Resource | CorrelationEngine |
| SubResource | Deployment |
| Message | deployment  <name> stopped |

### Rule Deployment is Modified

This event is sent out when an engine successfully reloads a rule deployment. This message is sent out regardless of the engine running state.

| Tag | Value |
|---|---|
| Severity | 1 |
| Event Name | DeploymentModified |
| Resource | CorrelationEngine |
| SubResource | Deployment |
| Message | Deployment <name> modified |

# WatchDog

### Controlled Process is started

Watchdog is run as a service. Its main purpose is to keep Sentinel processes running. If a process dies, Watchdog will automatically restart that process. This event is sent out when a process is started.

| Tag | Value |
|---|---|
| Severity | 1 |
| Event Name | ProcessStart |
| Resource | WatchDog |
| SubResource | Process |
| Message | Process <ProgramName> spawned (<pid>) |

### Controlled Process is stopped

This event is sent out when a process is stopped. The severity is set to 5 if the process was set to respawn (i.e. it is not expected to die). The severity is set to 1 if the process was set to run once.

| Tag | Value |
|---|---|
| Severity | 1/5 |
| Event Name | ProcessStop |
| Resource | WatchDog |
| SubResource | Process |
| Message | Process <ProgramName> exited with code <exit_code> |

### Watchdog Process is started

As the Watchdog process starts, the following internal event is generated.

| Tag | Value |
|---|---|
| Severity | 1 |
| Event Name | ProcessStart |
| Resource | WatchDog |
| SubResource | WatchDog |
| Message | WatchDog Service Starting |

### Watchdog Process is stopped

When the Watchdog service is stopped, the following internal event is generated.

| Tag | Value |
|---|---|
| Severity | 5 |
| Event Name | ProcessStop |
| Resource | WatchDog |
| SubResource | WatchDog |
| Message | WatchDog Service Ended |

# Agent Engine/Manager

### Port Start

Agent Manager sends this event when a port is started.

| Tag | Value |
|---|---|
| Severity | 1 |
| Event Name | PortStart |
| Resource | AgentManager |
| SubResource | AgentManager |
| Message | Processing started for port_<port id> |

### Port Stop

Agent Manager sends this event when a port is stopped.

| Tag | Value |
|---|---|
| Severity | 1 |
| Event Name | PortStop |
| Resource | AgentManager |
| SubResource | AgentManager |
| Message | Processing stopped for port_<port id> |

### Persistent Process Died

Agent Engine sends this event when the persistent process connector detects its controlled process has died.

| Tag | Value |
|---|---|
| Severity | 5 |
| Event Name | PersistentProcessDied |
| Resource | AgentManager |
| SubResource | AgentManager |
| Message | Persistent Process on port <port id> has died. |

### Persistent Process Restarted

Agent Engine sends this event when the persistent process connector is able to restart the controlled process that had died.

| Tag | Value |
|---|---|
| Severity | 1 |
| Event Name | PersistentProcessRestarted |
| Resource | AgentManager |
| SubResource | AgentManager |
| Message | Persistent Process on port <port id> has restarted. |

# Event Service

### Cyclical Dependency

Event Service sends this event when it detects a cycle in the Event Definition (in dependencies among tags due to referential map assignments). Check the event configuration in SDM and resolve the dependency.

| Tag | Value |
|---|---|
| Severity | 5 |
| Event Name | CyclicalDependency |
| Resource | EventService |
| SubResource | ObjectAttrInfos |
| Message | Cyclical dependency detected in event transformations. Check event configuration. |

# Active Views

### Active View Created

DAS_Binary sends this event when an Active View is created.

| Tag | Value |
|---|---|
| Severity | 1 |
| Event Name | RtChartCreated |
| Resource | RealTimeSummaryService |
| SubResource | ChartManager |
| Message | Creating new Active View with filter <filter> and attribute <attribute> for users with security filter <security filter>. Currently <n> Active View(s) Collecting. |

### Active View Joined

DAS_Binary sends this event when a user connects to an existing Active View.

| Tag | Value |
|---|---|
| Severity | 1 |
| Event Name | RtChartJoiningExistingData |
| Resource | RealTimeSummaryService |
| SubResource | ChartManager |
| Message | Joining existing Active View with filter <filter> and attribute <attribute> for users with security filter <security filter>. Currently <n> Active View(s) Collecting. |

### Idle Active View Removed

DAS_Binary sends this event when a non-permanent Active View is removed due to inactivity.

| Tag | Value |
|---|---|
| Severity | 1 |

| Tag | Value |
|---|---|
| Event Name | RtChartInactiveAndRemoved |
| Resource | RealTimeSummaryService |
| SubResource | ChartManager |
| Message | Removed idle Active View with filter <filter> and attribute <attribute> for users with security filter <security filter>. Currently <n> Active View(s) Collecting. |

### Idle Permanent Active View Removed

DAS_Binary sends this event when a permanent Active View is removed due to inactivity. Permanent Active Views are ones saved in user preferences and timeout after several days of inactivity by default.

| Tag | Value |
|---|---|
| Severity | 1 |
| Event Name | RtPermanentChartRemoved |
| Resource | RealTimeSummaryService |
| SubResource | ChartManager |
| Message | Removed idle permanent Active View with filter <filter> and attribute <attribute> for users with security filter <security filter>. Currently <n> Active View(s) Collecting. |

### Active View Now Permanent

DAS_Binary sends this event when it detects an Active View as newly permanent. This check happens periodically, so it may be several minutes after an Active View is saved to preferences before this event is generated.

| Tag | Value |
|---|---|
| Severity | 1 |
| Event Name | RtChartIsNowPermanent |
| Resource | RealTimeSummaryService |
| SubResource | ChartManager |
| Message | Active View with filter <filter> and attribute <attribute> for users with security filter <security filter> is now permanent. |

### Active View No Longer Permanent

DAS_Binary sends this event when it detects a formerly permanent Active View that is no longer permanent. This check happens periodically, so it may be several minutes after an Active View is removed from preferences before this event is generated.

| Tag | Value |
|---|---|
| Severity | 1 |
| Event Name | RtChartNotPermanent |
| Resource | RealTimeSummaryService |
| SubResource | ChartManager |

| Tag | Value |
| --- | --- |
| Message | Active View with filter <filter> and attribute <attribute> for users with security filter <security filter> is no longer permanent. |

| Event Name | Severity | Source | SubResource | Component |
|---|---|---|---|---|
| AuthenticationFailed | 4 | UserAuthentication | Authenticate | Authentication |
| NoSuchUser | 4 | UserAuthentication | Authenticate | Authentication |
| TooManyActiveUsers | 4 | UserAuthentication | Authenticate | Authentication |
| LockedUser | 4 | UserAuthentication | Authenticate | Authentication |
| UserLoggedOut | 1 | UserSessionManager | User | User Session |
| UserLoggedIn | 1 | UserSessionManager | User | User |
| UserLoggedIn | 1 | UserSessionManager | User | User |
| MoveArchiveFileFailed | 3 | *DAS Name* | ArchiveFile | Event |
| InsertEventsFailed | 5 | EventSubSystem | Events | Event |
| OpenArchiveFileFailed | 3 | *DAS Name* | ArchiveFile | Event |
| WriteArchiveFileFailed | 3 | *DAS Name* | ArchiveFile | Event |
| SummaryUpdateFailure | 4 | Aggregation | Summary | Aggregation |
| InsertIntoOverflowPartition | 5 | EventSubSystem | Events | Event |
| EventInsertionIsBlocked | 4 | EventSubSystem | Events | Event |
| EventInsertionResumed | 2 | EventSubSystem | Events | Event |
| EventRouterIsRunning | 1 | AgentManager | EventRouter | EventRouter |
| EventRouterInitializing | 1 | AgentManager | EventRouter | EventRouter |
| EventRouterStopping | 2 | AgentManager | EventRouter | EventRouter |
| EventRouterTerminating | 2 | AgentManager | EventRouter | EventRouter |
| ErrorNoSuchMap | 4 | MappingService | ReferentialDataObjectMap | Mapping |
| LoadingMapFromCache | 1 | MappingService | ReferentialDataObjectMap | Mapping |
| RefreshingMapFromServer | 1 | MappingService | ReferentialDataObjectMap | Mapping |
| TimeoutRefreshingMapData | 4 | MappingService | ReferentialDataObjectMap | Mapping |
| ErrorRefreshingMapData | 4 | MappingService | ReferentialDataObjectMap | Mapping |
| LoadedLargeMap | 0 | MappingService | ReferentialDataObjectMap | Mapping |
| LongTimeToLoadMap | 0 | MappingService | ReferentialDataObjectMap | Mapping |
| TimedoutWaitingForCallback | 2 | MappingService | ReferentialDataObjectMap | Mapping |
| ErrorApplyingIncrementalUpdate | 4 | MappingService | ReferentialDataObjectMap | Mapping |
| OutOfSyncDetected | 2 | MappingService | ReferentialDataObjectMap | Mapping |

| Event Name | Severity | Source | SubResource | Component |
|---|---|---|---|---|
| EngineRunning | 1 | CorrelationEngine | CorrelationEngine | |
| EngineStopped | 1 | CorrelationEngine | CorrelationEngine | |
| DeploymentStarted | 1 | CorrelationEngine | Deployment | |
| DeploymentStopped | 1 | CorrelationEngine | Deployment | |
| DeploymentModified | 1 | CorrelationEngine | Deployment | |
| ProcessStart | 1 | WatchDog | Process | |
| ProcessStop | 1/5 | WatchDog | Process | |
| ProcessStart | 1 | WatchDog | WatchDog | |
| ProcessStop | 5 | WatchDog | WatchDog | |
| PortStart | | AgentManager | AgentManager | |
| PortStop | | AgentManager | AgentManager | |
| PersistentProcessDied | 5 | AgentManager | AgentManager | |
| PersistentProcessRestarted | 1 | AgentManager | AgentManager | |
| SortDependencies | 5 | EventService | ObjectAttrInfo | EventService |
| DbSpaceReachedTimeThrshld | 0 | Database | Database | Event |
| DbSpaceReachedPercentThrshld | 0 | Database | Database | Event |
| DbSpaceVeryLow | 5 | Database | Database | Event |
| RtChartCreated | 1 | RealTimeSummaryService | ChartManager | Active Views |
| RtChartJoiningExistingData | 1 | RealTimeSummaryService | ChartManager | Active Views |
| RtChartInactiveAndRemoved | 1 | RealTimeSummaryService | ChartManager | Active Views |
| RtChartPermanentAndRemoved | 1 | RealTimeSummaryService | ChartManager | Active Views |
| RtChartIsNowPermanent | 1 | RealTimeSummaryService | ChartManager | Active Views |
| RtChartNotPermanent | 1 | RealTimeSummaryService | ChartManager | Active Views |

# Appendix B – Sentinel Copyright Information

## e-Security Sentinel™ 5



Sentinel 5 may contain the following third-party technologies:

- **Apache Axis** and **Apache Tomcat**, Copyright © 1999 to 2005, Apache Software Foundation. For more information, disclaimers and restrictions, see http://www.apache.org/licenses/
- **ANTLR.** For more information, disclaimers and restrictions, see http://www.antlr.org
- **Boost**, Copyright © 1999, Boost.org.
- **Bouncy Castle**, Copyright © 2000-2004, the Legion of Bouncy Castle. For more information, disclaimers and restrictions see http://www.bouncycastle.org.
- **Checkpoint.** Copyright © Check Point Software Technologies Ltd.
- **Concurrent**, utility package. Copyright © Doug Lea. Used without CopyOnWriteArrayList and ConcurrentReaderHashMap classes.
- **Crypto++ Compilation**. Copyright © 1995-2003, Wei Dai, incorporating the following copyrighted work: **mars.cpp** by Brian Gladman and Sean Woods. For more information, disclaimers and restrictions see http://www.eskimo.com/~weidai/License.txt.
- **Crystal Reports Developer and Crystal Reports Server**. Copyright © 2004 Business Objects Software Limited.
- **DataDirect Technologies Corp**. Copyright © 1991-2003.
- **edpFTPj**, licensed under the Lesser GNU Public License. For more information, disclaimers and restrictions see http://www.enterprisedt.com/products/edtftpj/purchase.html.
- **Enhydra Shark**, licensed under the Lesser General Public License available at: http://shark.objectweb.org/license.html.
- **ICEsoft ICEbrowser**. ICEsoft Technologies, Inc. Copyright © 2003-2004.
- **ILOG, Inc**. Copyright © 1999-2004.
- **Installshield Universal**. Copyright © 1996–2005, Macrovision Corporation and/or Macrovision Europe Ltd.
- **Java 2 Platform, Standard Edition**. Copyright © Sun Microsystems, Inc. For more information, disclaimers and restrictions see http://java.sun.com/j2se/1.4.2/j2re-1_4_2_10-license.txt.

  **The Java 2 Platform may also contain the following third-party products**:

  - CoolServlets © 1999
  - DES and 3xDES © 2000 by Jef Poskanzer
  - Crimson © 1999-2000 The Apache Software Foundation
  - Xalan J2 © 1999-2000 The Apache Software Foundation
  - NSIS 1.0j © 1999-2000 Nullsoft, Inc.
  - Eastman Kodak Company © 1992
  - Lucinda, a registered trademark or trademark of Bigelow and Holmes

  - Taligent, Inc.
  - IBM, some portions available at: http://oss.software.ibm.com/icu4j/

  For more information regarding these third-party technologies and their associated disclaimers and restrictions, see: http://java.sun.com/j2se/1.4.2/j2se-1_4_2-thirdpartylicensereadme.txt.

- **JavaBeans Activation Framework (JAF)**. Copyright © Sun Microsystems, Inc. For more information, disclaimers and restrictions see http://www.java.sun.com/products/javabeans/glasgow/jaf.html and click download > license.
- **JavaMail**. Copyright © Sun Microsystems, Inc. For more information, disclaimers and restrictions see http://www.java.sun.com/products/javamail/downloads/index.html and click download > license.
- **Java Ace**, by Douglas C. Schmidt and his research group at Washington University and **Tao (with ACE wrappers)** by Douglas C. Schmidt and his research group at Washington University, University of California, Irvine and Vanderbilt University. Copyright © 1993-2005. For more information, disclaimers and restrictions see http://www.cs.wustl.edu/~schmidt/ACE-copying.html and http://www.cs.wustl.edu/~pjain/java/ace/JACE-copying.html
- **Java Authentication and Authorization Service Modules**, licensed under the Lesser General Public License. For more information, disclaimers and restrictions see http://free.tagish.net/jaas/index.jsp.
- **Java Network Launching Protocol (JNLP).** Copyright © Sun Microsystems, Inc. For more information, disclaimers and restrictions, please see http://www.java.sun.com/products/javawebstart/download-jnlp.html and click download > license.
- **Java Service Wrapper**. Portions copyrighted as follows: Copyright © 1999, 2004 Tanuki Software and Copyright © 2001 Silver Egg Technology. For more information, disclaimers and restrictions, see http://wrapper.tanukisoftware.org/doc/english/license.html.
- **JIDE**. Copyright © 2002 to 2005, JIDE Software, Inc.
- **jTDS** is licensed under the Lesser GNU Public License. For more information, disclaimers and restrictions see http://jtds.sourceforge.net/.
- **MDateSelector**. Copyright © 2005, Martin Newstead, licensed under the Lesser General Public License. For more information, disclaimers and restrictions see http://web.ukonline.co.uk/mseries.
- **Monarch Charts**. Copyright © 2005, Singleton Labs.
- **Net-SNMP**. Portions of the code are copyrighted by various entities, which reserve all rights. Copyright © 1989, 1991, 1992 by Carnegie Mellon University; Copyright © 1996, 1998 to 2000, the Regents of the University of California; Copyright © 2001 to 2003 Networks Associates Technology, Inc.; Copyright © 2001 to 2003, Cambridge Broadband, Ltd.; Copyright © 2003 Sun Microsystems, Inc. and Copyright © 2003 to 2004, Sparta, Inc. For more information, disclaimers and restrictions, see http://net-snmp.sourceforge.net.
- **The OpenSSL Project**. Copyright © 1998-2004. the Open SSL Project. For more information, disclaimers and restrictions, see http://www.openssl.org.
- **Oracle Help for Java**. Copyright © 1994-2006, Oracle Corporation.
- **RoboHELP Office.** Copyright © Adobe Systems Incorporated, formerly Macromedia.

- **Skin Look and Feel (SkinLF)**. Copyright © 2000-2006 L2FProd.com. Licensed under the Apache Software License. For more information, disclaimers and restrictions see https://skinlf.dev.java.net/.
- **Sonic Software Corporation**. Copyright © 2003-2004. The SSC software contains security software licensed from RSA Security, Inc.
- **Tinyxml**. For more information, disclaimers and restrictions see http://grinninglizard.com/tinyxmldocs/index.html.
- **SecurityNexus**. Copyright © 2003 to 2006. SecurityNexus, LLC. All rights reserved.
- **Xalan** and **Xerces**, both of which are licensed by the Apache Software Foundation Copyright © 1999-2004. For more information, disclaimers and restrictions see http://xml.apache.org/dist/LICENSE.txt.
- **yWorks.** Copyright © 2003 to 2006, yWorks.

**NOTE**: As of the publication of this documentation, the above links were active. In the event you find that any of the above links are broken or the linked webpages are inactive, please contact e-Security's Office of the Counsel at 1921 Gallows Road, Vienna, VA 22182. 703-852-8000.