DO NOT REPRINT
© FORTINET



Advanced Analytics Study Guide

for FortiSIEM 6.7

FORTINET®
Training Institute

DO NOT REPRINT
© FORTINET

Fortinet Training Institute - Library

https://training.fortinet.com

Fortinet Product Documentation

https://docs.fortinet.com

Fortinet Knowledge Base

https://kb.fortinet.com

Fortinet Fuse User Community

https://fusecommunity.fortinet.com/home

Fortinet Forums

https://forum.fortinet.com

Fortinet Product Support

https://support.fortinet.com

FortiGuard Labs

https://www.fortiguard.com

Fortinet Training Program Information

https://www.fortinet.com/nse-training

Fortinet | Pearson VUE

https://home.pearsonvue.com/fortinet

Fortinet Training Institute Helpdesk (training questions, comments, feedback)

https://helpdesk.training.fortinet.com/support/home

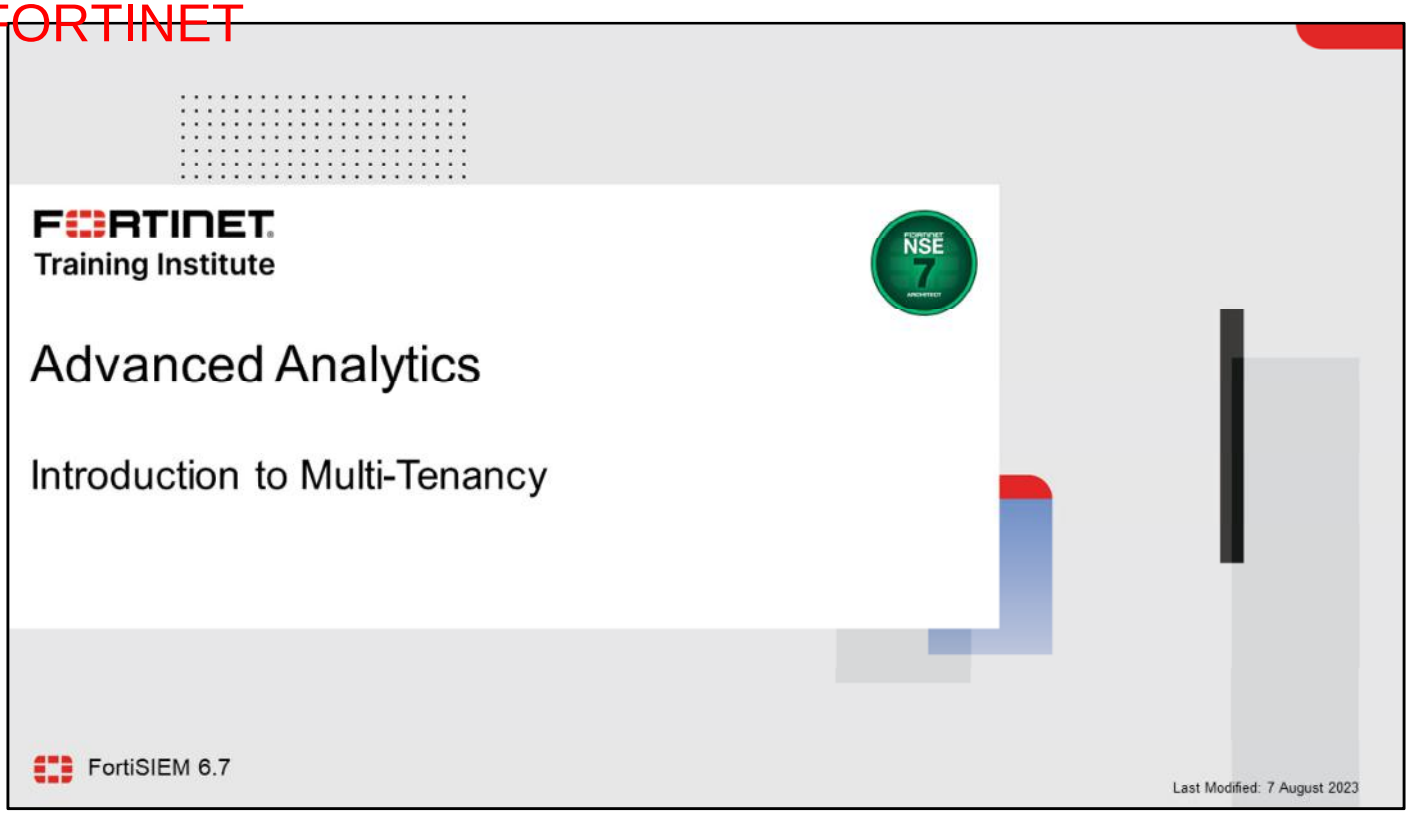
DO NOT REPRINT
© FORTINET

Brave-Dumps.com

TABLE OF CONTENTS

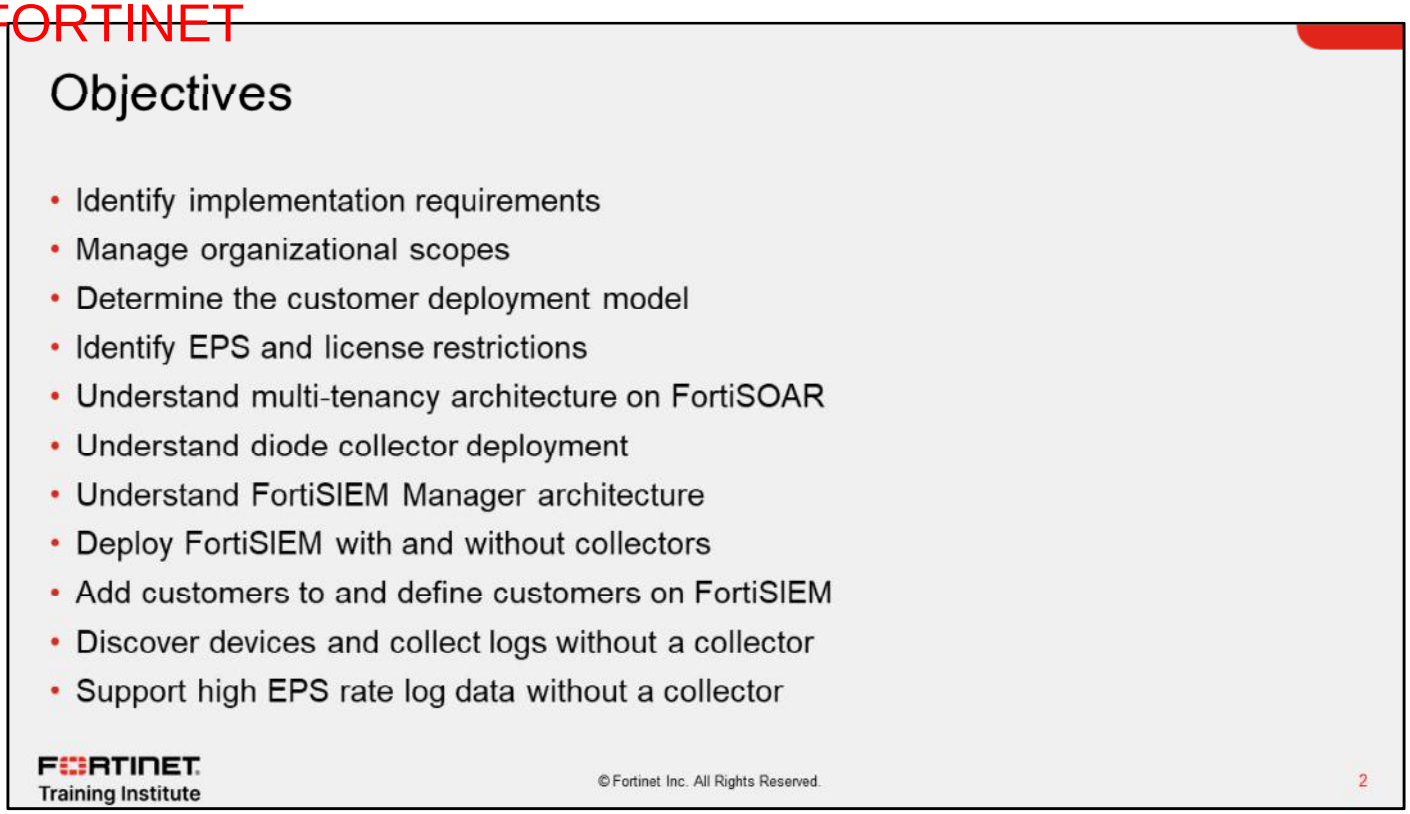
01 Introduction to Multi-Tenancy 4
02 Defining FortiSIEM Collectors and FortiSOAR Connectors 46
03 Operating Collectors 73
04 Windows and Linux Agents 123
05 Rules 172
06 Single Subpattern Security Rules 200
07 Multiple Subpattern Rules 241
08 Baselines 265
09 Baseline Rules 322
10 FortiSIEM UEBA 356
11 Nested Queries and Lookup Tables 390
12 Clear Conditions 447
13 Remediation 472

DO NOT REPRINT
© FORTINET



In this lesson, you will learn about deploying FortiSIEM in a multi-tenant environment, and identify key implementation requirements. You will also learn about deploying FortiSIEM with and without collectors, and how FortiSIEM collect logs without a collector.

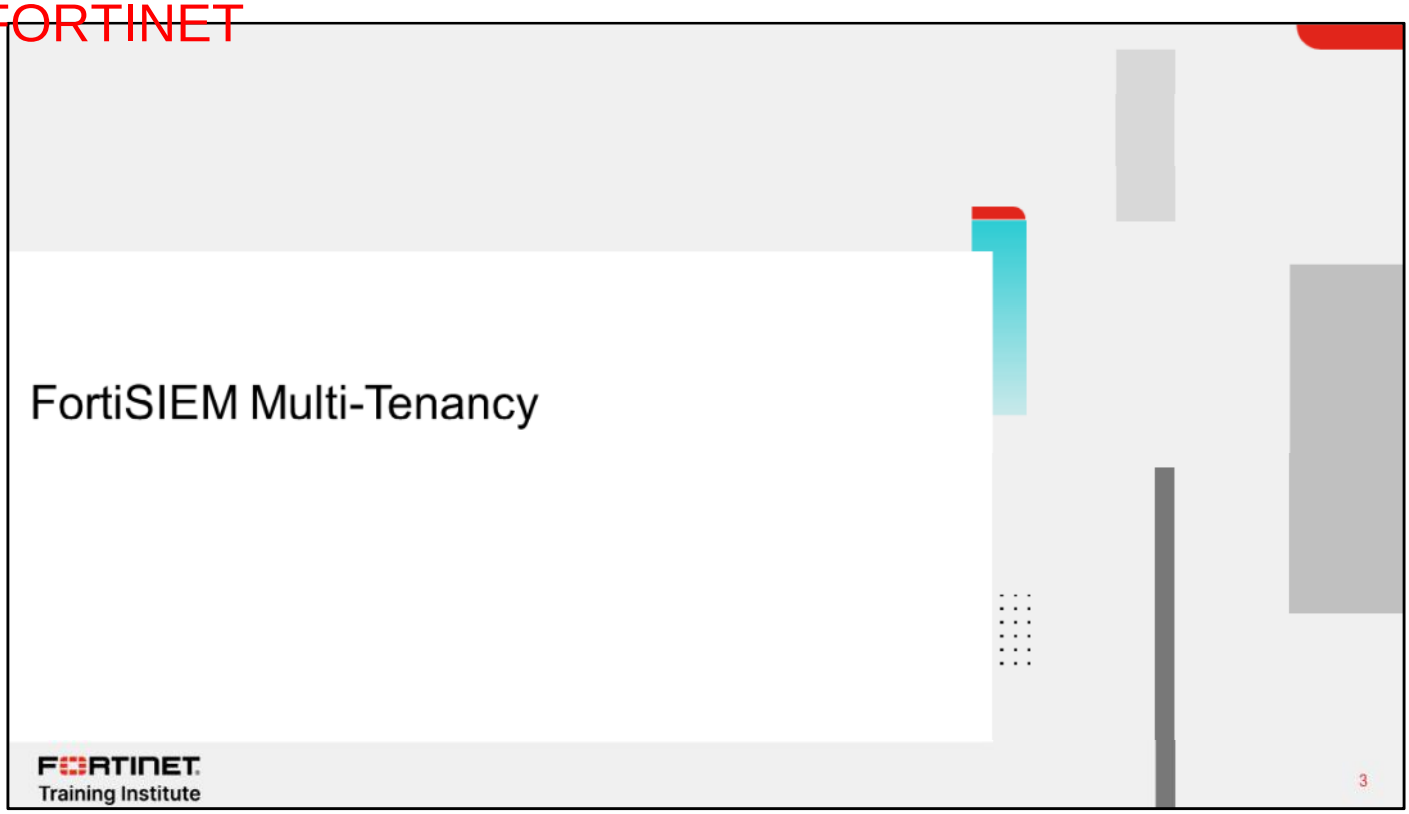
DO NOT REPRINT
© FORTINET



After completing this lesson, you should be able to achieve the objectives shown on this slide.

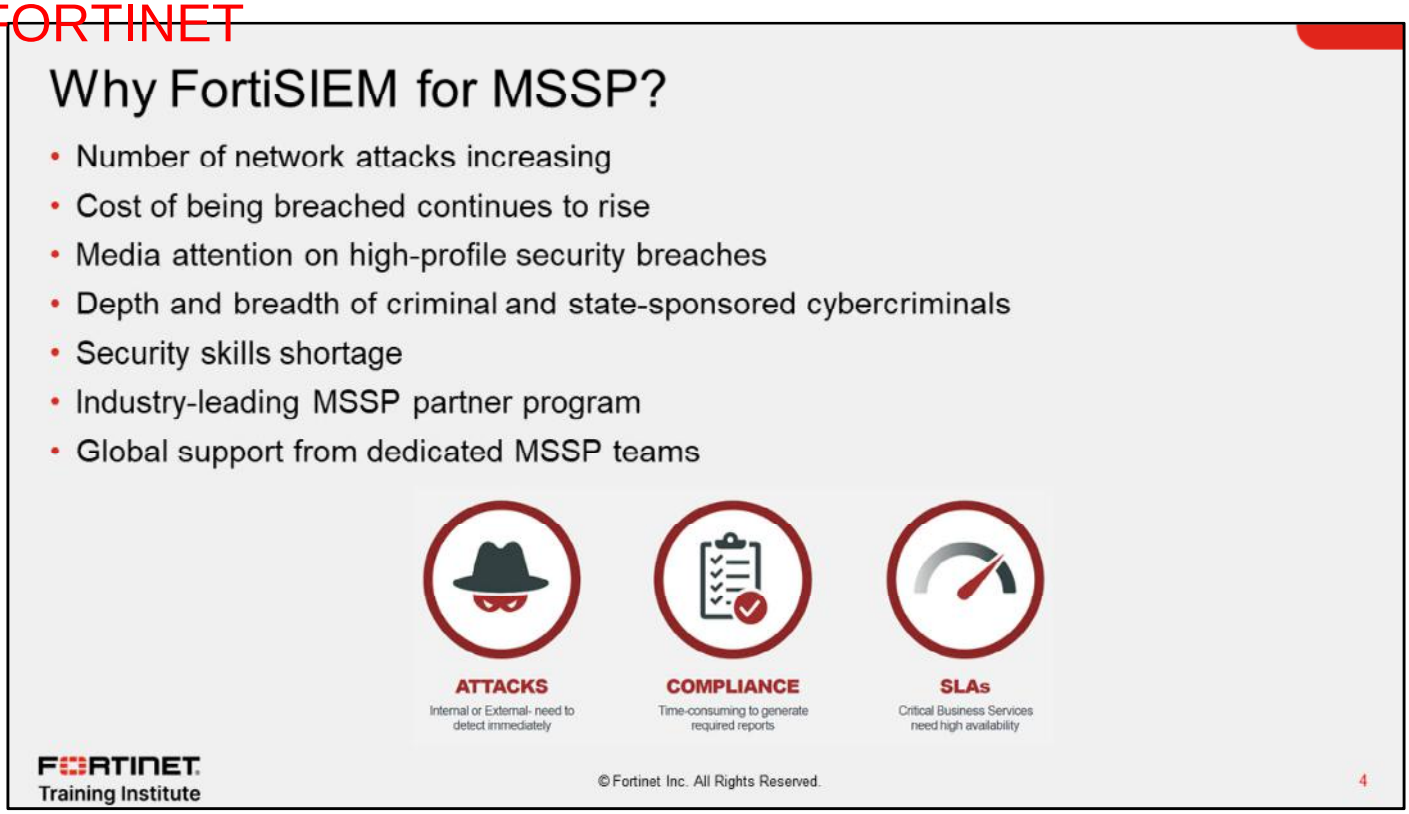
By demonstrating competence in multi-tenancy, you will be able to determine a customer deployment model, and deploy FortiSIEM with and without a collector.

DO NOT REPRINT
© FORTINET



In this section, you will learn about FortiSIEM multi-tenancy and how you can manage customer scopes on FortiSIEM.

DO NOT REPRINT
© FORTINET

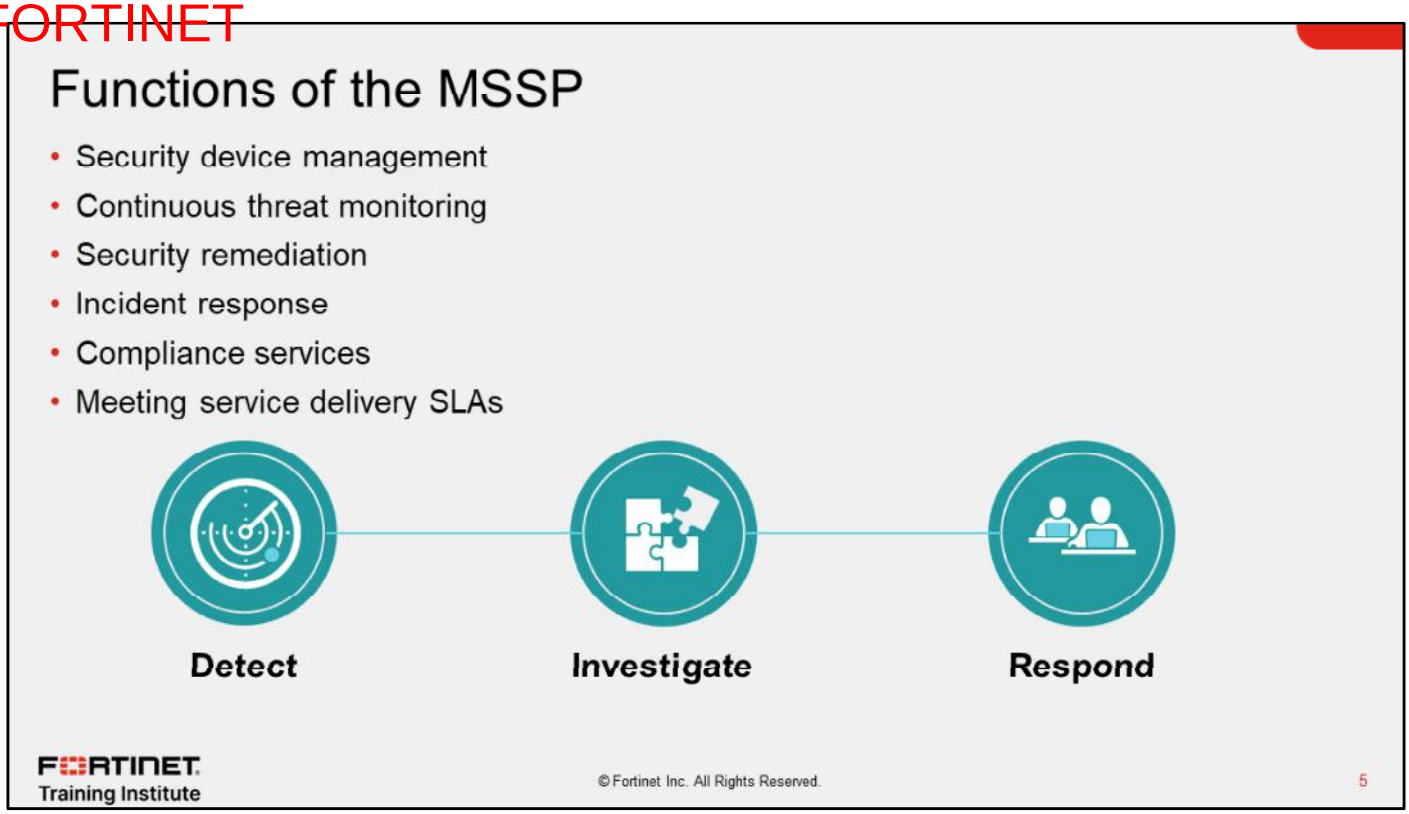


Businesses worldwide are struggling to secure their networks against determined attackers who are commonly two steps ahead of the defenders. There is a perfect storm of challenges facing organizations wanting to expand into the new digital market, without compromising their security or brand.

Studies show that not only is the number of network attacks increasing, but that the cost of being breached has also continued to rise with each incident. Media attention on high-profile security breaches has also created a higher level of awareness among business leaders, about the risks and costs of security breaches to businesses. The depth and breadth of criminal and state-sponsored cybercriminals, who are often well funded and are developing increasingly sophisticated attacks, is better understood by security buyers today than ever before, and at the same time, the security skills shortage continues to widen.

As a result, many CISOs are looking to migrate some or all of their risk out of their IT departments and into the hands of professionals such as managed security service providers (MSSPs).These MSSPs are expected to anticipate and secure networks against innovative and determined threat actors.

DO NOT REPRINT
© FORTINET

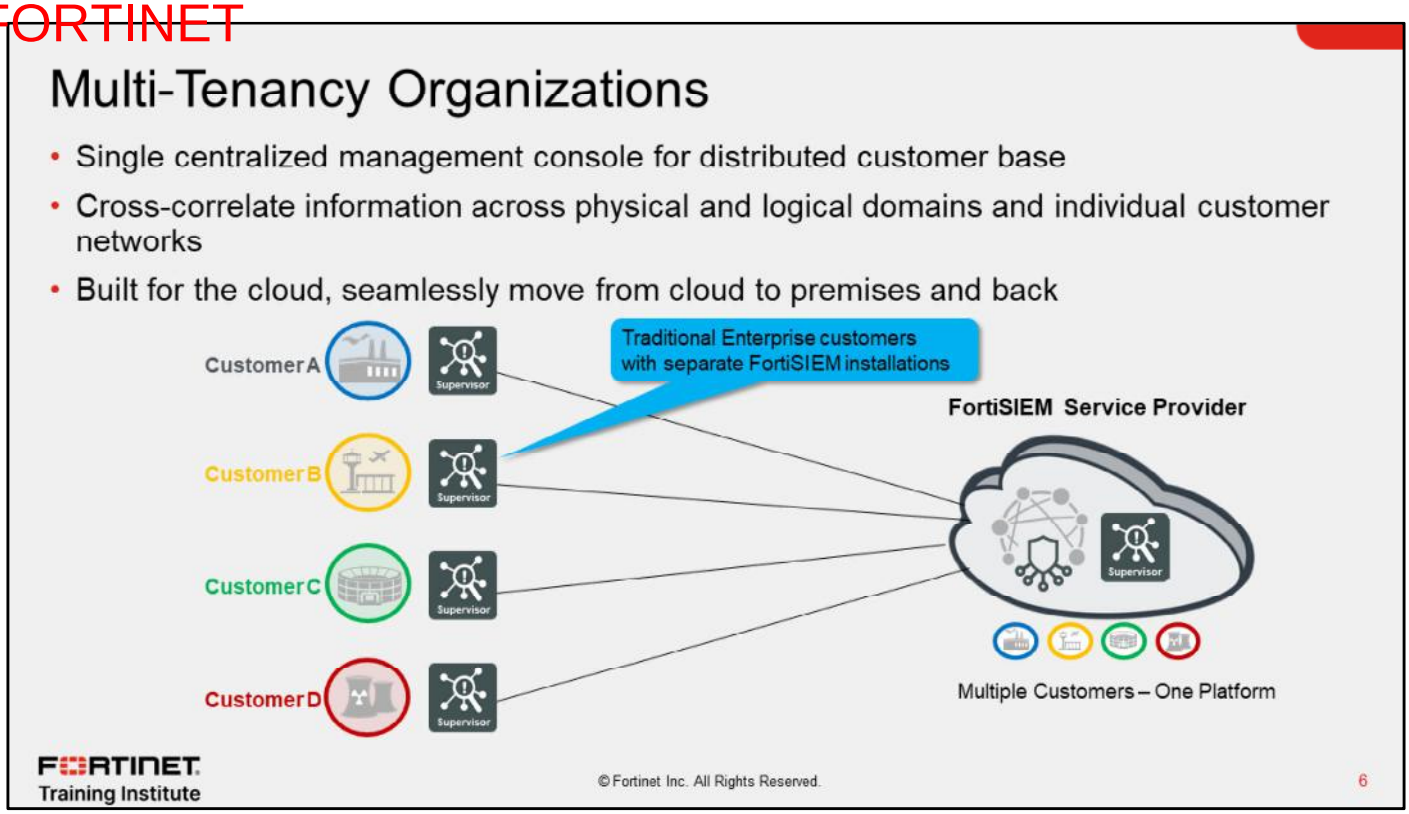


At its core, an MSSP must provide two kinds of basic services: security device management and continuous monitoring. Customers with multiple locations and business-critical applications running across the network have stringent uptime requirements; so, device management and monitoring are crucial to maintain business productivity.

Many service providers also add security analytics and threat intelligence services to help mitigate new attacks, including actionable intelligence and a comprehensive view of the distributed security infrastructure. Going forward, these firms will likely differentiate themselves in new areas such as security analytics, threat intelligence, information portals, and customer service.

Many MSSPs will also want to provide security remediation, incident response, compliance services, or loss prevention, to further differentiate their organization in the market. Securing thousands of customer locations, meeting service-level agreements (SLAs), and balancing profitability with engineering headcount and infrastructure are ongoing challenges facing MSSPs. So MSSPs need specialized security vendors that understand their business model and can help them meet customer requirements.

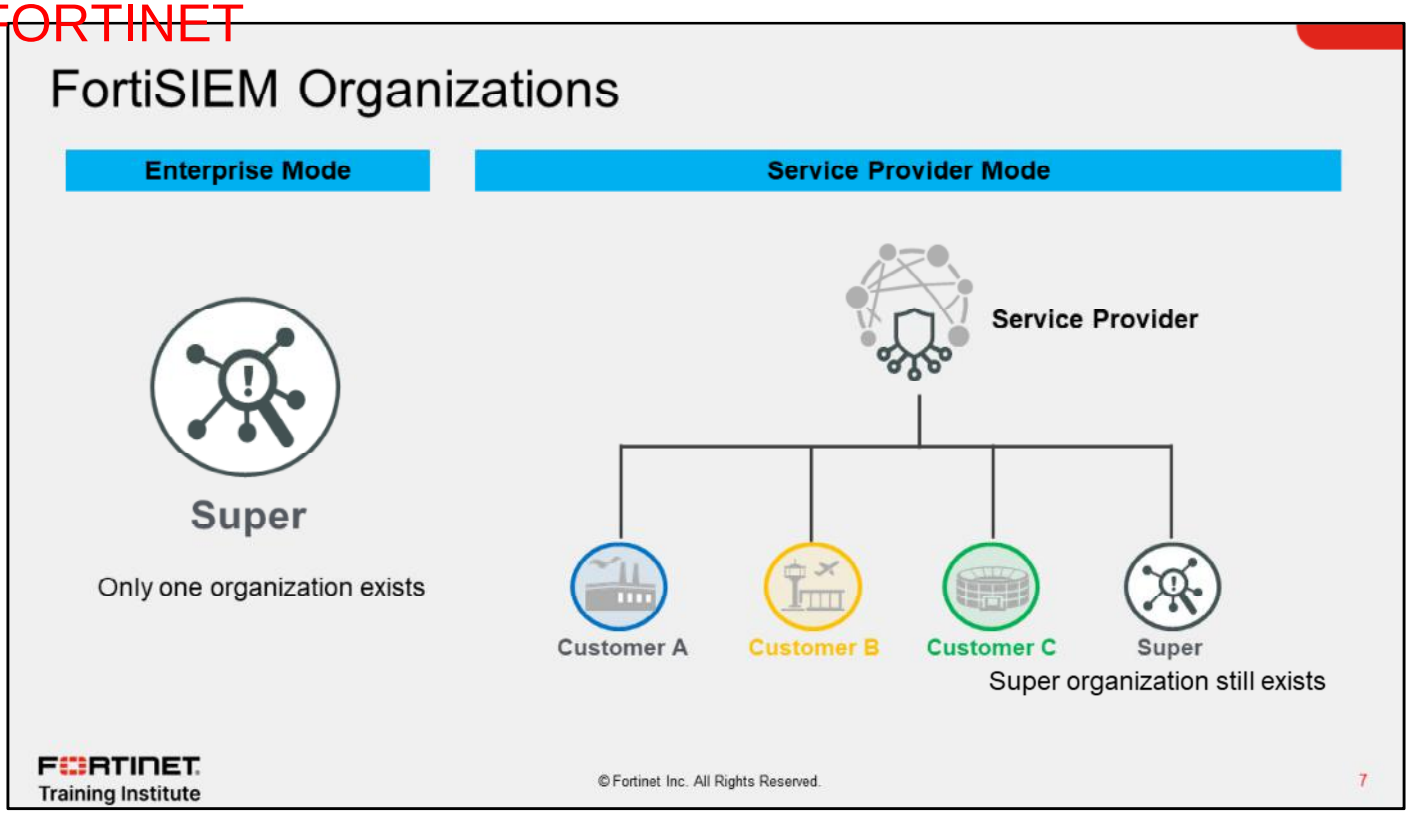
DO NOT REPRINT
© FORTINET



A key requirement for MSSPs with a large, distributed customer base is support for multi-tenancy, or the ability to support multiple customers from a single centralized management console. FortiSIEM is a highly customizable, multi-tenant architecture that enables MSSPs to manage a large number of physical domains, logical domains, and overlapping systems and networks, from a single console. In this environment, it is very easy to cross-correlate information across physical and logical domains and individual customer networks. Not only is FortiSIEM natively multi-tenant, it is also built for the cloud. It is easy to install and manage either as a virtual machine (VM) in an MSSP or customer data center, or as an instance in an Amazon Web Services (AWS) or Microsoft Azure environment. FortiSIEM can even be run as a hybrid model, and is the only product with the ability to seamlessly move from cloud to premises and from premises to cloud.

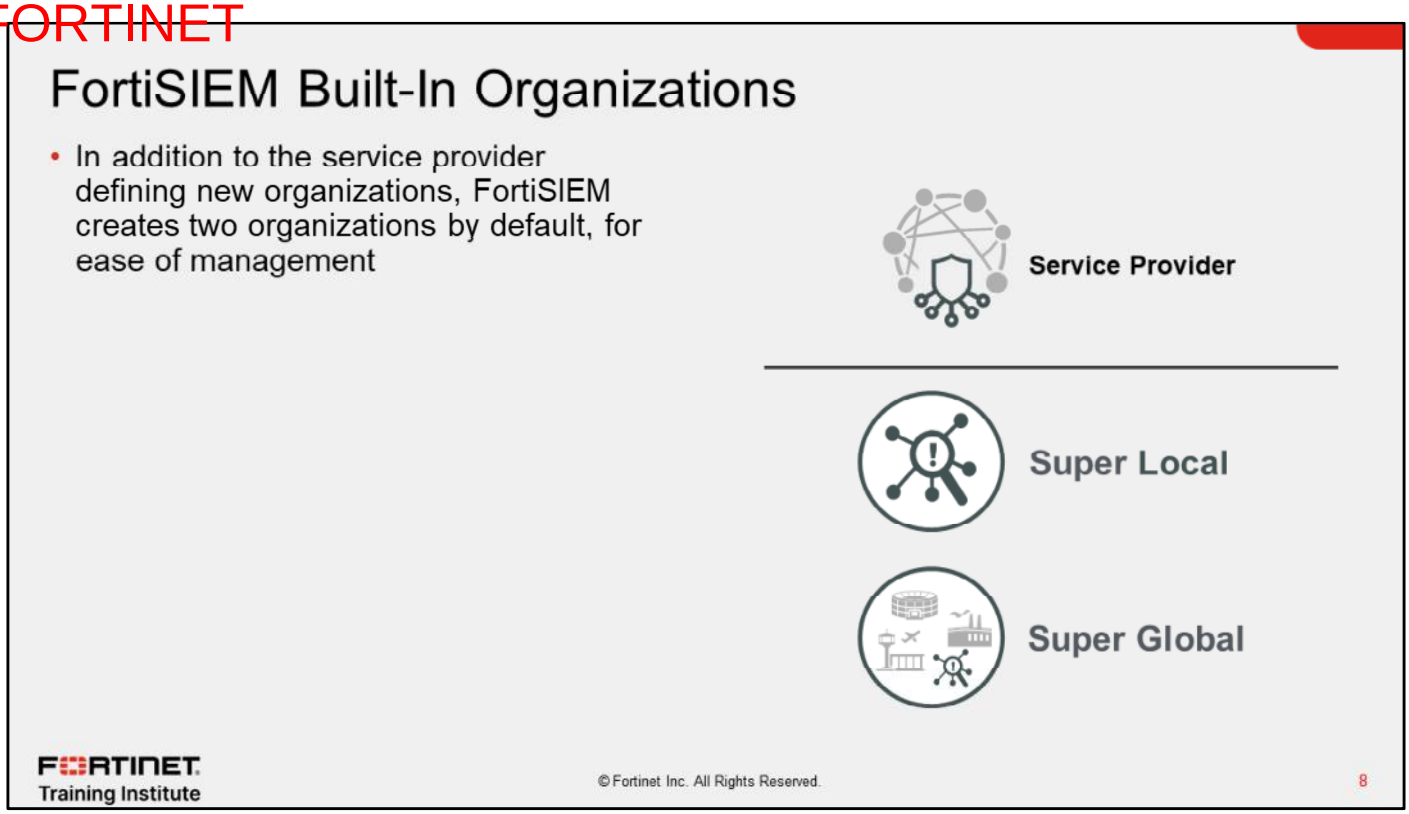
Compartmentalizing SIEM operations can be done in a variety of ways. Just as administrative domains (ADOMs) can be broken into sections based on customer, location, and so on, FortiSIEM can be divided into what are known as organizations. This allows the MSSP to overlay an aggregate view of all tenants, that allows data and events to be managed based on business requirements.

DO NOT REPRINT
© FORTINET



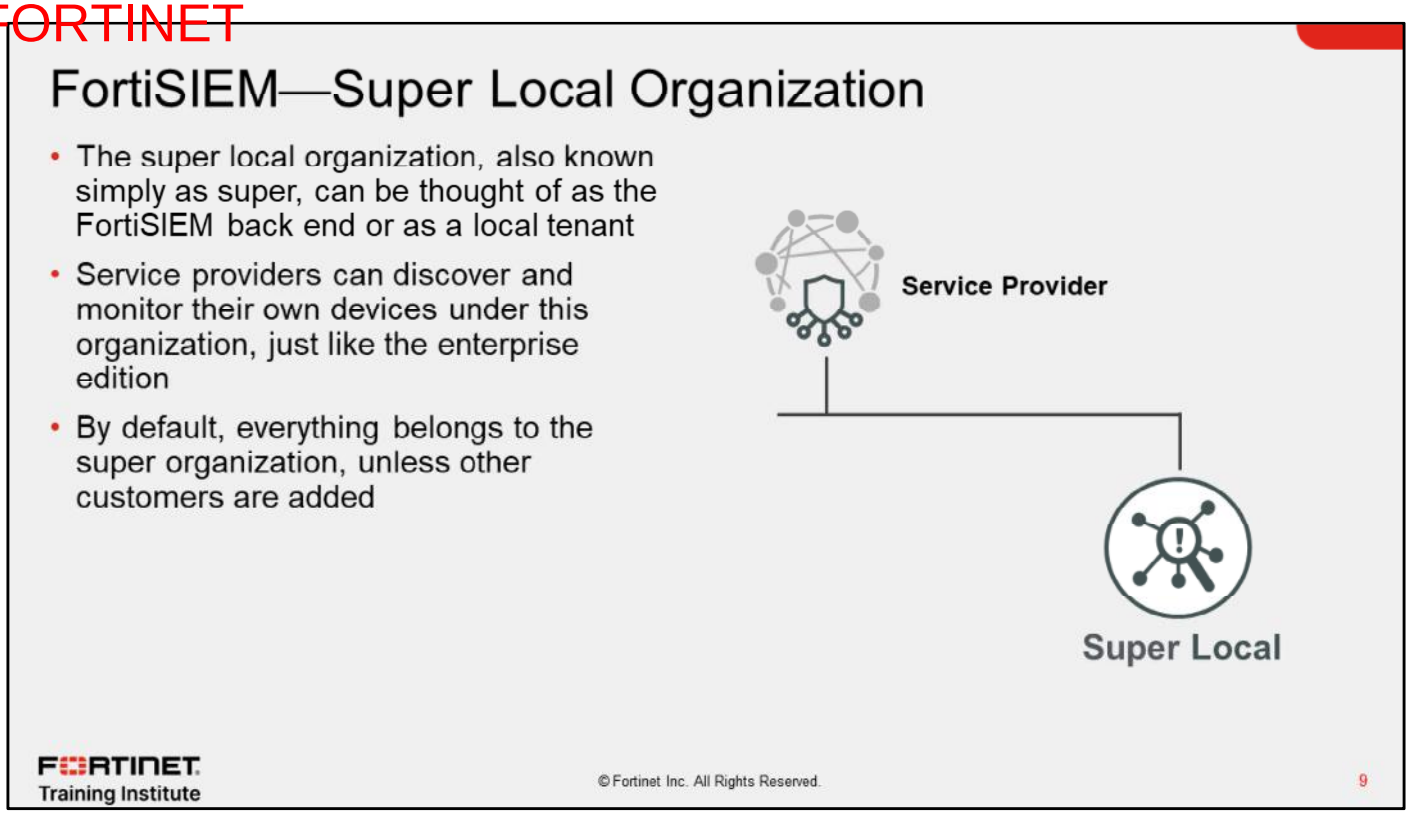
In the enterprise mode of FortiSIEM, there is just one organization, called the super organization. In a service provider deployment, FortiSIEM supports multi-tenancy through multiple organizations, alongside the super organization.

DO NOT REPRINT
© FORTINET



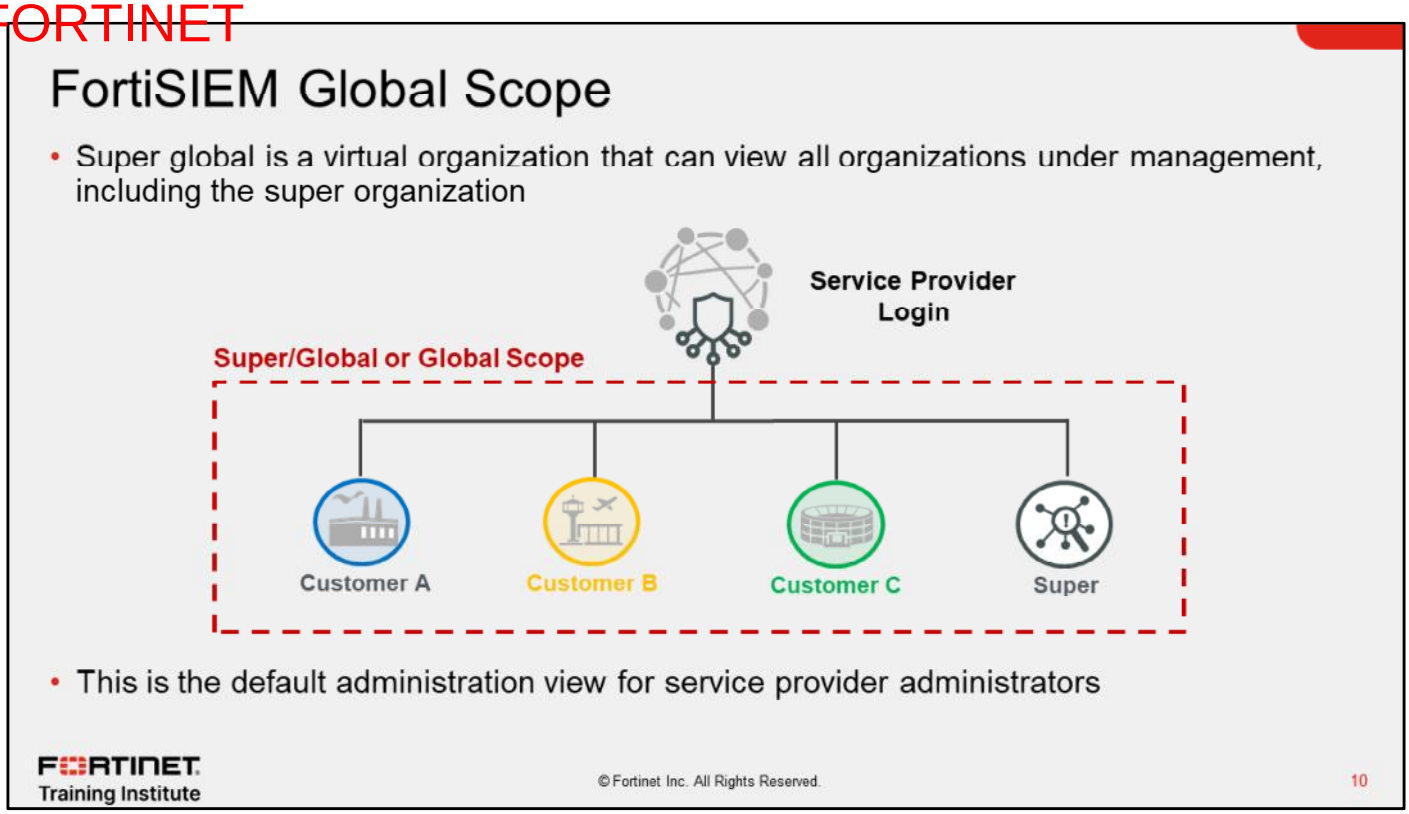
In addition to the service provider defining new organizations, FortiSIEM creates two organizations by default, for ease of management. The super local and super global organizations are built in, and are not shown under the **Organizations** tab.

DO NOT REPRINT
© FORTINET



The super local organization, also known simply as super, can be thought of as the FortiSIEM back end, or a local tenant. Service providers can discover and monitor their own devices under this organization just like the enterprise edition. By default, everything belongs to the super organization, unless other customers are added.

DO NOT REPRINT
© FORTINET



Super global is a virtual organization that can view all organizations under management, including the super organization. Users belonging to the super global organization can see other organizations and their data.

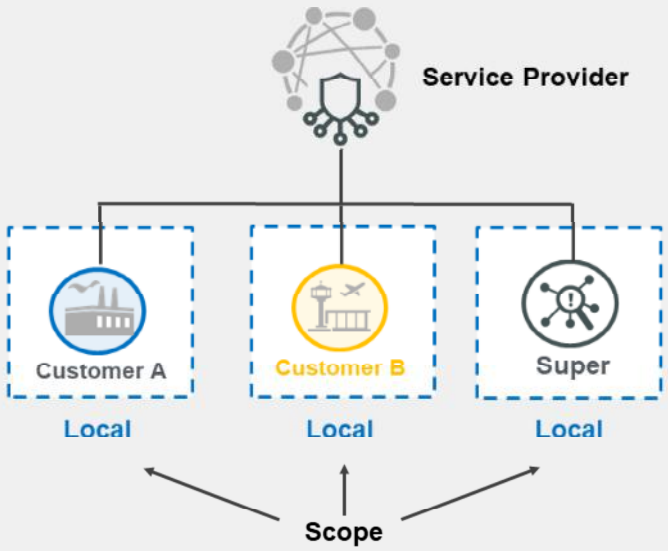
After installation, FortiSIEM automatically creates an administrator user with full administrator rights for the super global and super local organizations. When the user creates a new organization, FortiSIEM creates an administrator user for that organization. These are accounts with full administrator rights. FortiSIEM users with full administrator rights can create roles and then create users and assign them a role.

You can configure roles on FortiSIEM to restrict what users can see and do. You can restrict GUI visibility, organization visibility, and data by writing filters on device type, event type, and any parsed event attribute.

DO NOT REPRINT
© FORTINET

Views and Scopes

- Scopes are organization views
- Local is just the view of an individual tenant
- Local views are available for all customers under the super global management



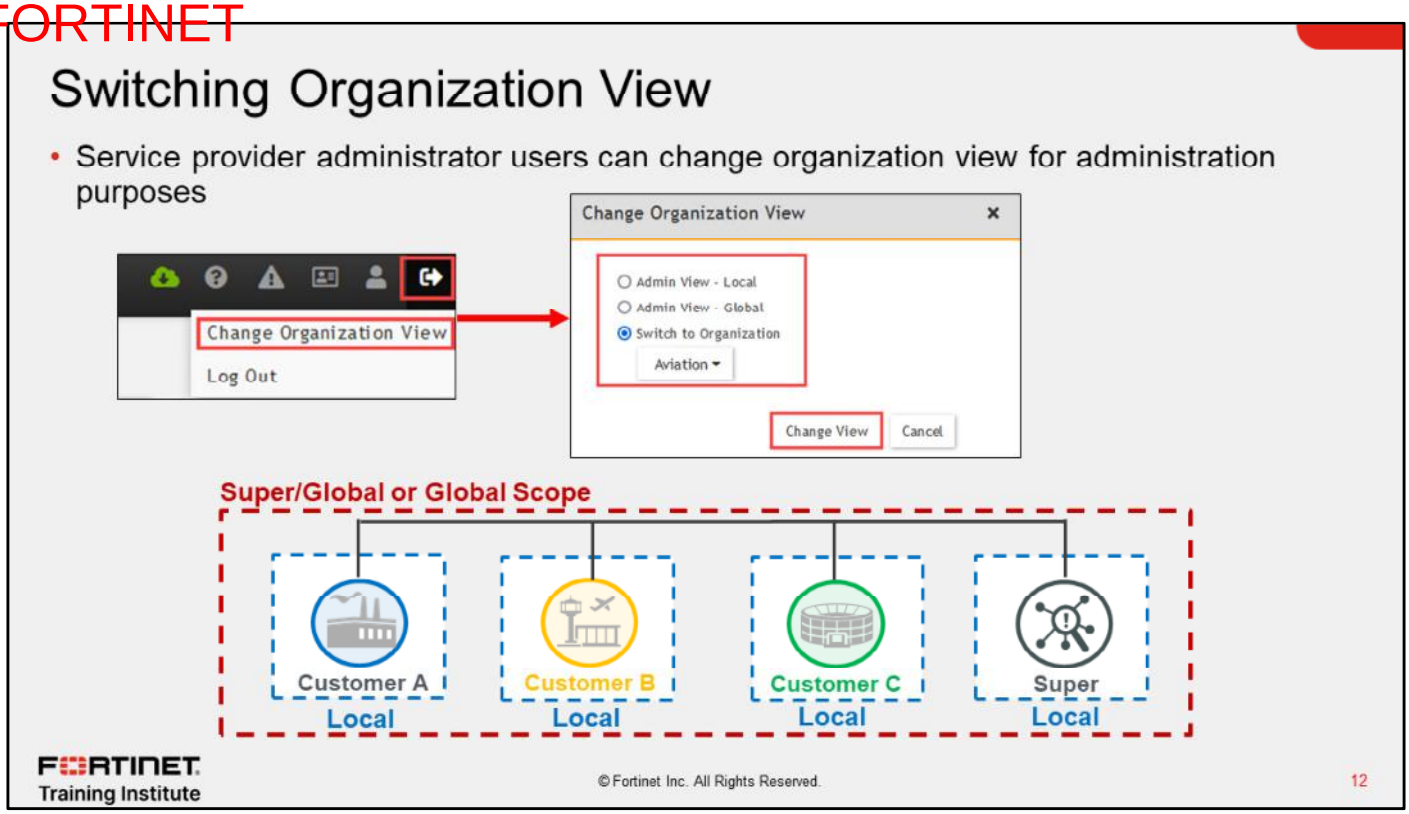
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

11

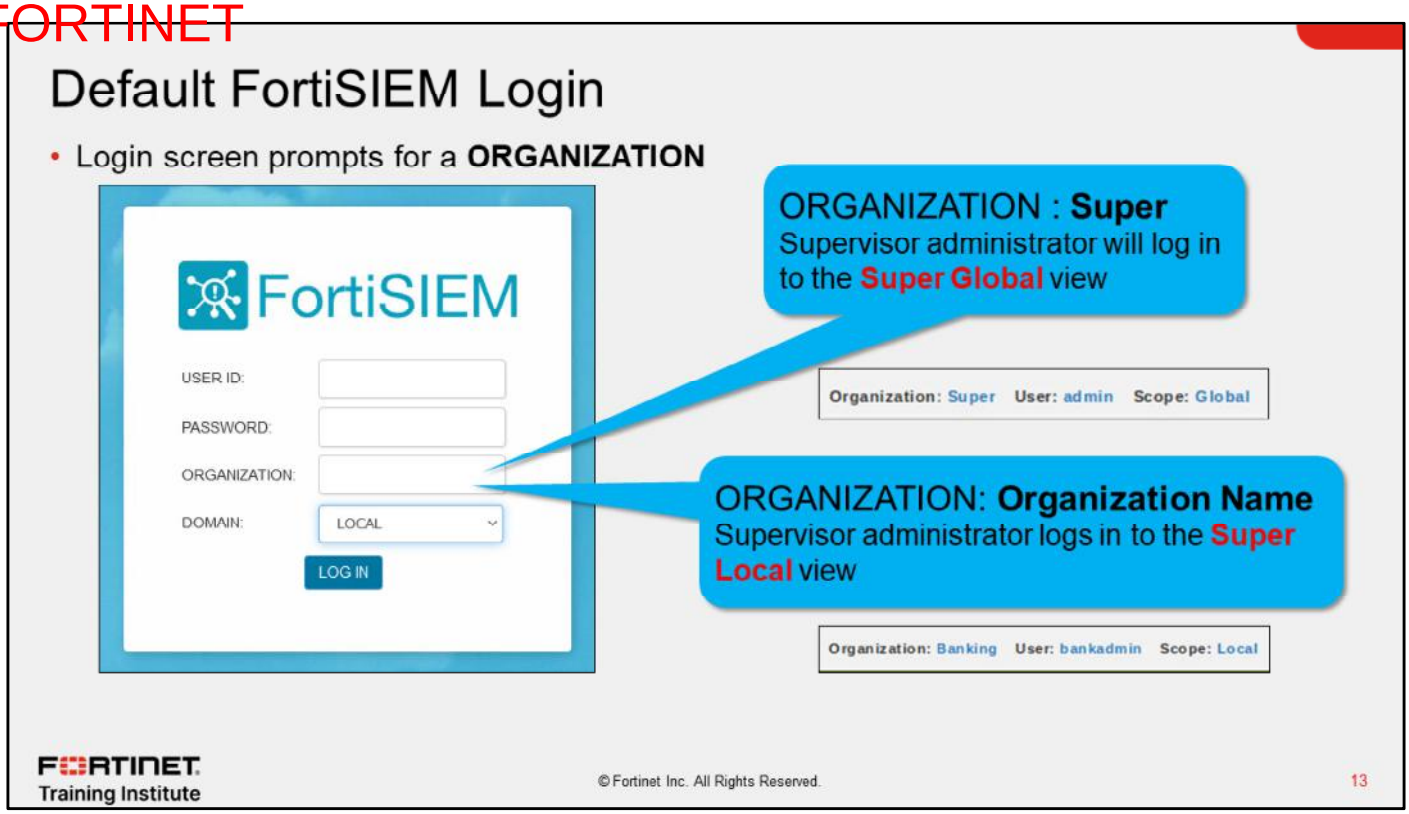
Scopes on FortiSIEM are administrative views where logs sent by a collector from a customer location can be viewed locally. Therefore, the scope is sometimes called local for an individual customer. Users belonging to local organizations and user-defined organizations can see only their own organization.

DO NOT REPRINT
© FORTINET



Service provider administrator users can change scopes for administration purposes. This allows the administrator to change the organization view.

DO NOT REPRINT
© FORTINET



You can log in to the supervisor node as a super global user, if you are a service provider administrator. To log in to the supervisor node as a super global user, in the **ORGANIZATION** field, type `super`. You can also log in to the individual organization by typing the organization name in the **ORGANIZATION** field.

DO NOT REPRINT
© FORTINET

Logging in to the Customer View

- A customer administrator can view their organization by logging in using the customer administrator credentials from the default FortiSIEM login



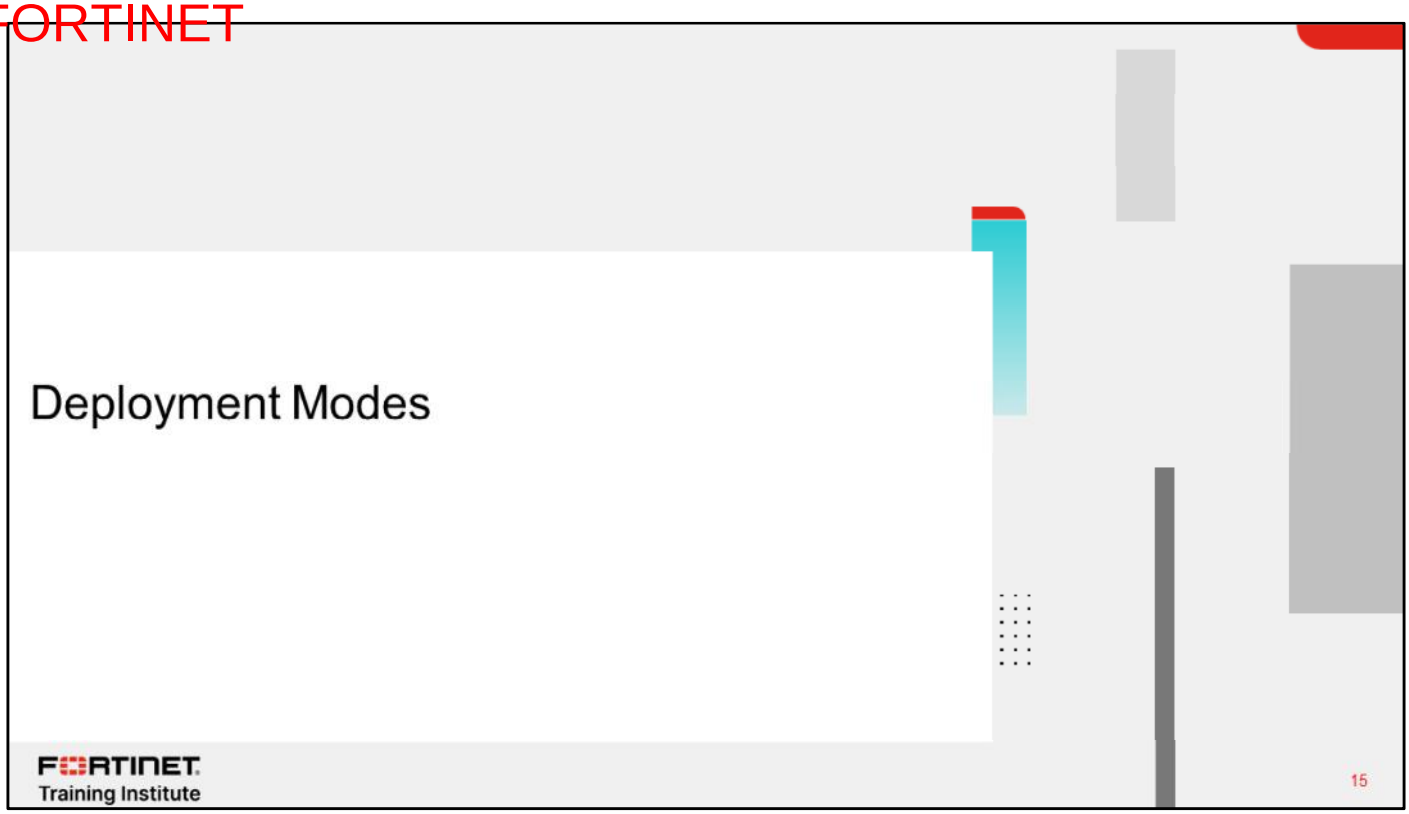
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

14

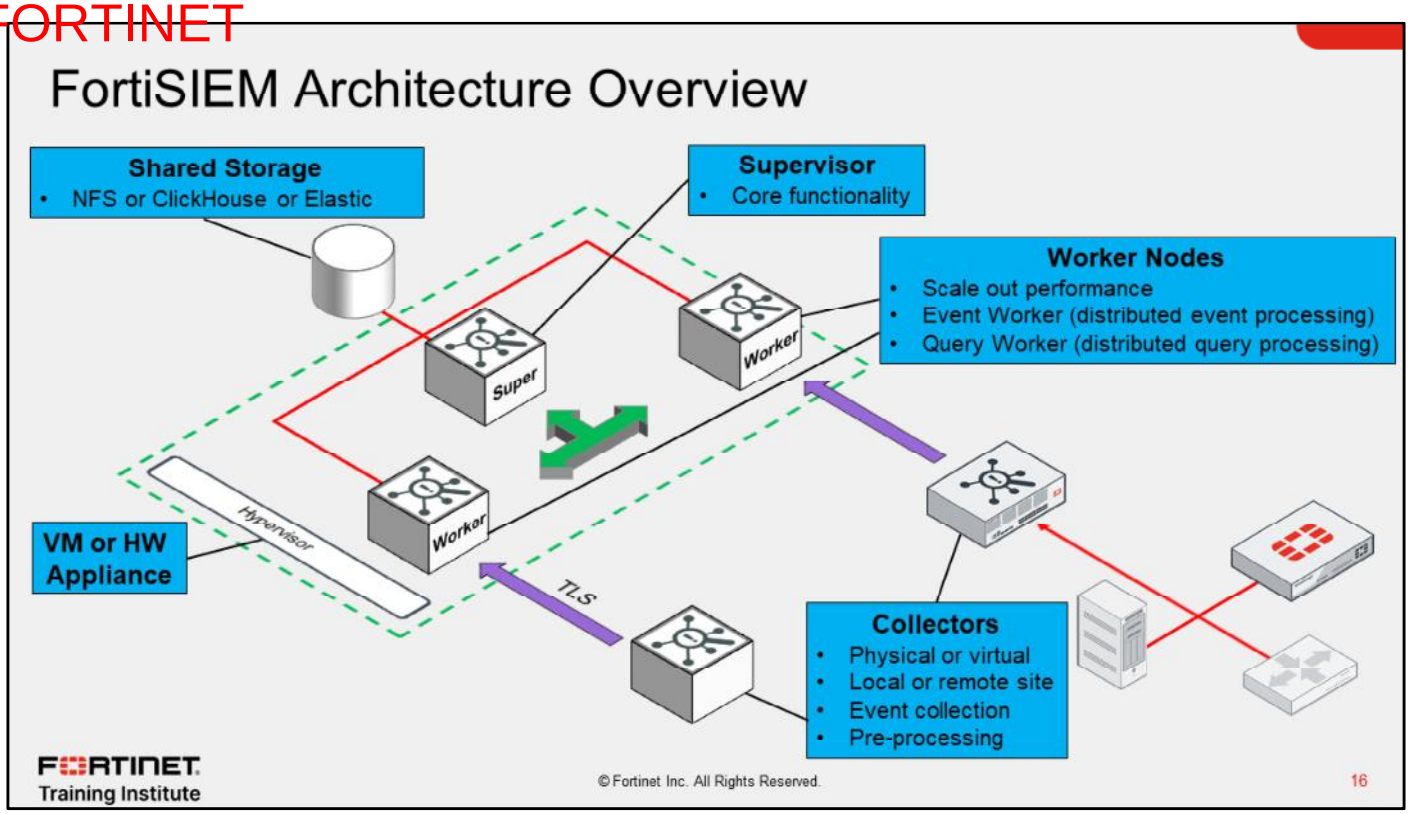
If you are the administrator for a customer, you can access your organization view by logging in to the default login screen with an administrator account that was set up for your organization.

DO NOT REPRINT
© FORTINET



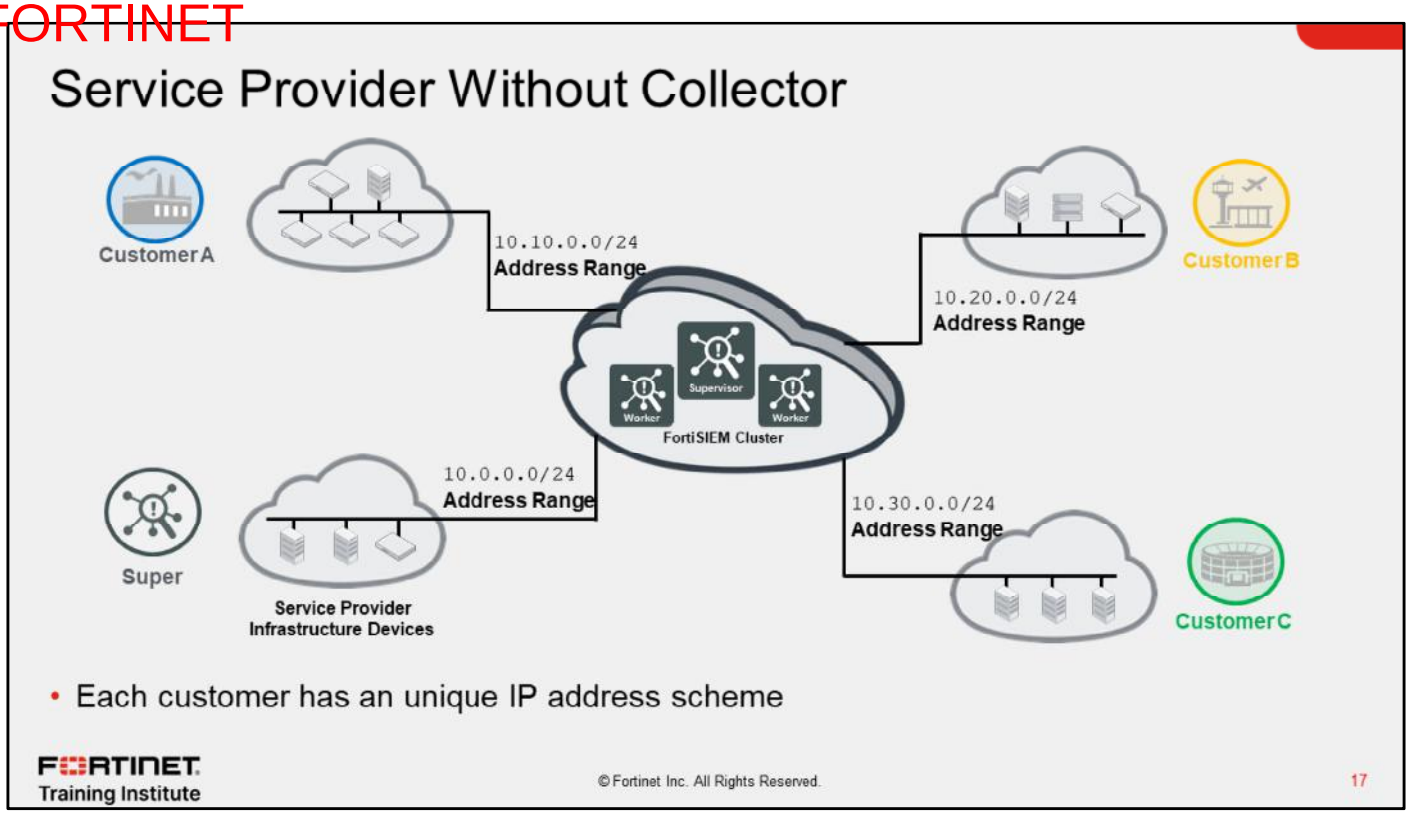
In this section, you will learn about various common deployment modes for FortiSIEM in a multi-tenant environment.

DO NOT REPRINT
© FORTINET



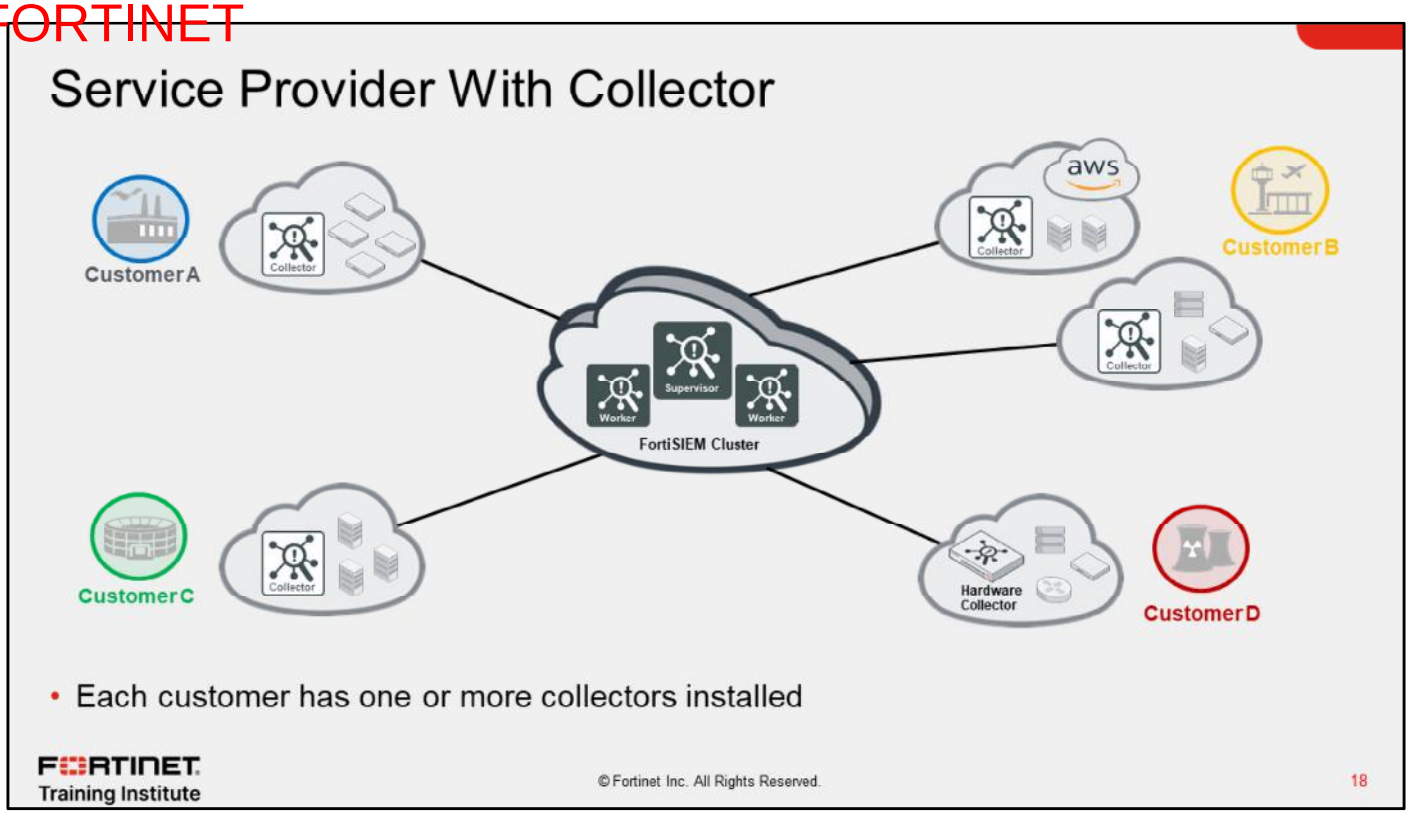
Our virtual appliance architecture uses a three tier structure to provide an easily scalable solution that can start as a small single node deployment, and rapidly scale to a large, high-performance system.
The supervisor node provides core functionality, and in a smaller solution can deliver an all-in-one system.
Worker nodes are used in conjunction with the supervisor node to scale event processing and reporting.
Collectors can be used to provide remote site log collection, and to offload log collection from the supervisor or worker nodes for increased scalability.

DO NOT REPRINT
© FORTINET



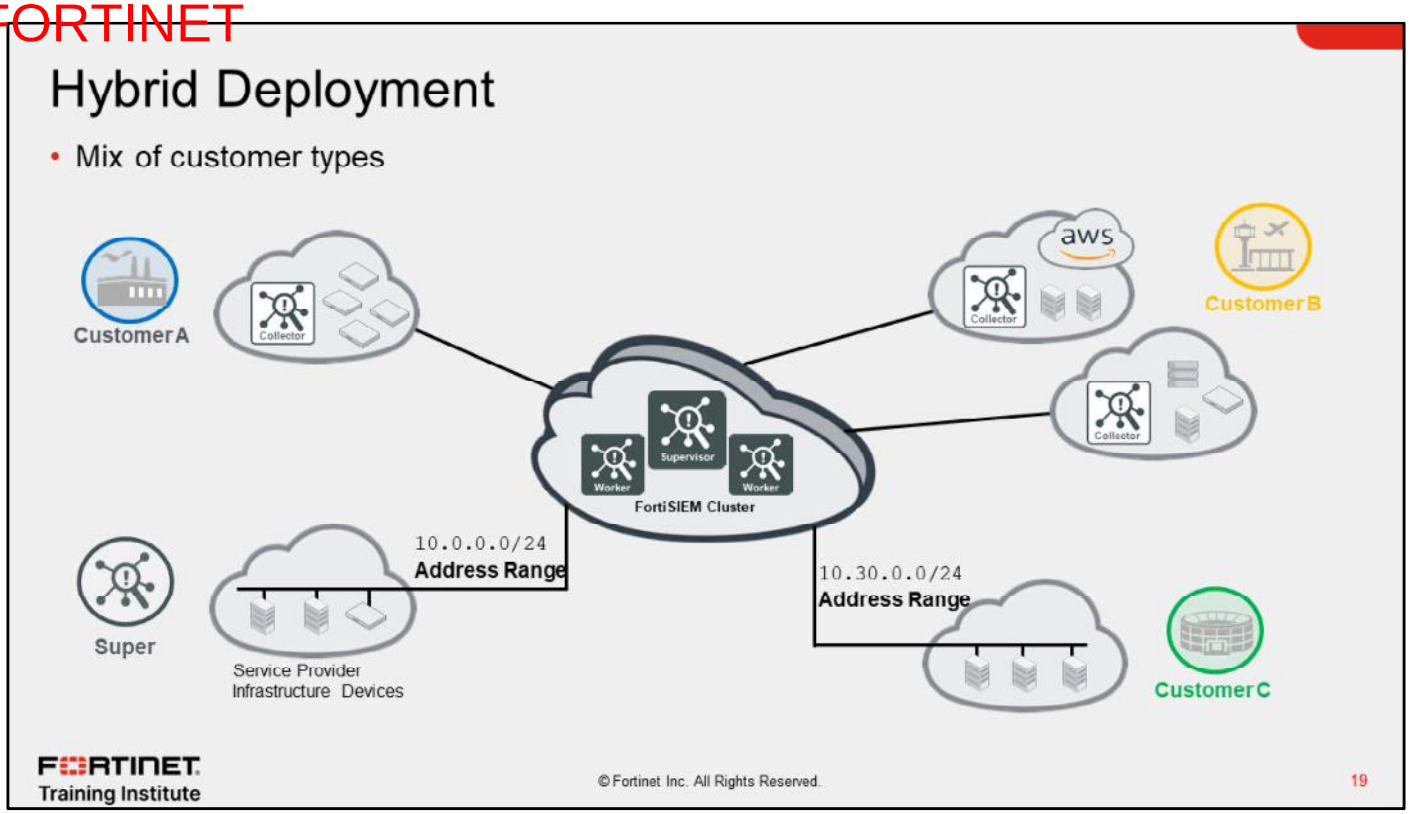
Service providers can deploy FortiSIEM without a collector, which is best suited for a hosting type environment. The key is that each customer is on a unique IP address scheme, with no overlap allowed. Typically, each customer device is local to the FortiSIEM cluster, and you can distinguish events and incidents by filtering with the reporting IP address of devices that belong to individual customers.

DO NOT REPRINT
© FORTINET



A service provider with a collector is the most common deployment type for service providers or very large enterprises using multi-tenancy features. This deployment model allows for overlapping IP address ranges. Each customer can have one or more collectors defined. Collectors can be placed anywhere on the LAN, WAN, DMZ, or remote sites across the internet or in the cloud. Additional benefits, such as remote administration of customer devices, is possible if collectors are used.

DO NOT REPRINT
© FORTINET

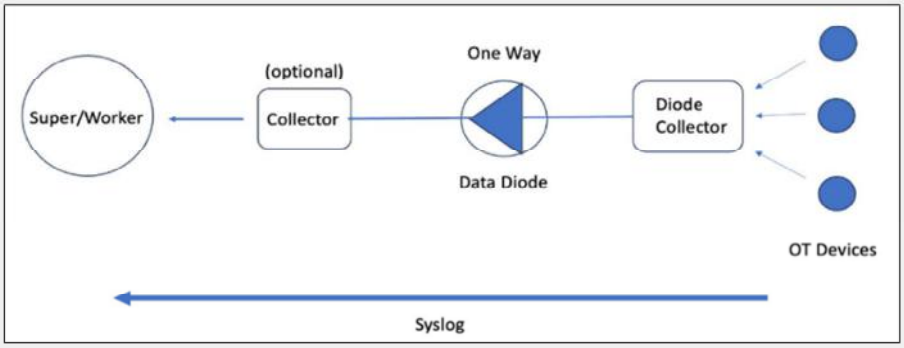


You can also deploy FortiSIEM in a hybrid manner. In hybrid deployments, some customers will have collectors, which are responsible for collecting and sending logs to the FortiSIEM cluster, while other customers can send logs directly to the FortiSIEM cluster. The rules of an overlapping IP address schema still apply in a deployment without a collector. Customers who do not have collectors will need to be on a distinct IP subnet.

DO NOT REPRINT
© FORTINET

Unidirectional Network Deployment — Diode Collector

- Install without internet connectivity
- Work without registering with supervisor node
- Collect syslog, SNMP traps, and Windows logs via WMI or OMI protocol using local configuration
- Send events to another Collector or Worker via UDP 514 using syslog protocol



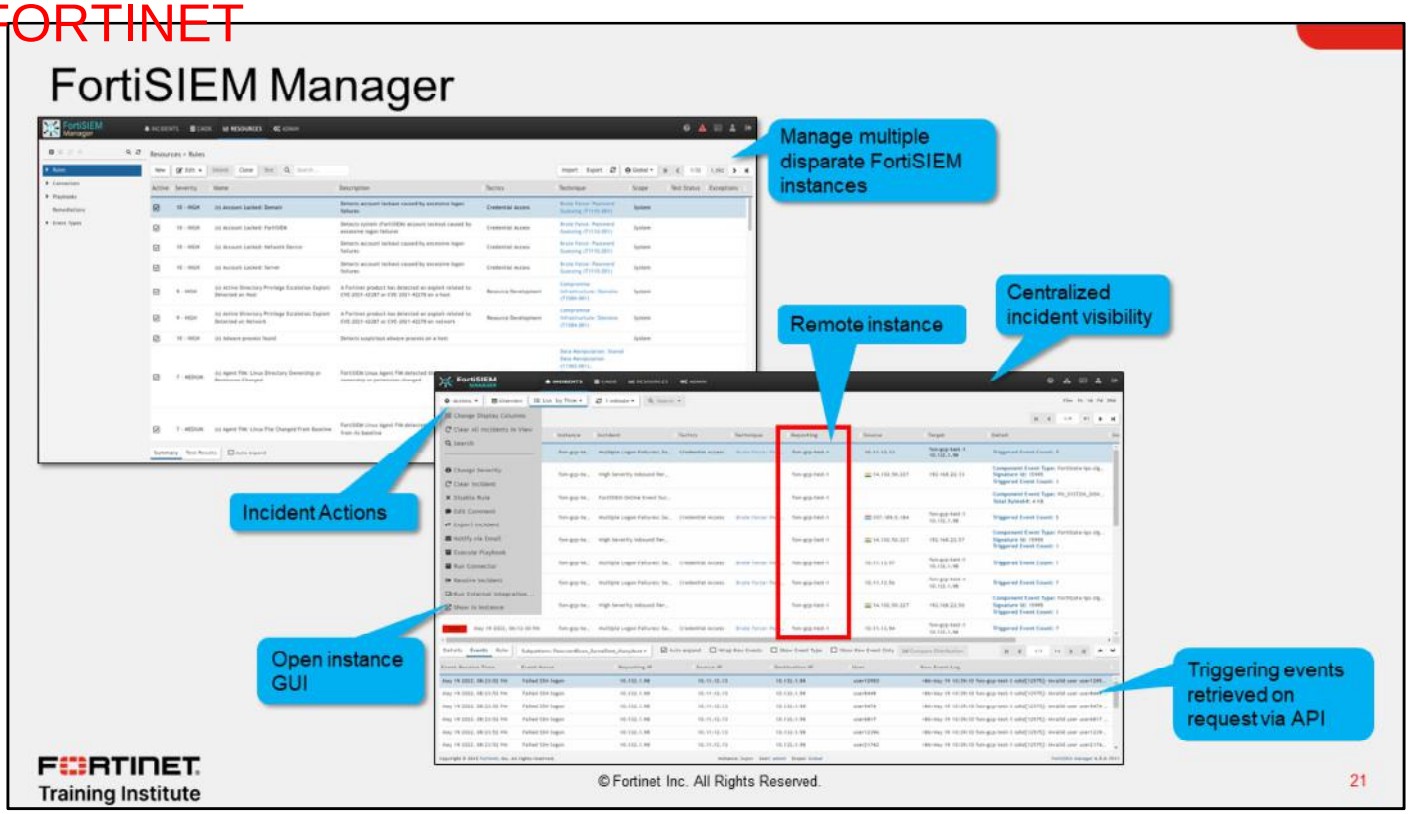
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

20

A diode collector only requires a strictly one-way communication from itself to another collector or worker.
Typically deployed in OT networks or in highly secured environments.
There are two deployment modes:
1. Diode collector to worker
2. Diode collector to regular collector, and then to worker
A Diode collector sends events to worker or regular collector using syslog over port UDP 514, while a regular collector uses HTTPS over port TCP 443 to register with supervisor and to send events to workers.

DO NOT REPRINT
© FORTINET



FortiSIEM Manager is a dedicated management node. FortiSIEM Manager cannot ingest events directly but you can use it to monitor and manage multiple FortiSIEM instances. FortiSIEM Manager needs to be installed on a separate virtual machine and requires a separate license. Each FortiSIEM instance needs to register on the FortiSIEM Manager.

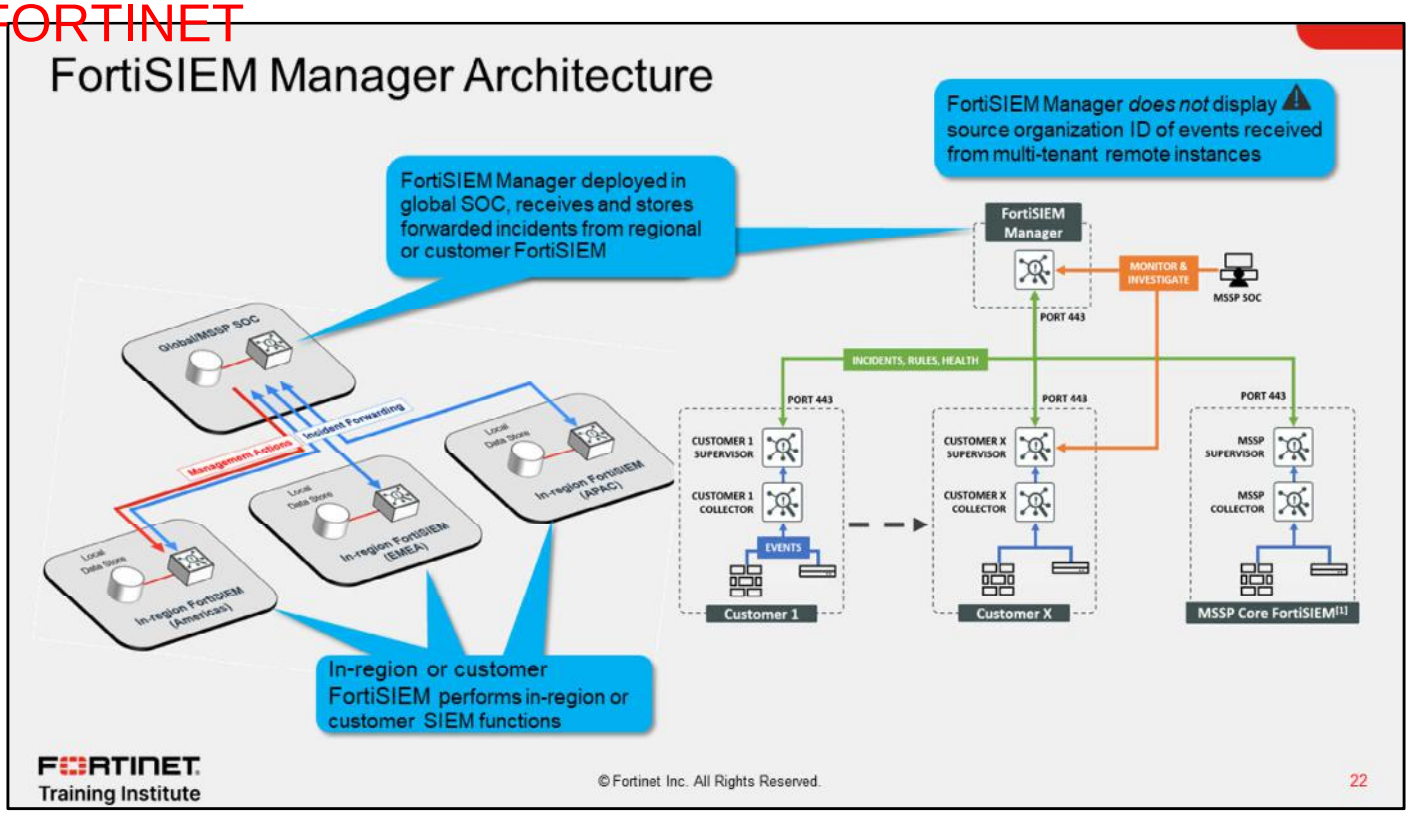
Each FortiSIEM instance forwards incidents, license, and health information to FortiSIEM Manager. Incidents are forwarded in near-real time, health information is forwarded once every minute, and license information is forwarded once every hour. FortiSIEM Manager also stores incidents and the latest license information in a local PostgreSQL database. The number of incidents stored depends on the size of the local PostgreSQL database.

FortiSIEM does not store raw events. It fetches raw events on demand (when a user visits the triggering event tab on the **INCIDENTS** page) from the appropriate FortiSIEM instance. All incident status changes in each FortiSIEM instance are forwarded to FortiSIEM Manager.

If you create a new rule or make changes to a rule in a FortiSIEM instance, the instance forwards the changes to FortiSIEM Manager. From FortiSIEM Manager, you can conduct the following operations, and FortiSIEM Manager propagates the changes to the correct FortiSIEM instances with the correct FortiSIEM Manager logged-in-user context:

- Clear, resolve, and add comments to one or more incidents. You can also change the severity of an incident, and disable one or more rules and change their severity.
- Run FortiSOAR playbooks and connectors and update incident status and comments.
- Log in with one click to the appropriate FortiSIEM instance where an incident occurred. This enables you quickly to investigate an incident in depth.

DO NOT REPRINT
© FORTINET

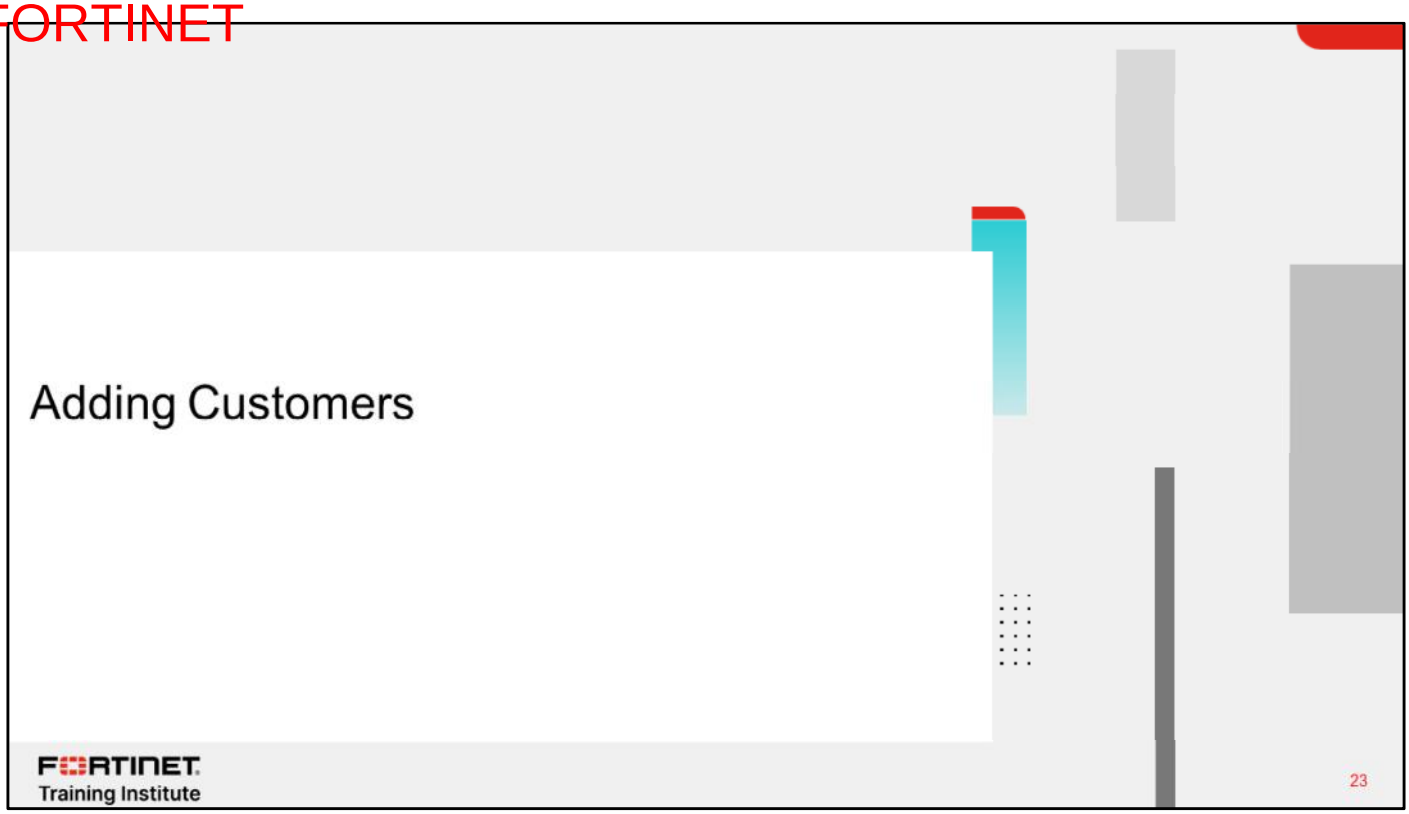


FortiSIEM Manager, deployed in a global SOC, receives and stores forwarded incidents from regional or customer FortiSIEM instances. FortiSIEM Manager does not display source organization ID of events received from multi-tenant remote instances. Each FortiSIEM Manager node can manage from 5-50 remote FortiSIEM instances.

Communication between FortiSIEM Manager and remote instances is through REST API over HTTP(S). You must upgrade FortiSIEM Manager first before upgrading all FortiSIEM instances—this applies to both content update and software image updates.

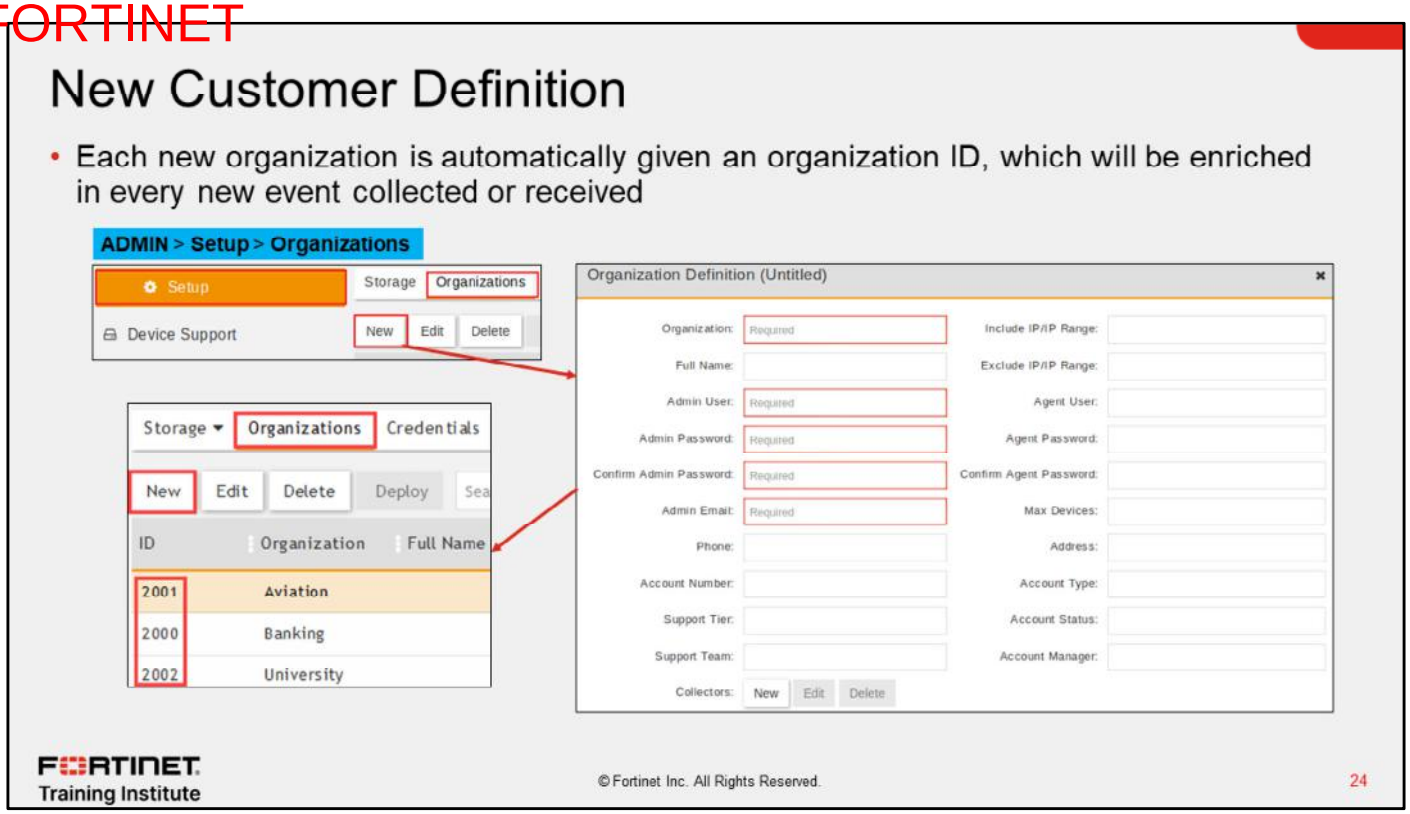
Only FortiSIEM Manager and FortiSIEM Supervisor instances 6.5.0 or later are supported.

DO NOT REPRINT
© FORTINET



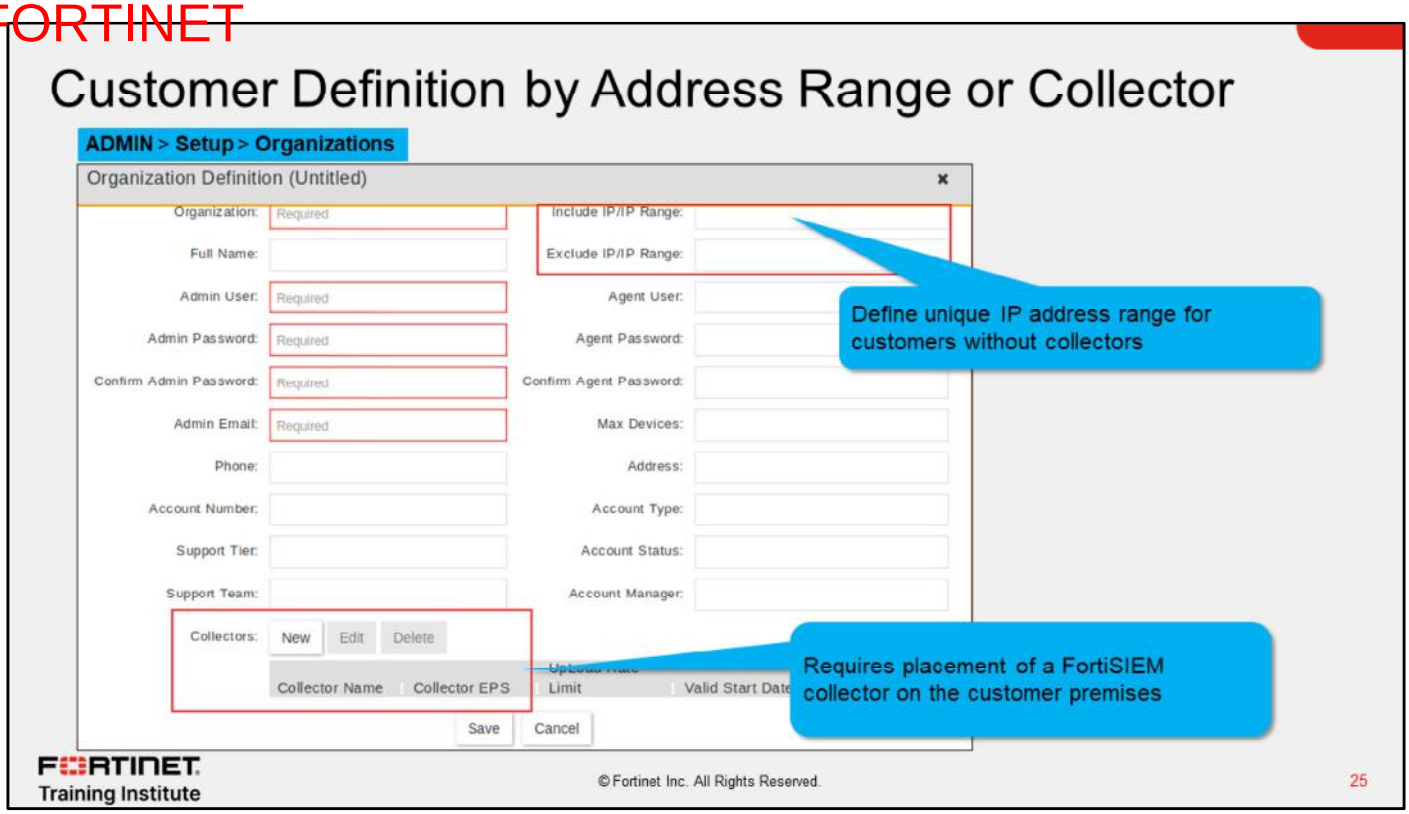
In this section, you will learn about adding customers on FortiSIEM, and defining collectors for those customers. You will also learn about adding an IP range for customers without collectors.

DO NOT REPRINT
© FORTINET



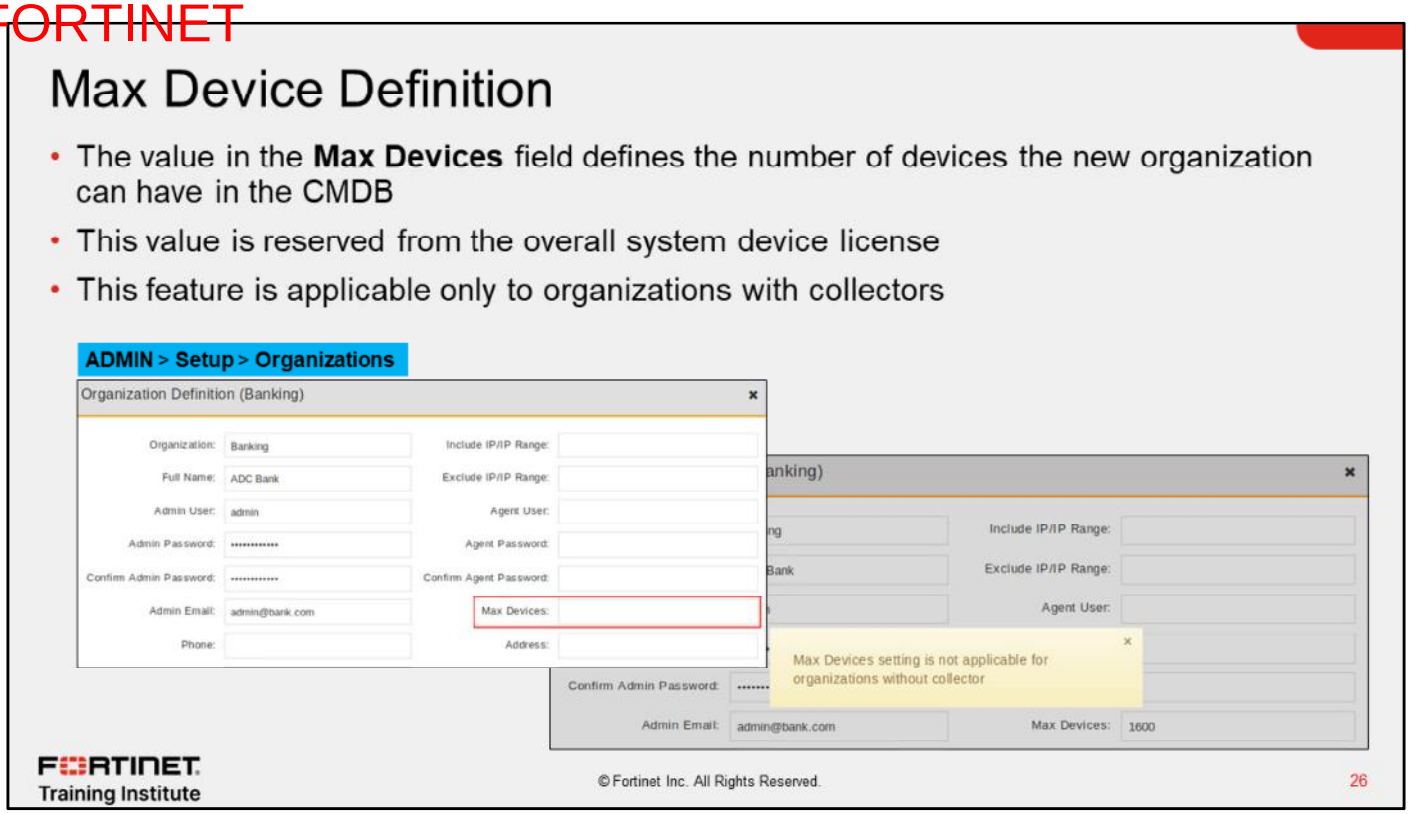
When you create a new organization, the organization name is the name of the new customer, which will be referenced on the GUI. Typing a value in the **Full Name** field is optional, and what you type here is not displayed elsewhere. The values that you type in the **Admin User** and **Admin Password** fields define the local administrator account username and password for this customer. The value that you type in the **Admin Email** filed is the email address that will be set for the administrator user for that customer, which is particularly useful for sending alerts and incidents to the administrator user for that organization. Each new organization is automatically given an organization ID, which will be enriched in every new event collected or received for that organization.

DO NOT REPRINT
© FORTINET



Organizations can be defined in one of two ways. You can associate one or more collectors with an organization. The devices monitored by the collector or the events sent to the collector, automatically belong to the associated organization. You can also define an IP range for an organization. If the sending IP of a device belongs in the IP range, then the device and logs belong to that organization.

DO NOT REPRINT
© FORTINET

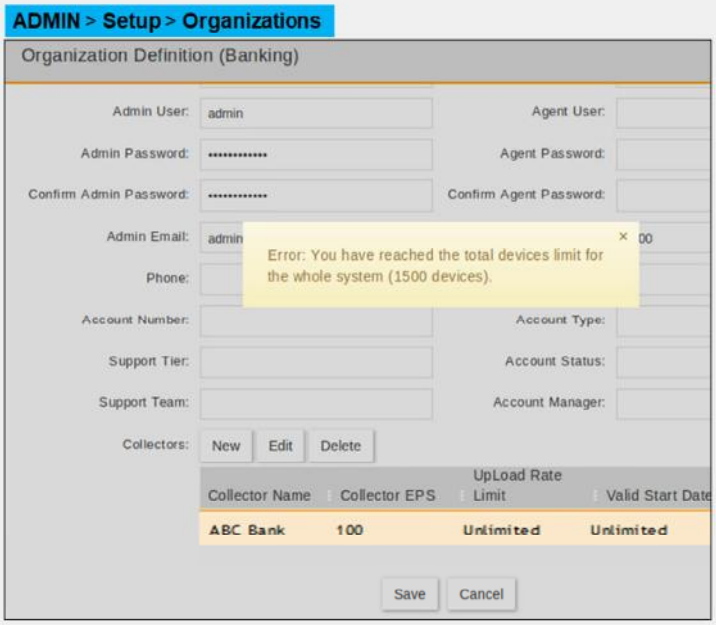


The maximum number of monitored configuration management database (CMDB) devices for the organization can be defined in the **Max Devices** field. This value is reserved from the overall system device license. The max device feature is applicable only to organizations with collectors.

DO NOT REPRINT
© FORTINET

Total Devices Limit

- This value acts as a reservation for the new organization
- Care should be taken because you cannot define a max devices value greater than the total system device license count



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

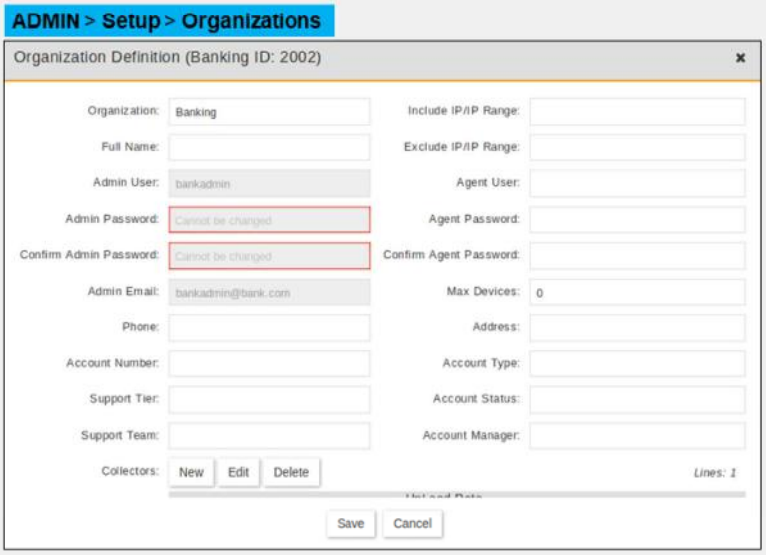
27

There is a limit to the maximum number of devices that you can assign to an organization, and that limit depends on the license that was purchased from Fortinet. You cannot exceed the total devices limit for the whole system.

DO NOT REPRINT
© FORTINET

Editing an Organization Definition

- Various fields, including the organization name, can be edited after definition



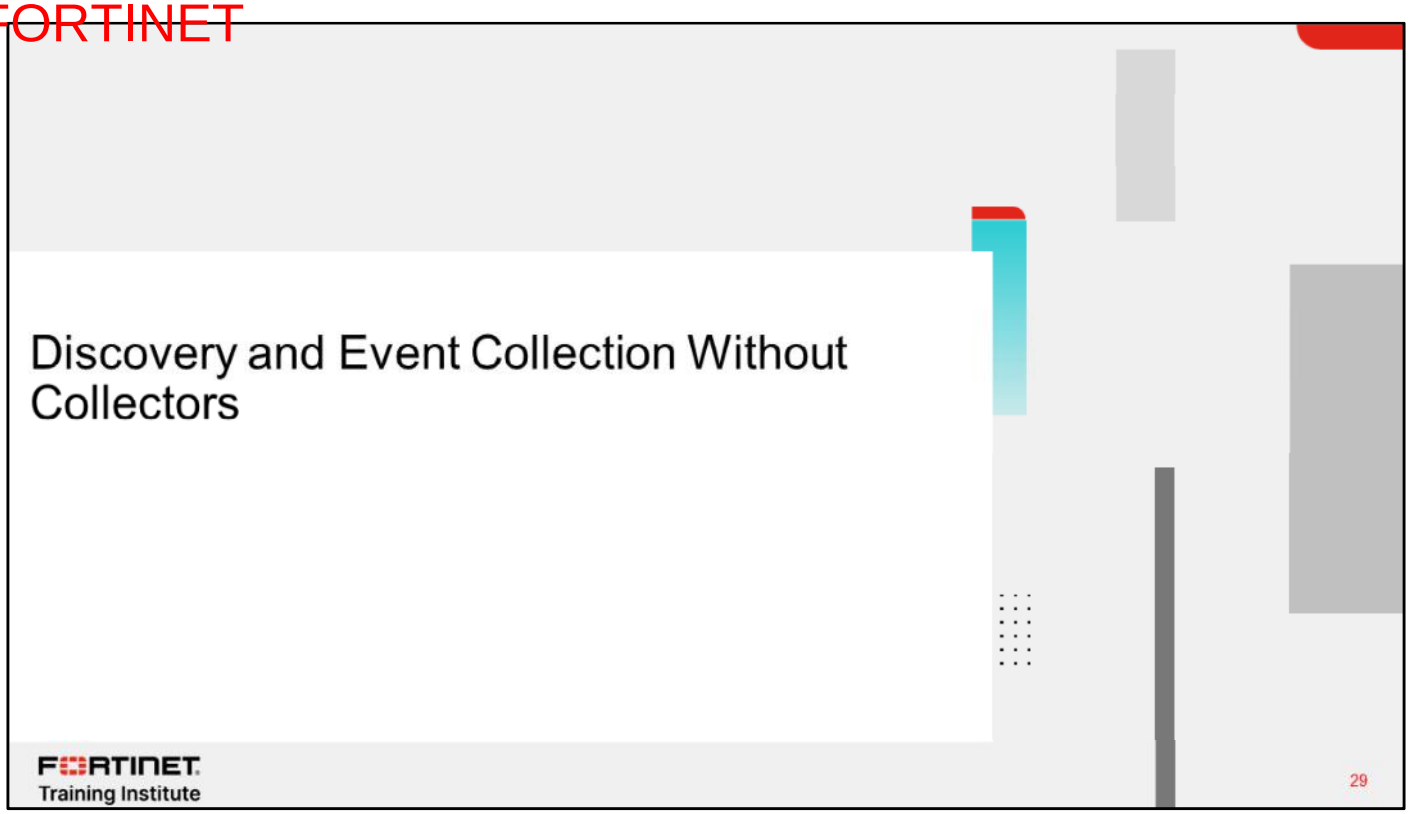
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

28

Various fields, including the organization name, can be edited after definition. However, fields such as **Admin Password**, cannot be changed.

DO NOT REPRINT
© FORTINET



In this section, you will learn about discovering devices and collecting events for customers without collectors.

DO NOT REPRINT
© FORTINET

Defining Customers Without Collectors

- Organizations are defined by a unique IP range which can be:
  - A single IP address
  - A multiple comma separated IP address
  - An IP address range



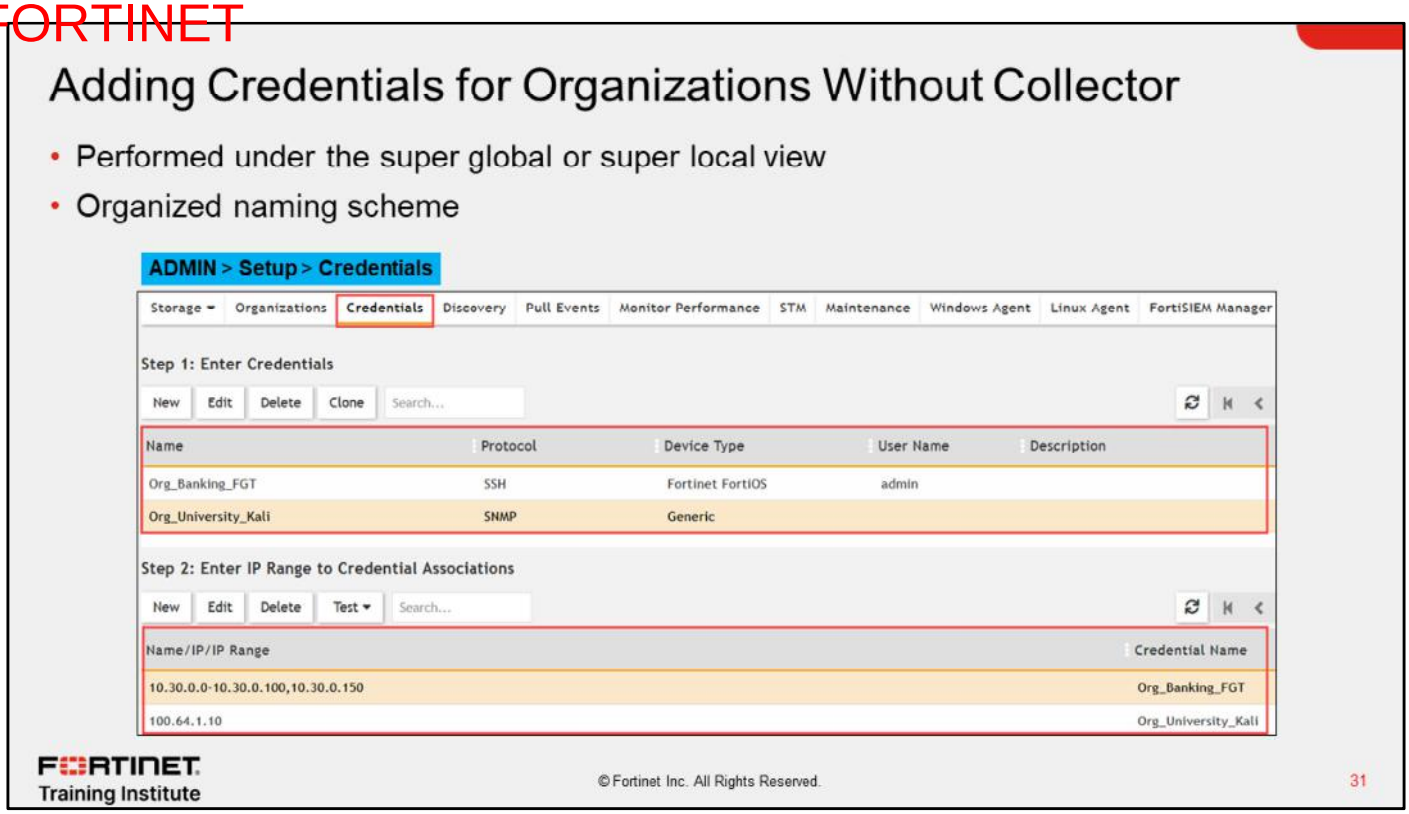
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

30

Organizations without collectors are defined by unique IP addresses, which can be a single IP address, multiple comma separated IP addresses, or an IP address range. Note that classes inter-domain routing (CIDR) notation is not supported here. If you have multiple networks defined on a device like a router, then you need to define the subnet or IP range for each interface on FortiSIEM. All interface IP addresses count during discovery, unless an exclusion is defined.

DO NOT REPRINT
© FORTINET

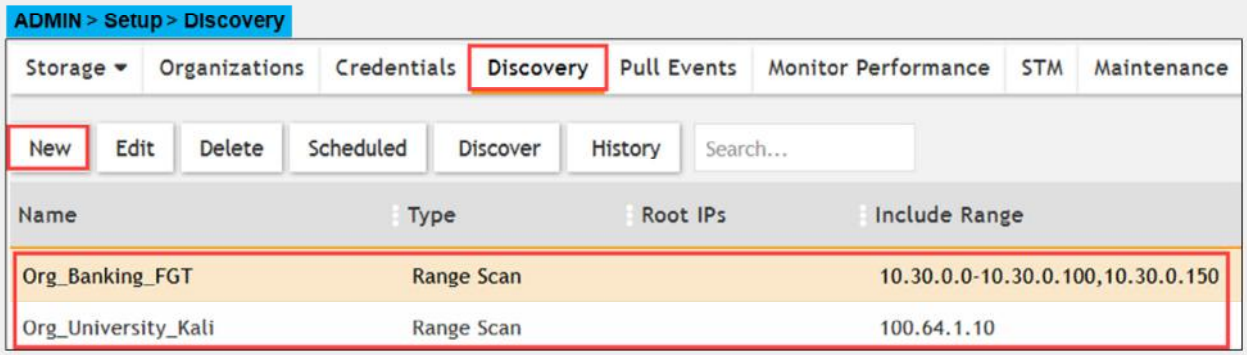


You can add credentials for organizations without collectors on the **Credentials** tab. You will need to be logged in as a super global user. It is a best practice to name credentials in an organized manner, so that you can identify the credentials for each organization.

DO NOT REPRINT
© FORTINET

Without Collector—Discovery Definitions

- Performed under the super global or super local view
- Organized naming scheme



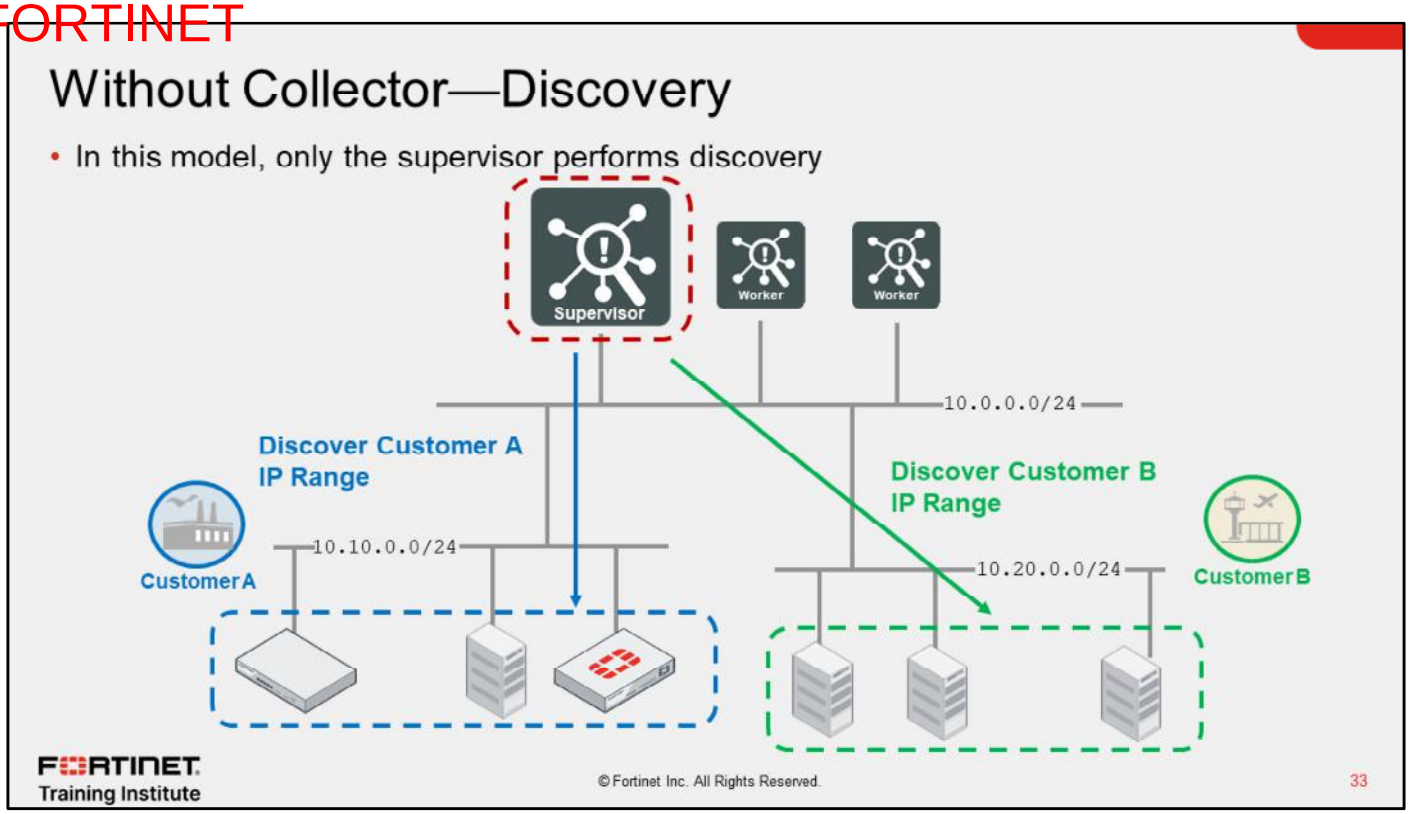
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

32

You can define discovery definitions for deployment without collectors on the **Discovery** tab. You will have to log in as a super global user. Since multiple discovery entries will be defined under the same view, you should develop a good naming scheme.

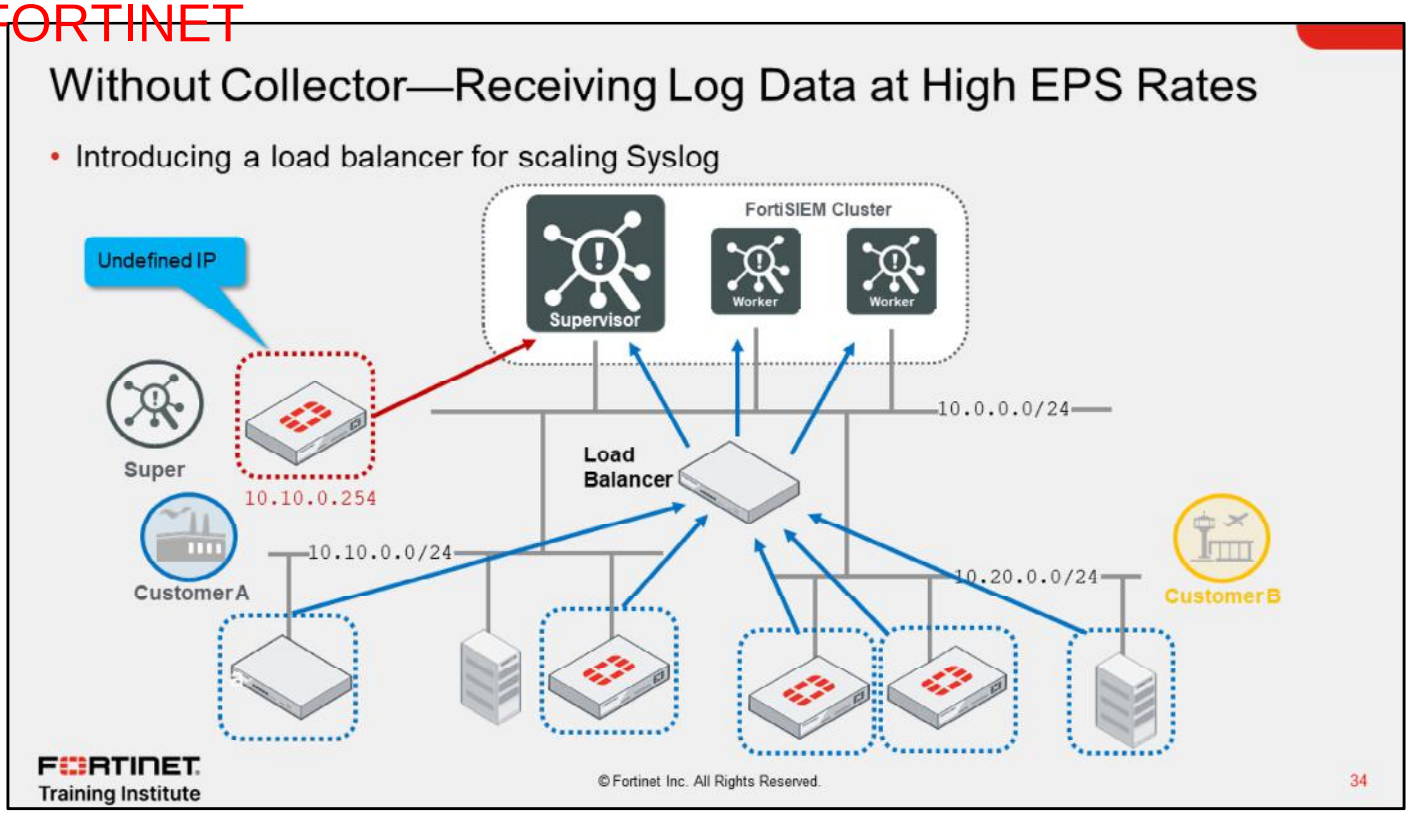
DO NOT REPRINT
© FORTINET



The supervisor node performs all discovery, if there are no defined collectors. The supervisor will automatically discover devices, applications, and users in the customer's IT infrastructure and starts monitoring them.

Any discoveries that do not match an organization IP range belong to the super organization.

DO NOT REPRINT
© FORTINET

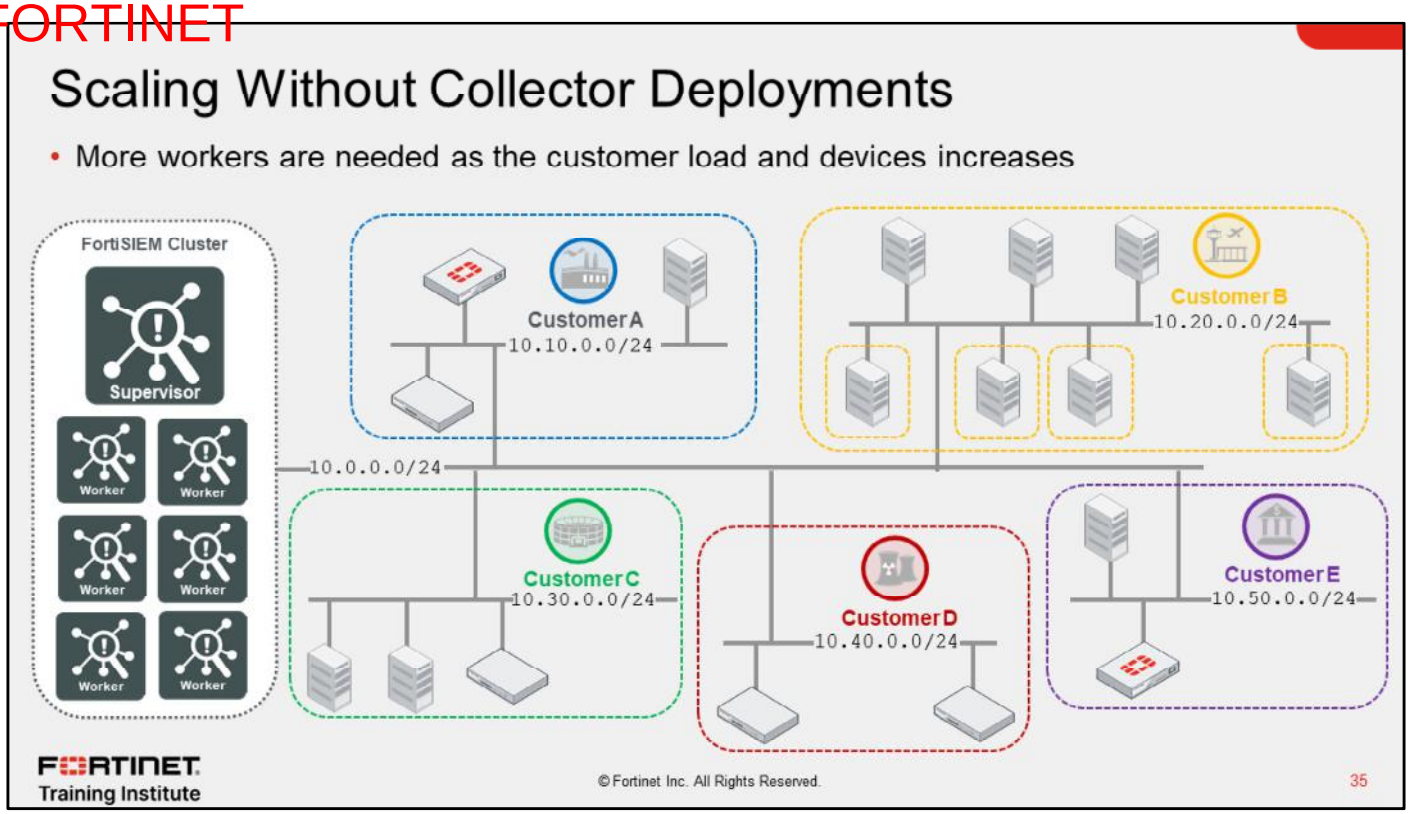


To scale a deployment without collectors, you need to add a load balancer so that logs can be balanced in an orderly way, and distributed across all workers and the supervisor, without choking the performance of any single node in the cluster. Remember that not all protocols are supported when you introduce a load balancer to the environment. Syslog and SNMP traps from the network and servers are fine, but protocols such as WMI for performance monitoring, or Windows event log collection, are actually polls from FortiSIEM itself, so those requests cannot be load balanced, and needs to go directly to the device that is being polled.

To identify the customer data, the receiving node in a FortiSIEM cluster without a collector looks up which customer the reporting IP belongs to, and tags the event with the customer name and customer ID. This way, you can filter events using the customer ID or name to perform analytics. If an incident is generated for a customer, you can identify the customer by looking at the events that generated the incident.

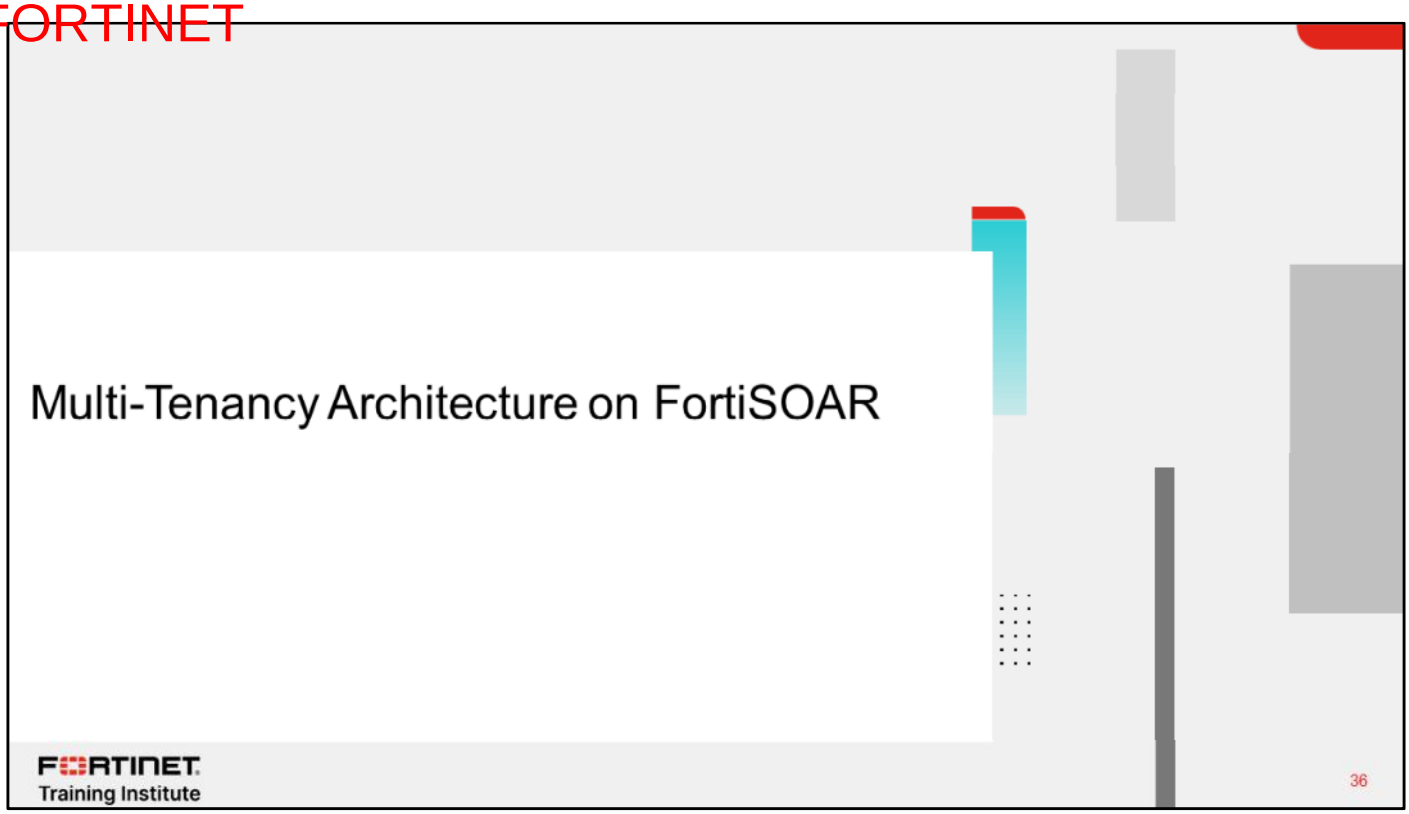
Any undefined reporting IP addresses belong to the super organization, and will be tagged as a local asset of the service provider.

DO NOT REPRINT
© FORTINET



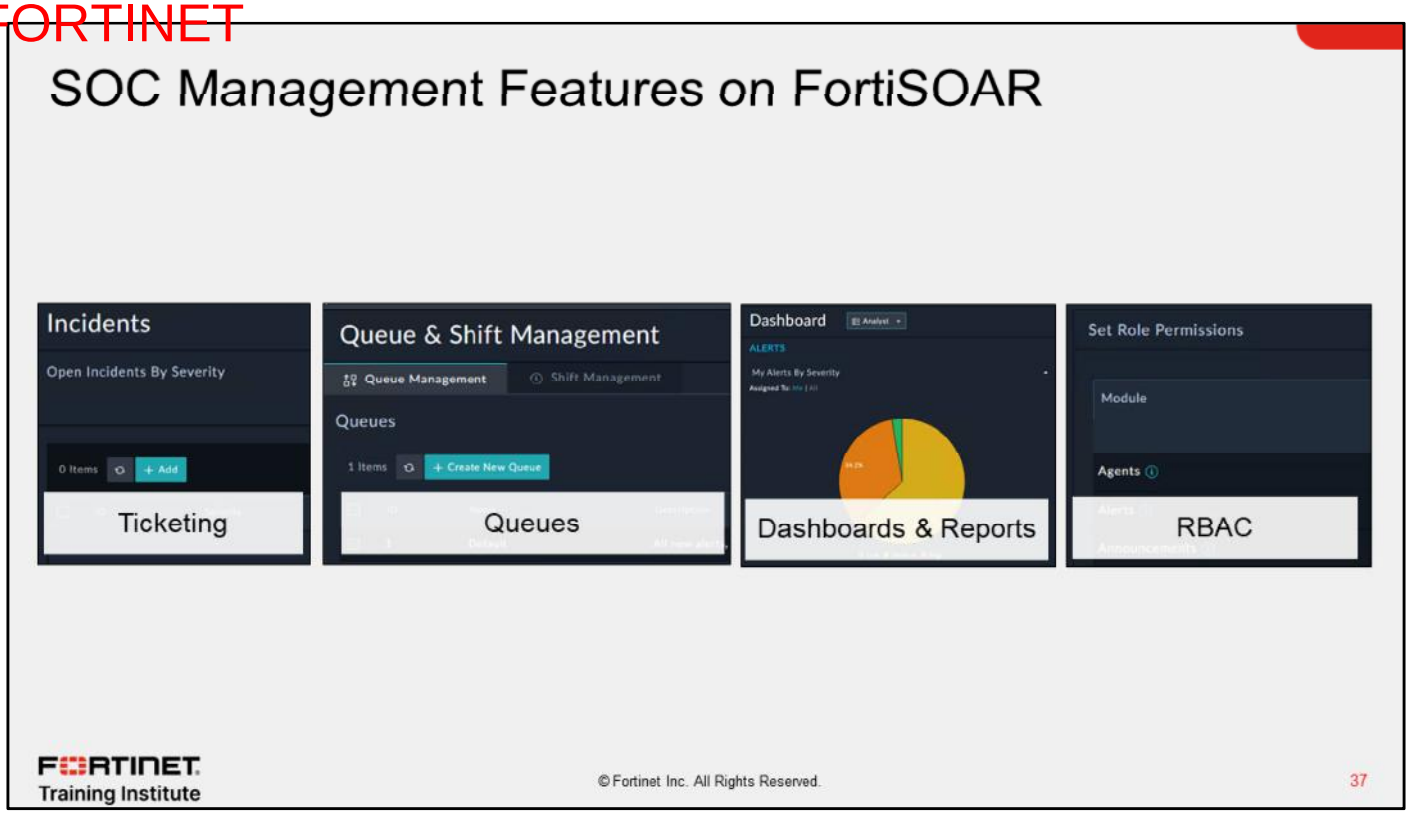
As the customer base grows and more devices are added to the customer infrastructure, the number of events per second (EPS) also increases. As the EPS increases, the resources on the FortiSIEM cluster can be under strain. More workers need to be added to the FortiSIEM cluster to balance the load.

DO NOT REPRINT
© FORTINET



In this section, you will learn about multi-tenancy architecture on FortiSOAR.

DO NOT REPRINT
© FORTINET



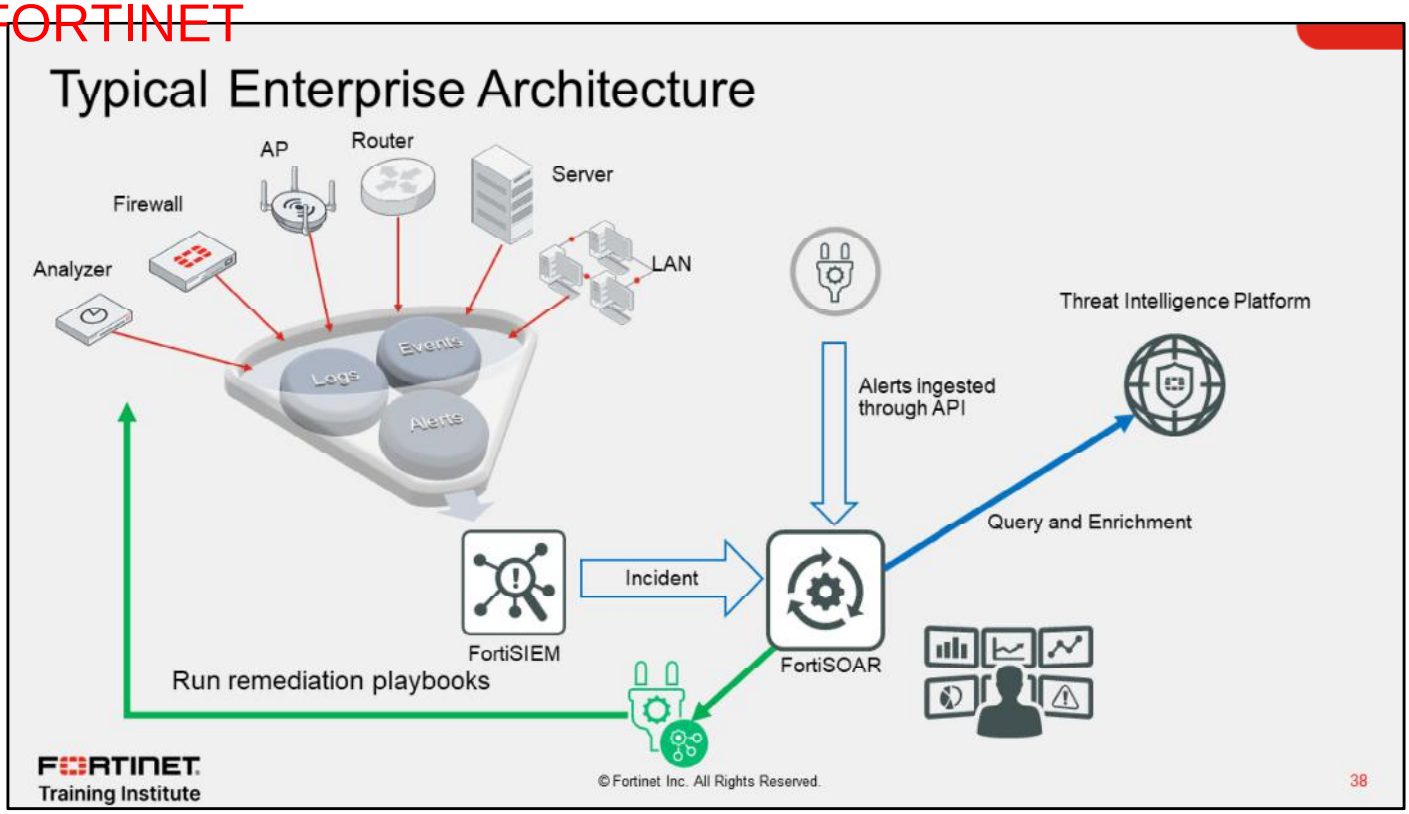
When you escalate an alert on FortiSOAR, it becomes an incident. Every incident on FortiSOAR can be considered as a ticket that an analyst will need to work on until they close the incident.

The FortiSOAR queue management feature provides you with an overview of the work (records) that must be completed and enables you to assign pending work to users. You can also reassign assignments in case of absence or analyst shift changes.

Administrators can create applicable dashboards throughout the application. The dashboards are assigned to users based on their roles. For example, you can create a dashboard that displays alerts that are critical and high, and then assign them to users who have the role of handling alerts. Users can prioritize their work by looking at their dashboard, which displays critical and high alerts.

FortiSOAR gives you the option to assign levels of accessibility to users using role-based access control (RBAC) combined with team membership. You can grant users access to specific modules in FortiSOAR based on their role permissions. Users exercise their permissions in conjunction with their team membership(s).

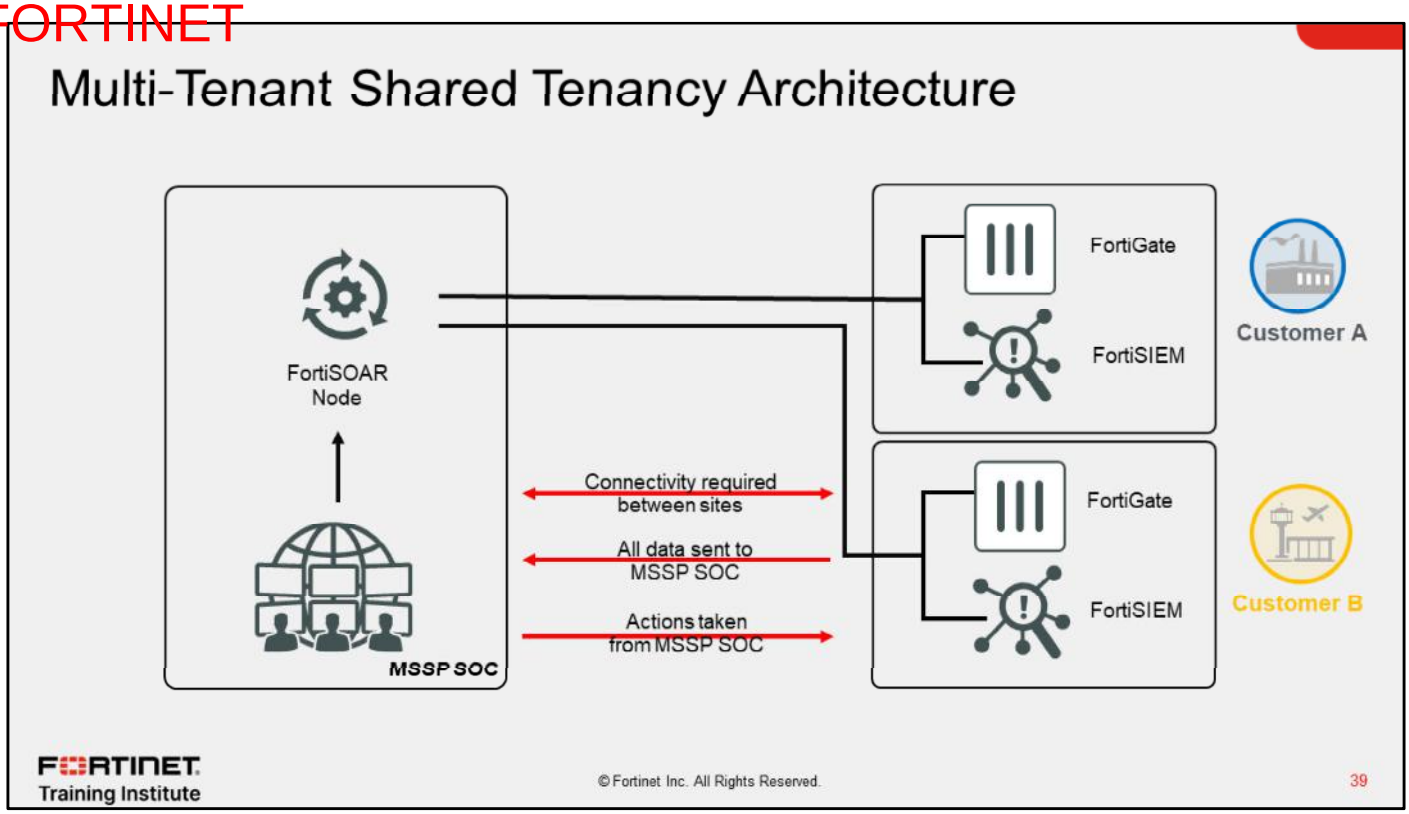
DO NOT REPRINT
© FORTINET



In an enterprise architecture, there is one single instance of FortiSOAR. You can ingest data using connectors into FortiSOAR from various devices in your infrastructure. When available, it is recommended that most logs be sent through the SIEM, rather than through direct connectors. This ensures the SIEM is a central point of log aggregation and can be used for analytics and reporting.

After you integrate your FortiSIEM, or any other SIEM solution, with FortiSOAR, incidents generated by FortiSIEM are ingested by FortiSOAR. Every incident that is sent from FortiSIEM to FortiSOAR is a unique record on FortiSOAR. You can run remediation playbooks from FortiSOAR against those incidents, and perform remediation action on the target devices. For example, if the logs that are sent by FortiGate to FortiSIEM generate an incident that indicates an external malicious actor is trying to access corporate resources, then FortiSOAR evaluates that incident and rates that external IP against various threat intelligence platforms. If the IP is malicious, then you can run a remediation playbook to take action against that IP address. In the scenario shown on this slide, the solution is to block that IP on the FortiGate firewall. The playbook can automatically log in to FortiGate and put that IP on the quarantine list for an indefinite period of time.

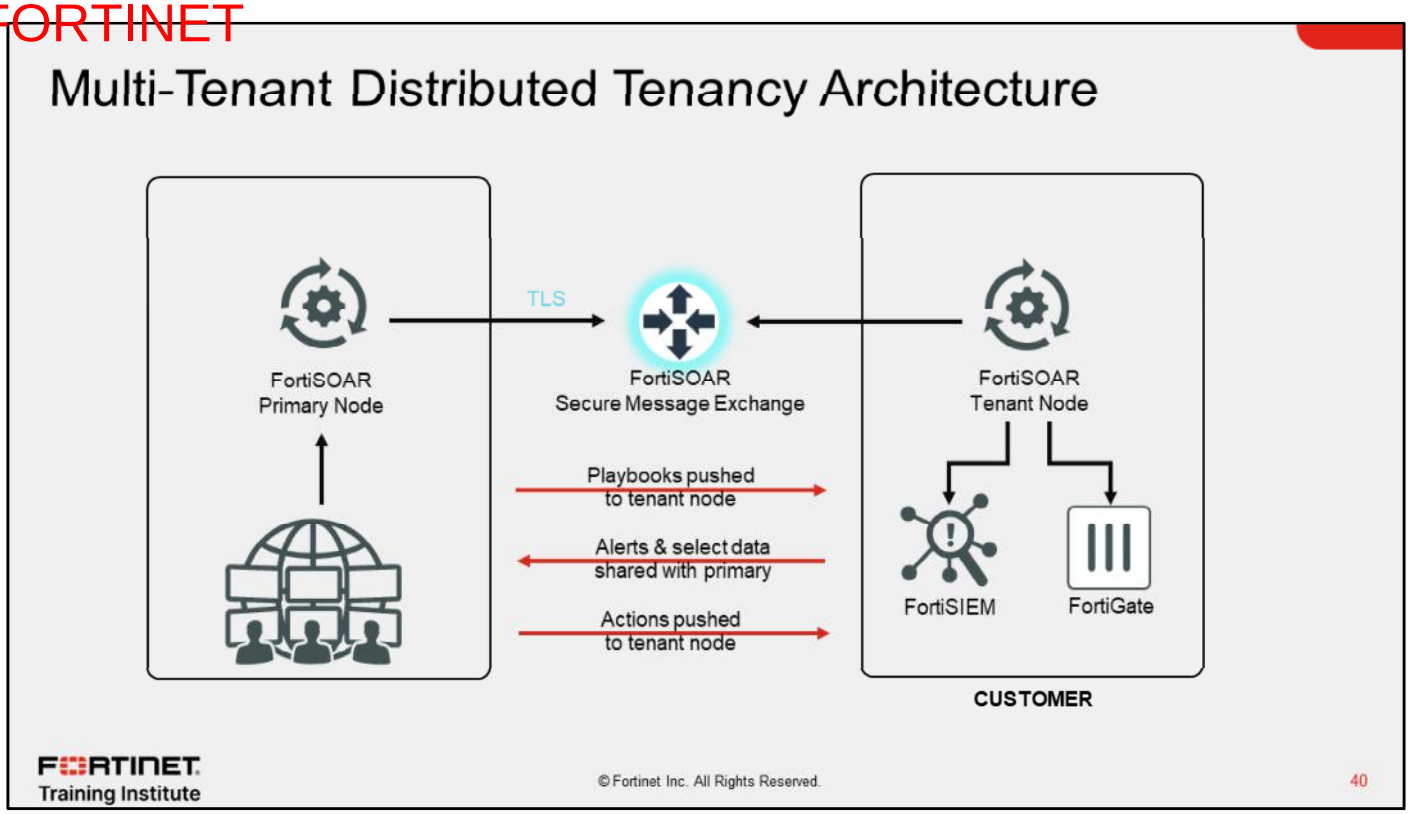
DO NOT REPRINT
© FORTINET



In the case of shared tenancy, tenants share the same system as the primary device; that is, tenants are local, but with restricted access on the system. The SOC team provides cybersecurity monitoring and management to various tenants in a single FortiSOAR instance. The shared tenancy model ensures that the data belonging to different tenants is segregated, and data access is controlled using RBAC. Therefore, a tenant can view only their own data or record, and not the data of other tenants.

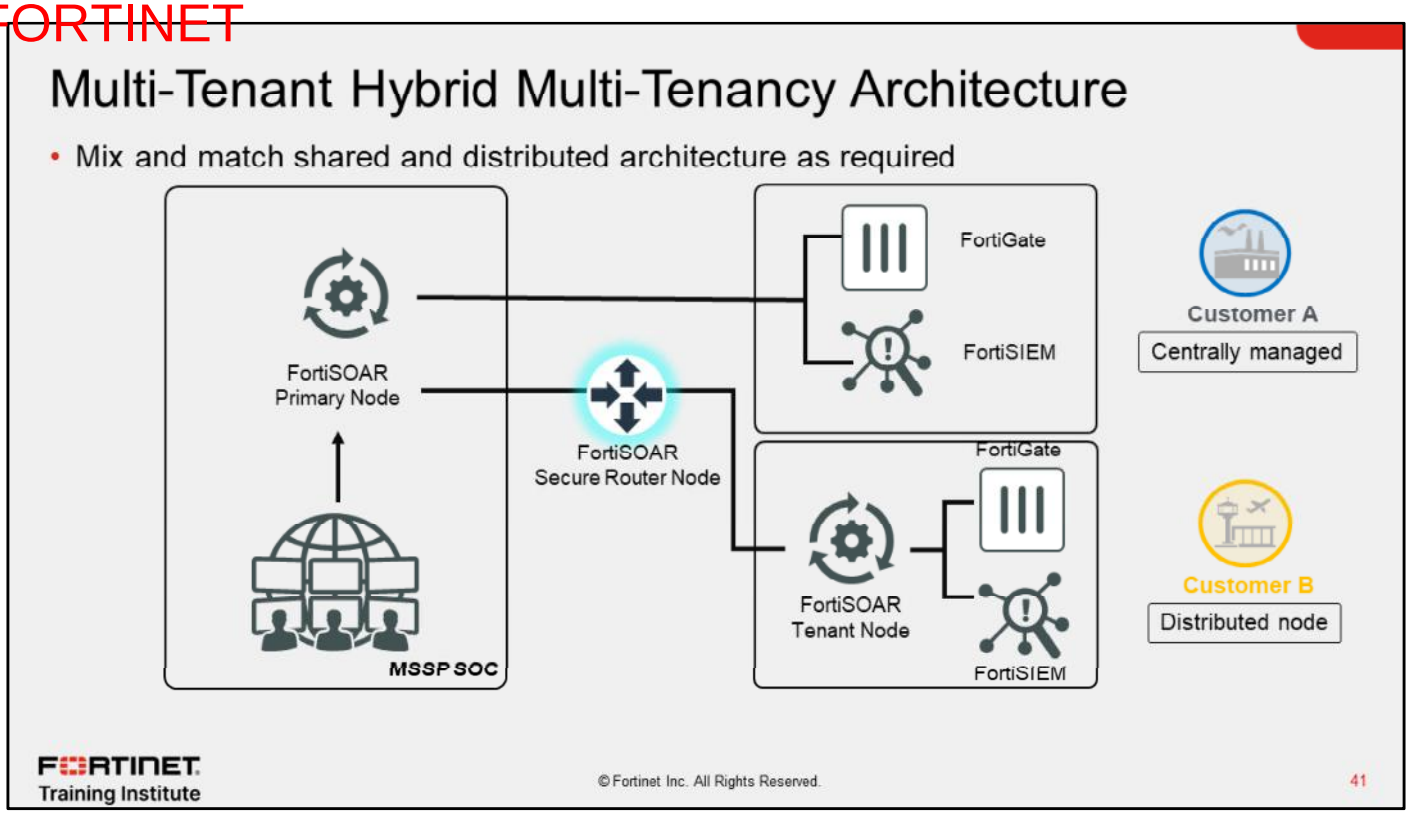
You can give each tenant their own login, which they can use to view their dashboards, report, check the actions taken on their records, check their SLA management, and so on.

DO NOT REPRINT
© FORTINET



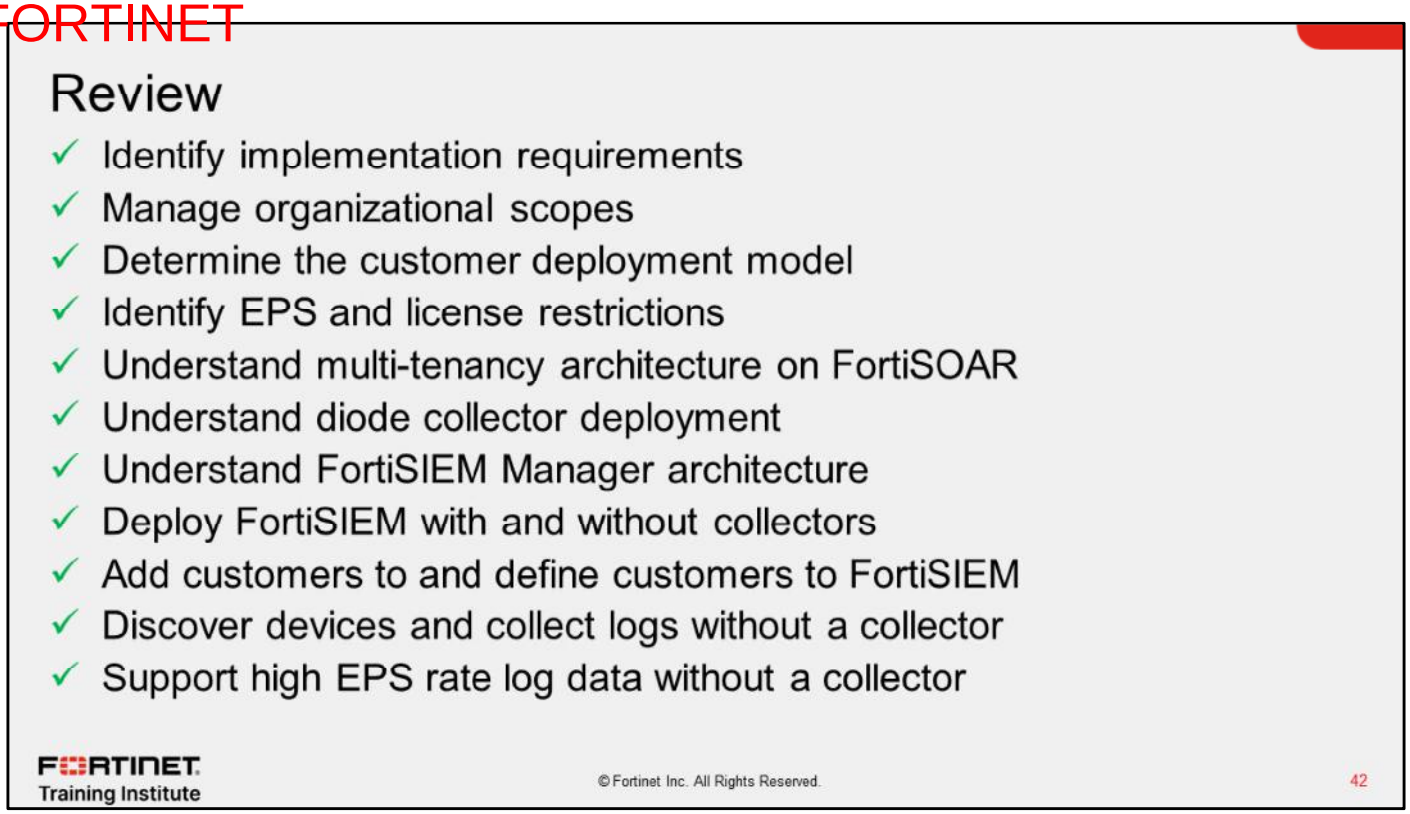
In the case of a distributed tenancy model, the tenant node instance of FortiSOAR is remote and every tenant has their own instance of FortiSOAR. The primary FortiSOAR node resides at the MSSP location and communicates with the tenant node through a secure channel. Tenant data remains in the tenant environment, and they control how much data they want to share with the primary node. All sensitive information stays with the tenant node. Since the actual workflow execution happens at the tenant node itself, the primary node requires only the summary of information to help identify what investigations should run. The primary node pushes any action that needs to be executed to the tenant node. Similarly, any playbook that needs to be executed is pushed by the primary node to the tenant node.

DO NOT REPRINT
© FORTINET



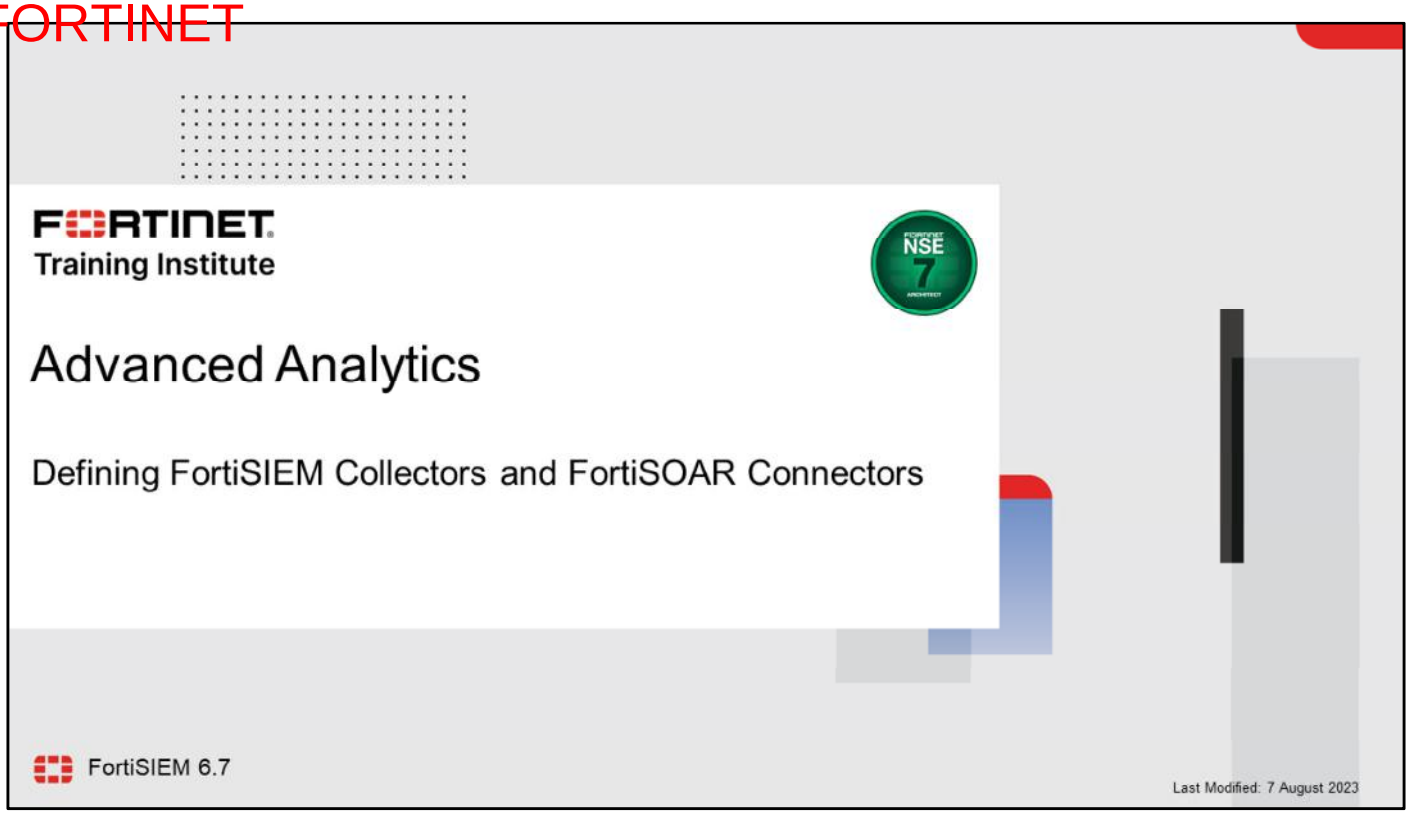
In a multi-tenant hybrid model, the primary node centrally manages some customers. If the customer is managed by the primary node, then there is no requirement for a tenant node for that customer. Then there are other customers who use a distributed method. For those customers, you must install a tenant node.

DO NOT REPRINT
© FORTINET



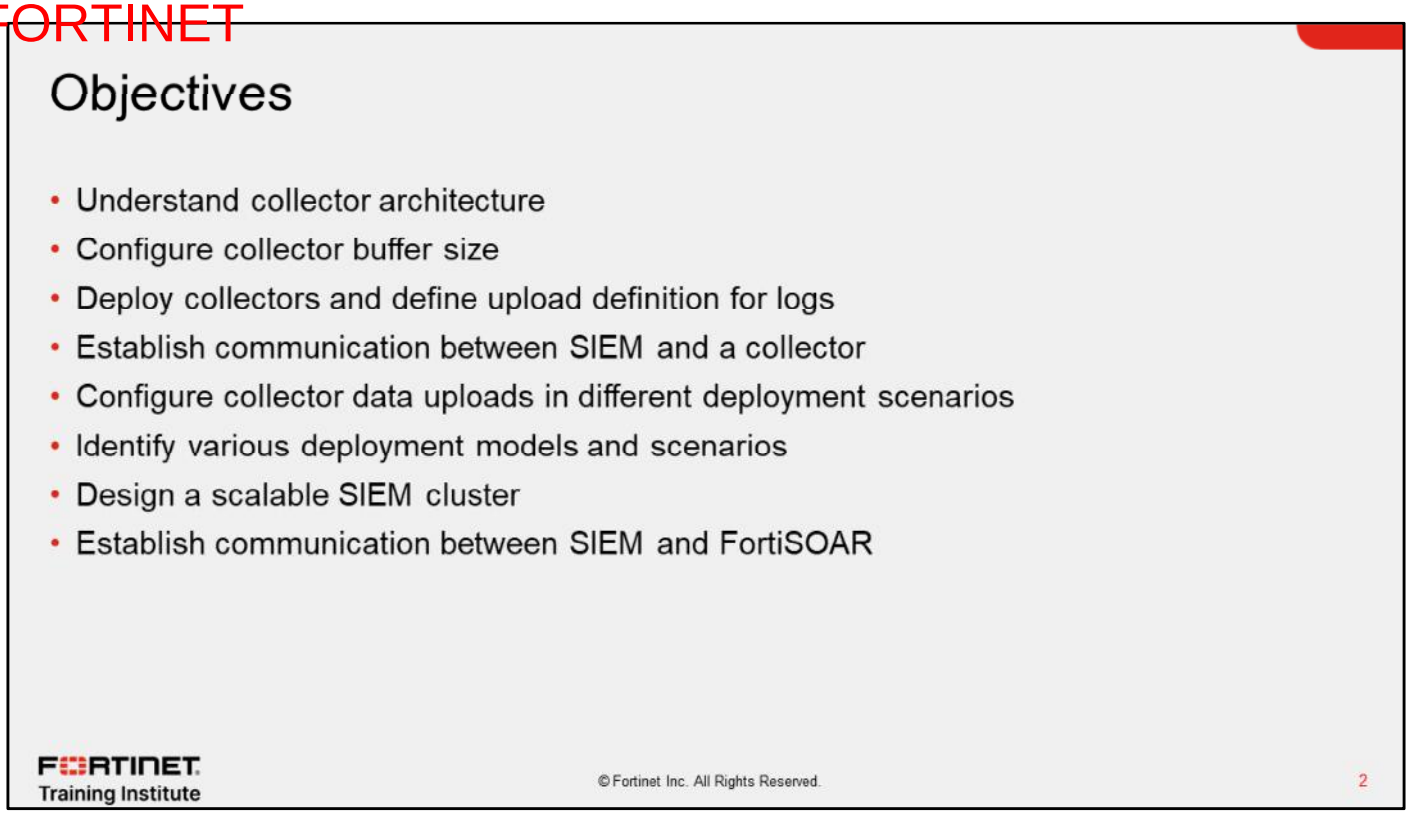
By mastering the objectives covered in this lesson, you learned how to deploy FortiSIEM in a multi-tenant environment and identify key implementation requirements. You also learned about deploying FortiSIEM with and without collectors, and how FortiSIEM collect logs without a collector.

DO NOT REPRINT
© FORTINET



In this lesson, you will learn about defining and deploying collectors. You will also learn about collector operation, and scaling collectors as organizations grow. Additionally, you will learn to ingest data from FortiSIEM into FortiSOAR.

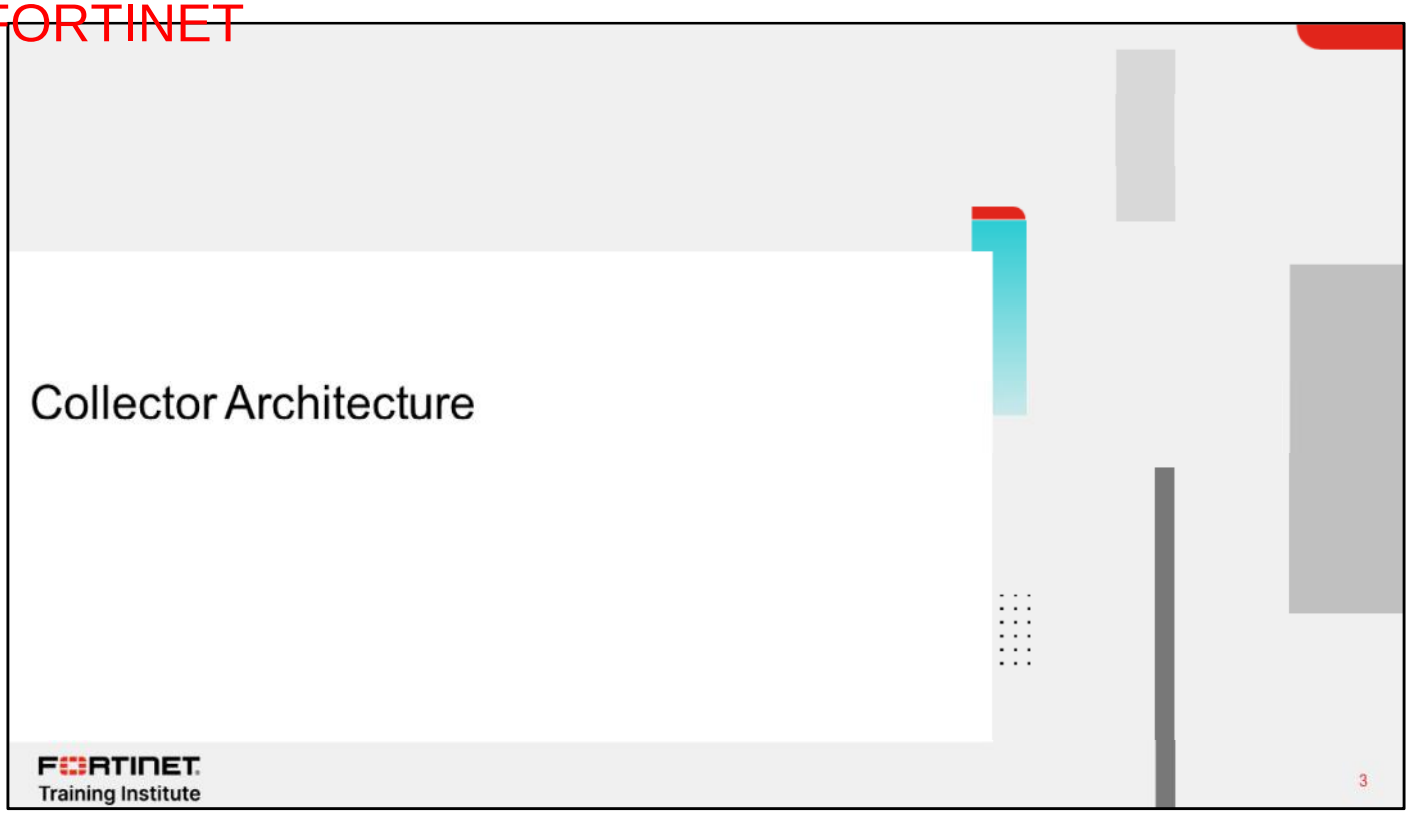
DO NOT REPRINT
© FORTINET



After completing this lesson, you should be able to achieve the objectives shown on this slide.

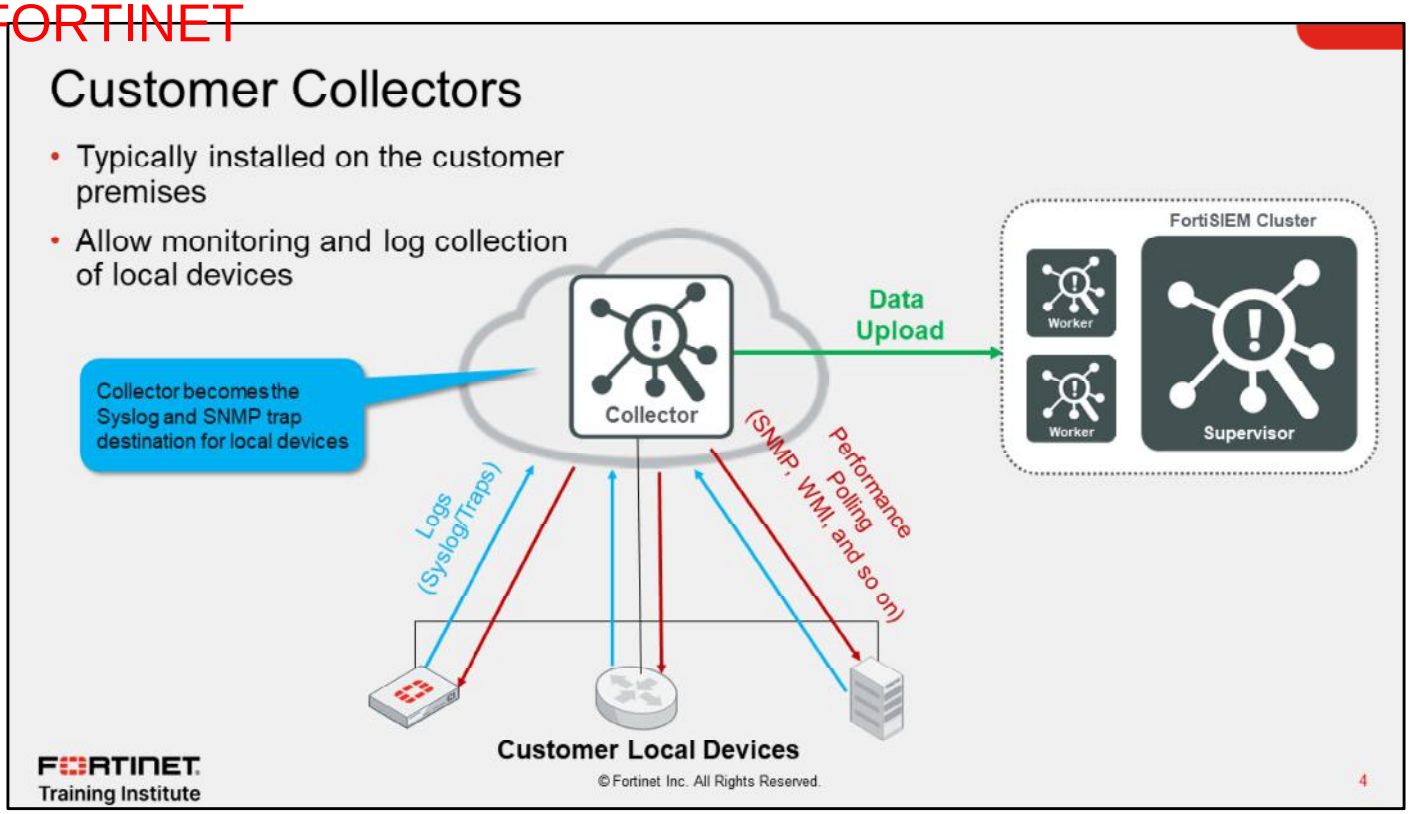
By demonstrating competence in understanding collector architecture, you will be able to configure collector buffer sizes, define upload definitions for logs, and establish communication with supervisor nodes and collector nodes.

DO NOT REPRINT
© FORTINET



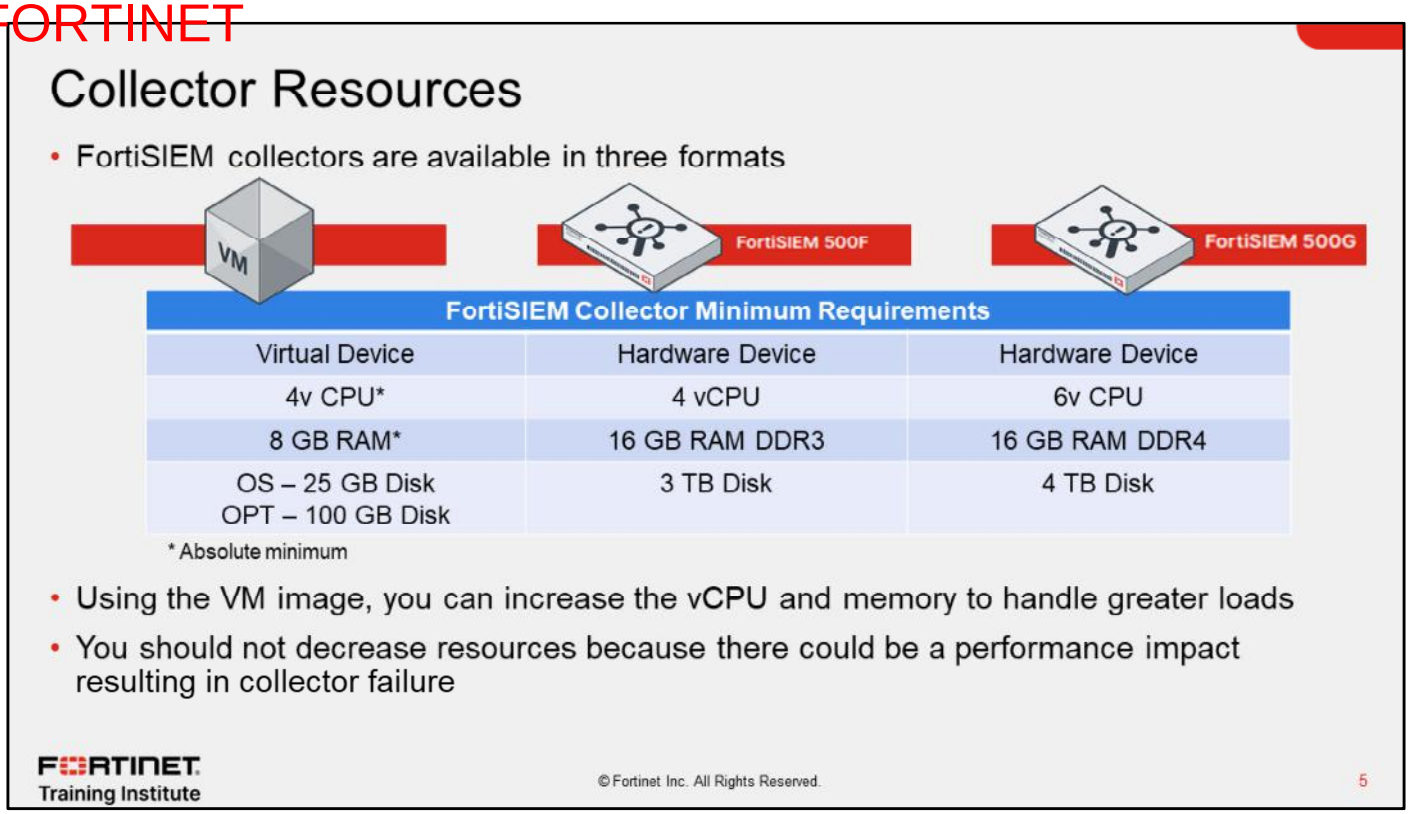
In this section, you will learn about collector architecture and various processes involved in collector architecture.

DO NOT REPRINT
© FORTINET



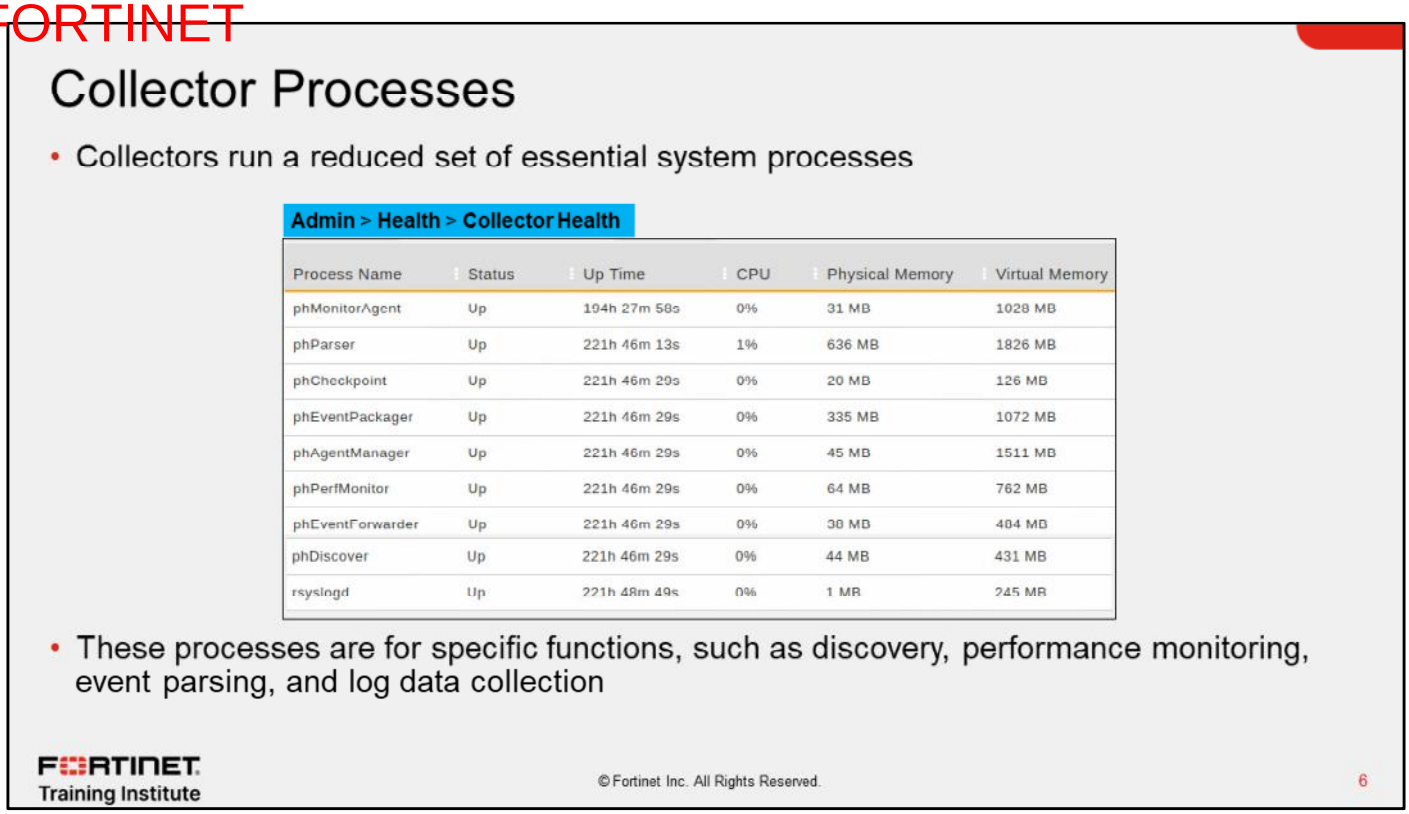
A collector enables FortiSIEM to collect logs and performance metrics from geographically disparate networks. Data collection protocols, such as SNMP and WMI, are often chatty, and the devices may be reachable only from the supervisor node through the internet, through a firewall. The Syslog protocol, especially over UDP, is unreliable and unsecure. You can deploy a collector behind the firewall to solve these issues. The collector registers with the FortiSIEM supervisor node and then receives commands from the supervisor regarding discovery and data collection.

DO NOT REPRINT
© FORTINET



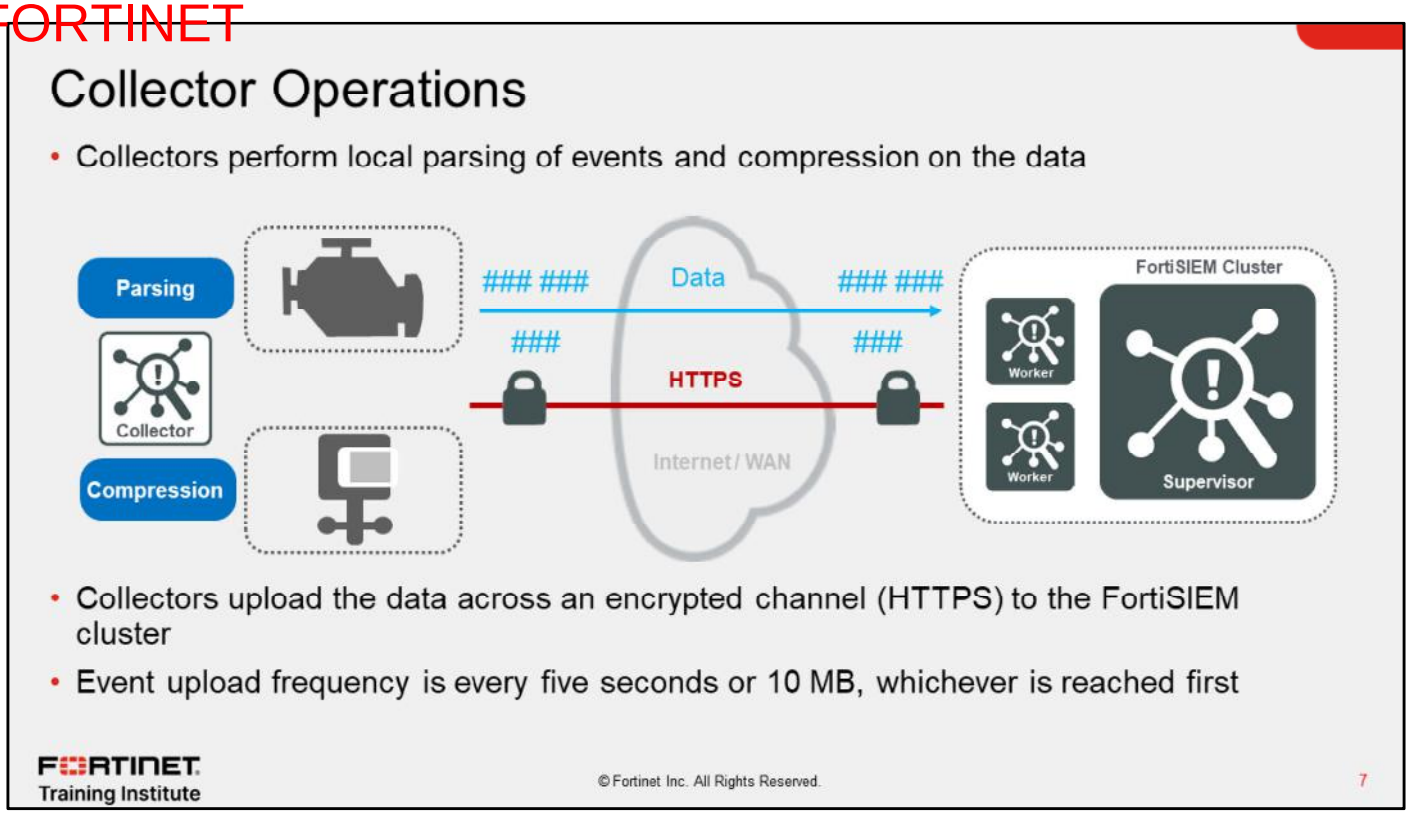
Collectors are available as virtual devices and hardware devices. If you use the VM image, you can increase the vCPU and memory to handle greater loads. FortiSIEM becomes an essential core tool in your security operations center (SOC) environment, and its use will likely grow over time as the corporate infrastructure grows. Designing the FortiSIEM deployment to accommodate this increase in use over time helps to ensure continued system performance and minimizes the need to re-architect the system in the future. For that reason, you should not decrease resources because there could be a performance impact resulting in collector failure.

DO NOT REPRINT
© FORTINET



You can monitor the status of the collectors on the page shown on this slide. Collectors run a reduced set of system processes, compared to a supervisor or a worker. The processes that run on a collector are for specific functions, such as discovery, performance monitoring, event parsing, and log data collection.

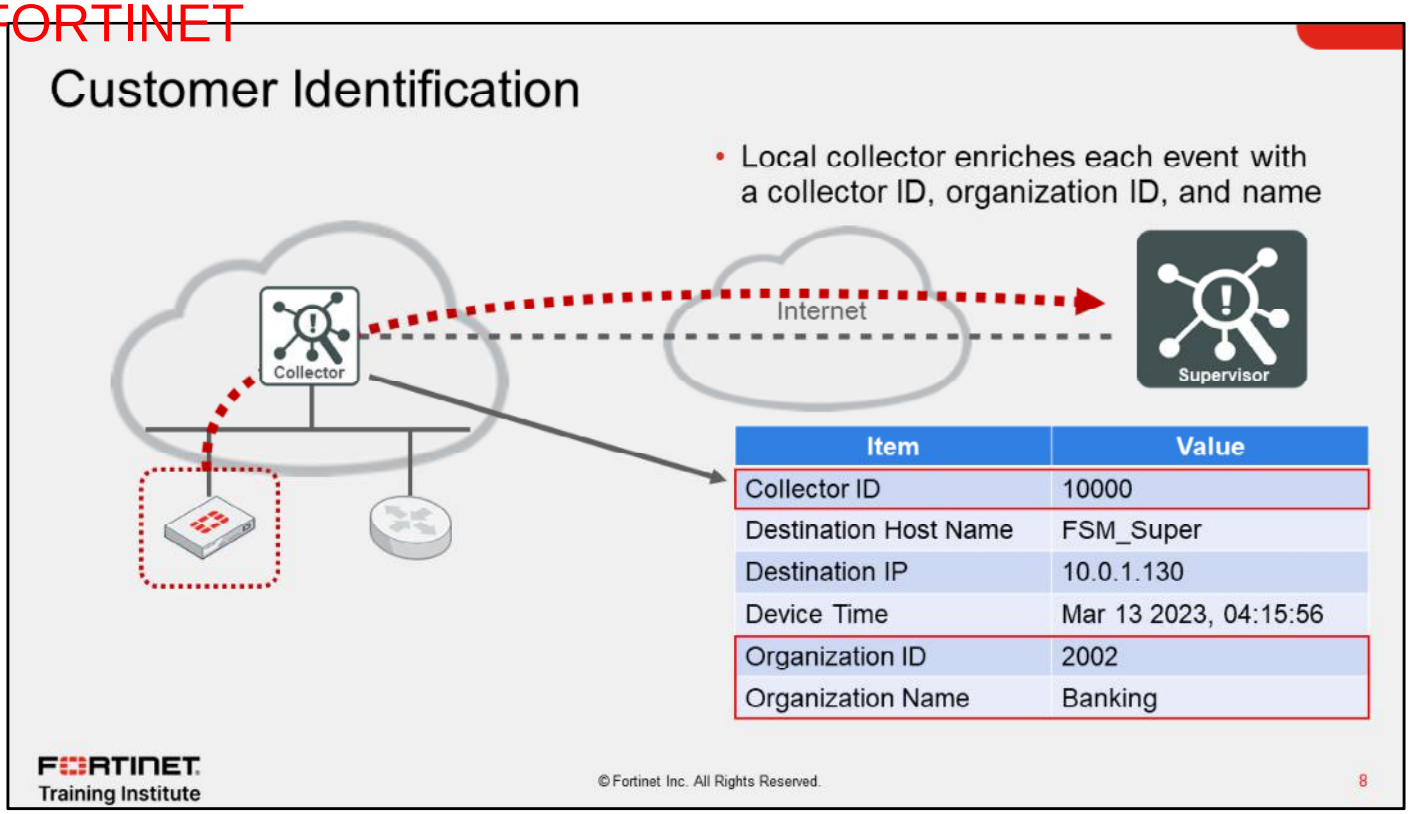
DO NOT REPRINT
© FORTINET



The collector parses the logs and forwards the compressed logs to the supervisor or worker nodes over an encrypted HTTPS channel. The collector also buffers the logs locally for a period of time, if the network connection to the supervisor or worker is not available.

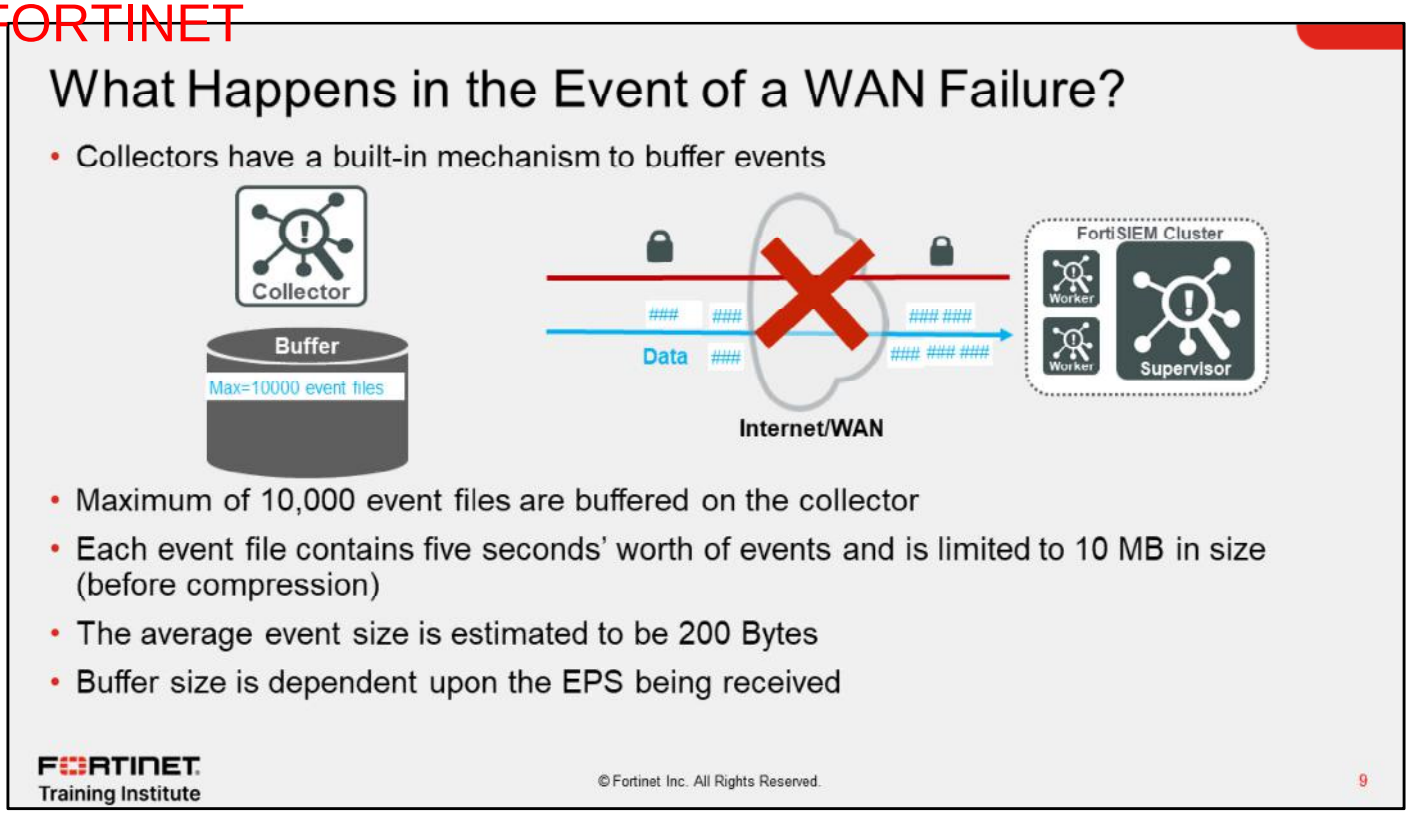
The collector uploads events every five seconds or 10 MB, whichever is reached first. FortiSIEM is a real-time system, so it attempts to minimize event delay and optimize the send. It is five seconds for low EPS environments, and 10 MB for high EPS environments. The compression ratio is 8:1, which is accomplished by standard zlib, which is an algorithm used for data compression.

DO NOT REPRINT
© FORTINET



The collector enriches each event with a collector ID, organization ID, and name. The collector name is not added during the event enrichment process. Because the collector performs the enrichment process, it saves resources on the supervisor node. Having a collector in your environment makes it convenient for MSSP administrators to query for logs and incidents related to your organization, because all logs being generated from your organization are tagged with the unique collector ID, organization ID, and organization name.

DO NOT REPRINT
© FORTINET



Collectors have a built-in mechanism to buffer events in case there is WAN link failure and you lose connection to the FortiSIEM cluster. Buffer size is dependent upon the EPS being received. Collectors will ship the logs automatically when the link becomes available again.

By default, there are a maximum of 10,000 event files that are buffered on the collector. Each event file contains five seconds' worth of events and is limited to 10 MB in size before compression. The average event size is estimated to be 200 Bytes, which depends on the device type and event.

DO NOT REPRINT
© FORTINET

Collector Buffer Size

- Estimated buffer time:

EPS	Estimated Buffer Time	EPS	Estimated Buffer Time
500	55.55 Hours	4000	6.94 Hours
1000	27.77 Hours	5000	5.55 Hours
2000	13.88 Hours	6000	4.62 Hours
3000	9.25 Hours	7000	3.96 Hours

- Estimated amount of time to reach the maximum buffer size on the collector:
  - Default bytes per second = 10 MB/5 sec = 2 MBPS
  - Total buffer storage = 100,000 MB
  - Total time to fill that storage = 100,000/2 = 50,000 seconds
  - Total time in hours = 13.8

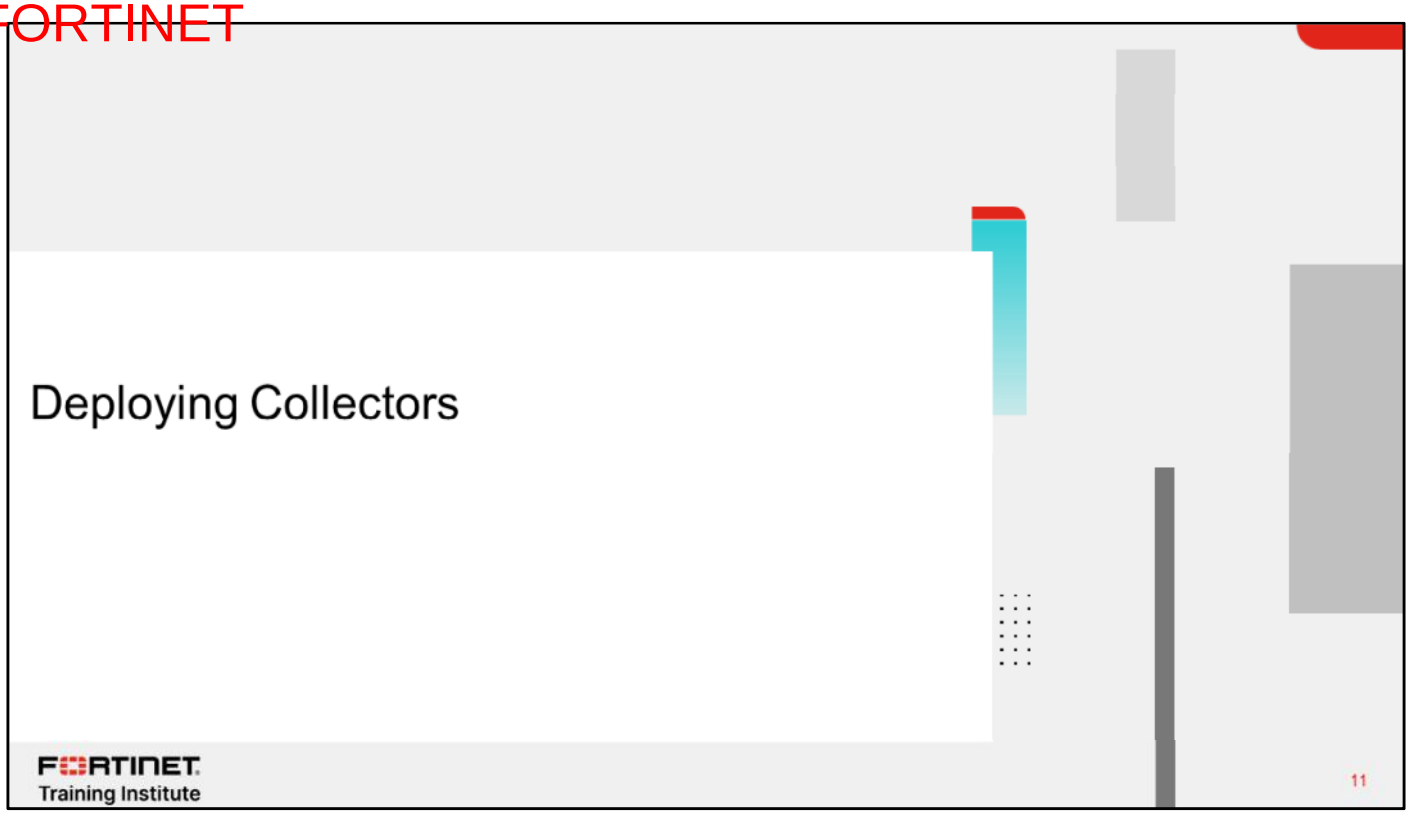
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

10

This slide shows an example of the estimated time that it would take for the collector to reach the maximum buffer size for a 2000 EPS license. As you increase the EPS, the buffer storage fills up. The collector drops events when the buffer is full.

DO NOT REPRINT
© FORTINET

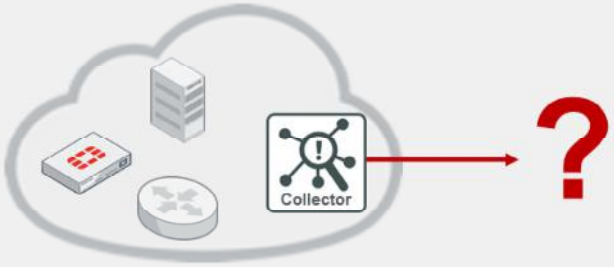


In this section, you will learn about deploying collectors in a multi-tenant environment.

DO NOT REPRINT
© FORTINET

Establish Collector Connectivity With a Supervisor

- When initially deployed, the collector is in an unconfigured state
- During the collector registration process, you must gather the following information about the collector:
  - The identity of the customer it belongs to
  - The location where it will upload data
  - The length of time that it is valid
- You can define all the required information during the collector registration process



FORTINET Training Institute

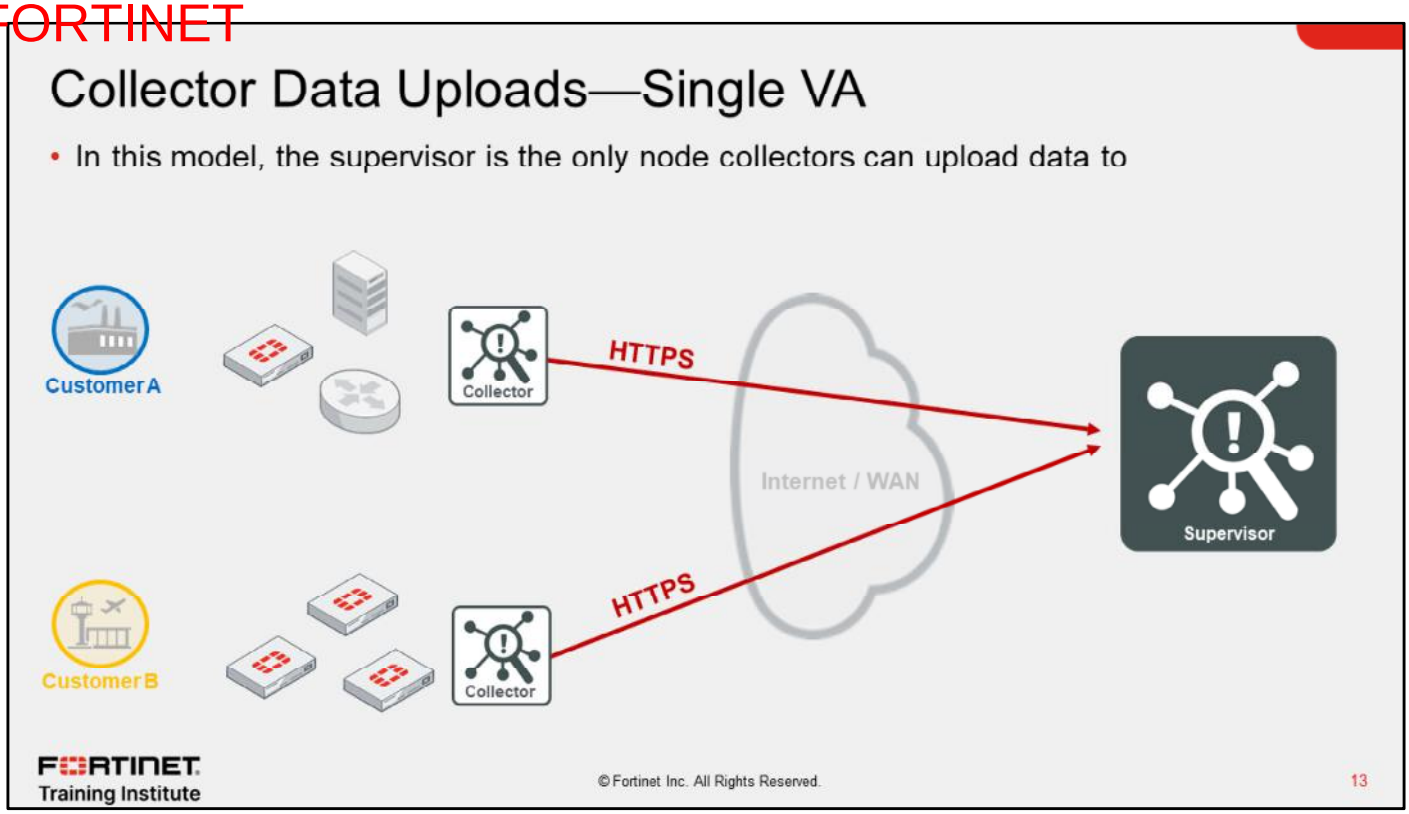
© Fortinet Inc. All Rights Reserved.

12

When a collector is initially deployed, it is in an unconfigured state. To deploy the collector, you must know which customer it belongs to, where it will upload data, and how long it is valid.

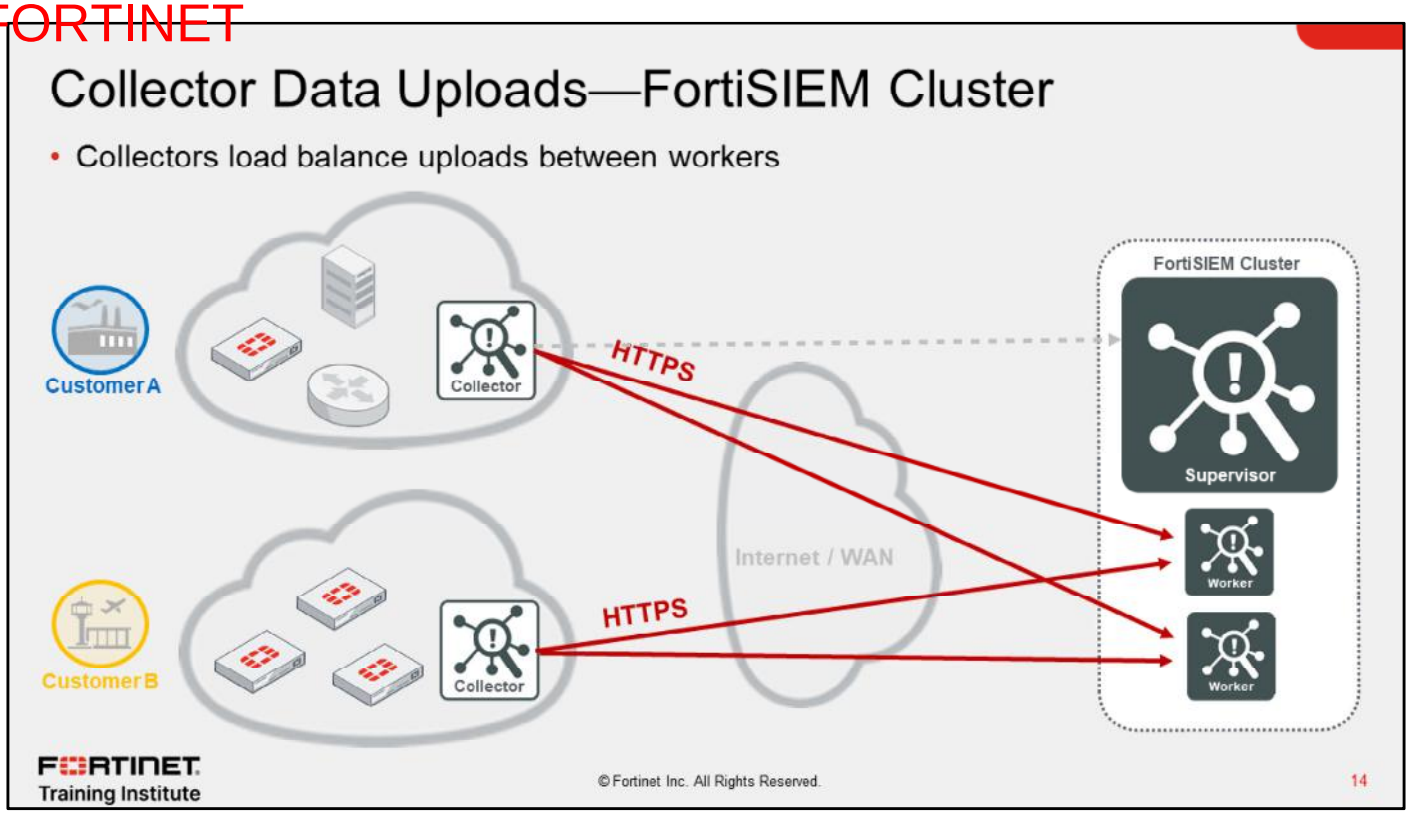
You can define all the required information during the collector registration process.

DO NOT REPRINT
© FORTINET



In a single virtual appliance (VA) setup with one supervisor node, specify the supervisor node as the data upload destination for collectors. This is not recommended in larger setups because it can potentially overload the single supervisor node.

DO NOT REPRINT
© FORTINET

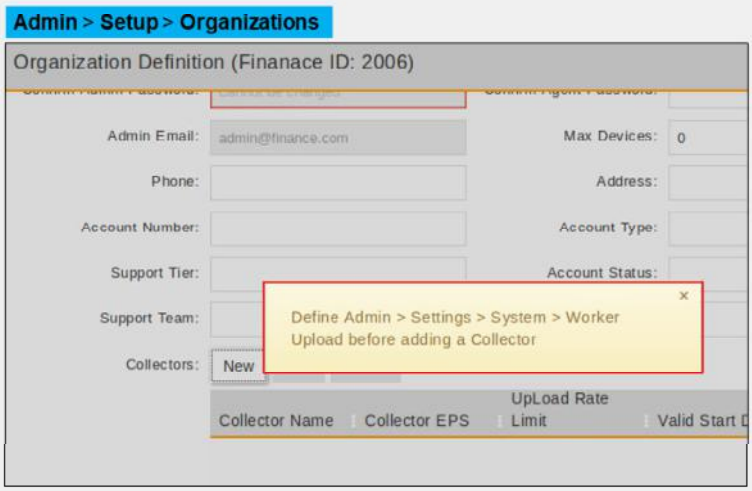


In a FortiSIEM cluster, you can specify one or more worker nodes using the worker IP addresses. The collectors load balance across the specified worker nodes. In this manner, streaming analytics, like inline reports and rules, are distributed over the worker nodes. The best practice is to upload all the data to the worker nodes and leave the supervisor node for performing other important tasks.

DO NOT REPRINT
© FORTINET

Worker Upload Settings

- You cannot define collectors for an organization before you set the worker upload address
- Collectors receive this information during registration and this value tells the collector where to upload data to



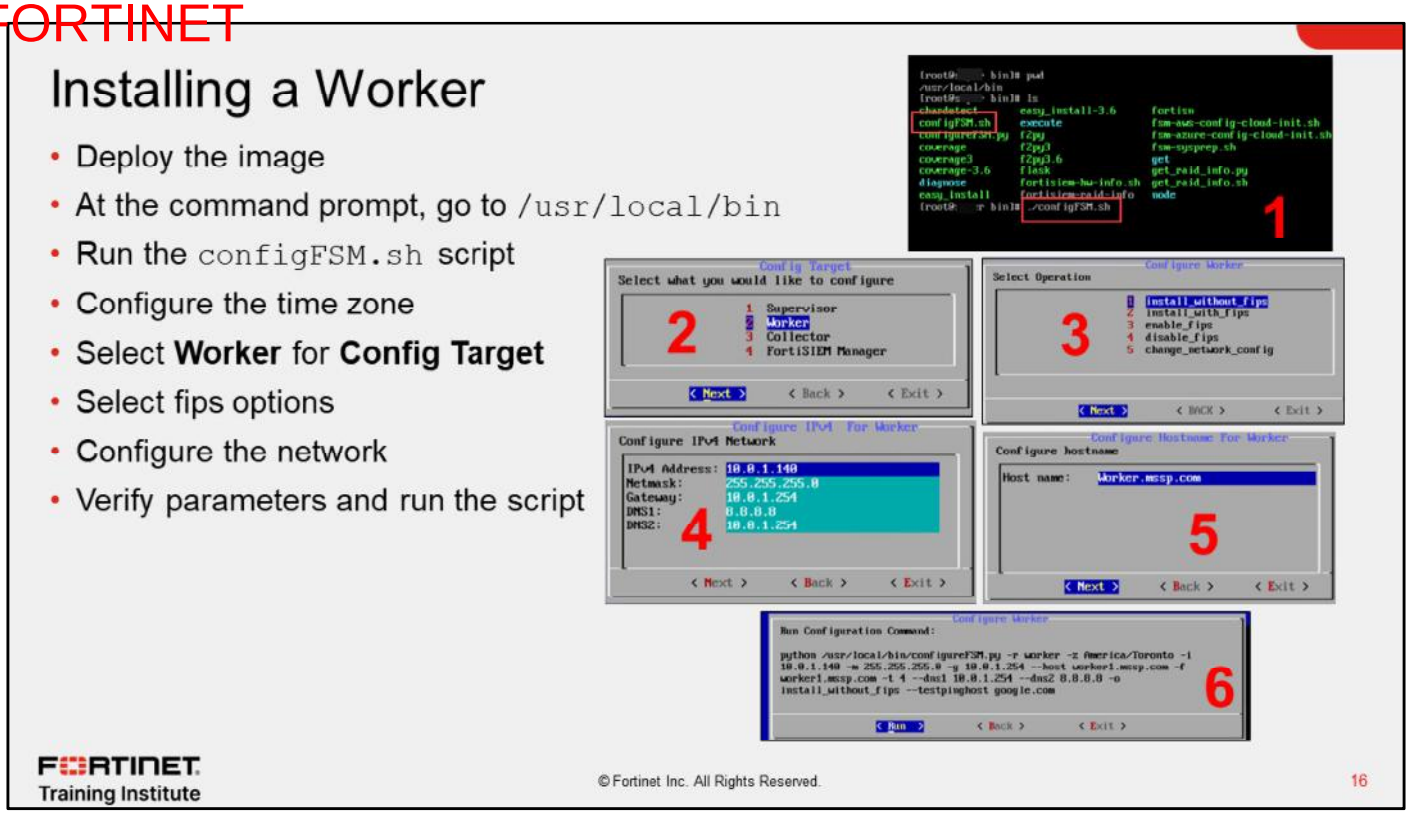
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

15

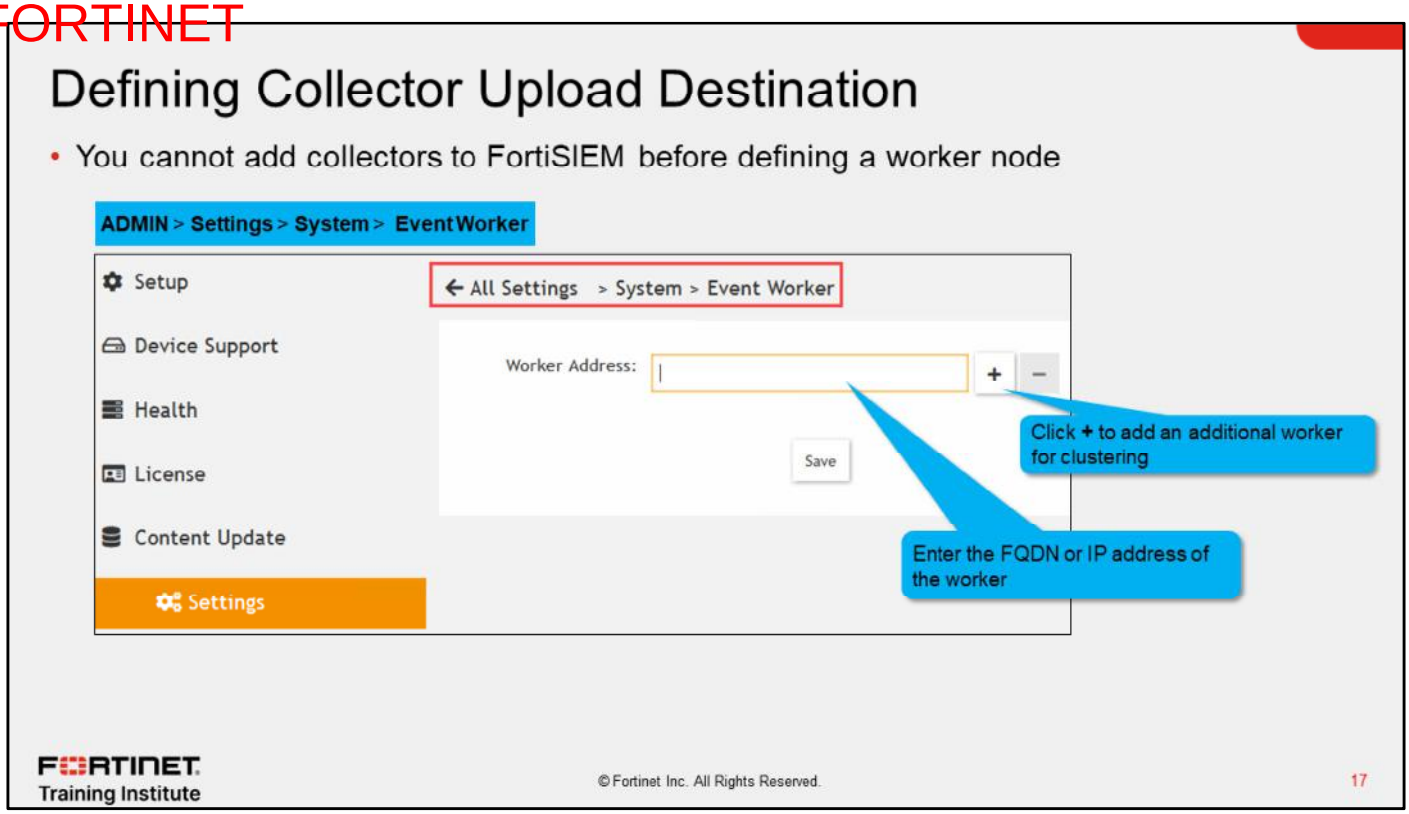
You cannot define collectors before you set the worker upload address. Collectors receive the worker upload address during registration and this value tells the collector to which node it should upload the data. If you don’t have a worker, then you must at least define the supervisor as a worker.

DO NOT REPRINT
© FORTINET



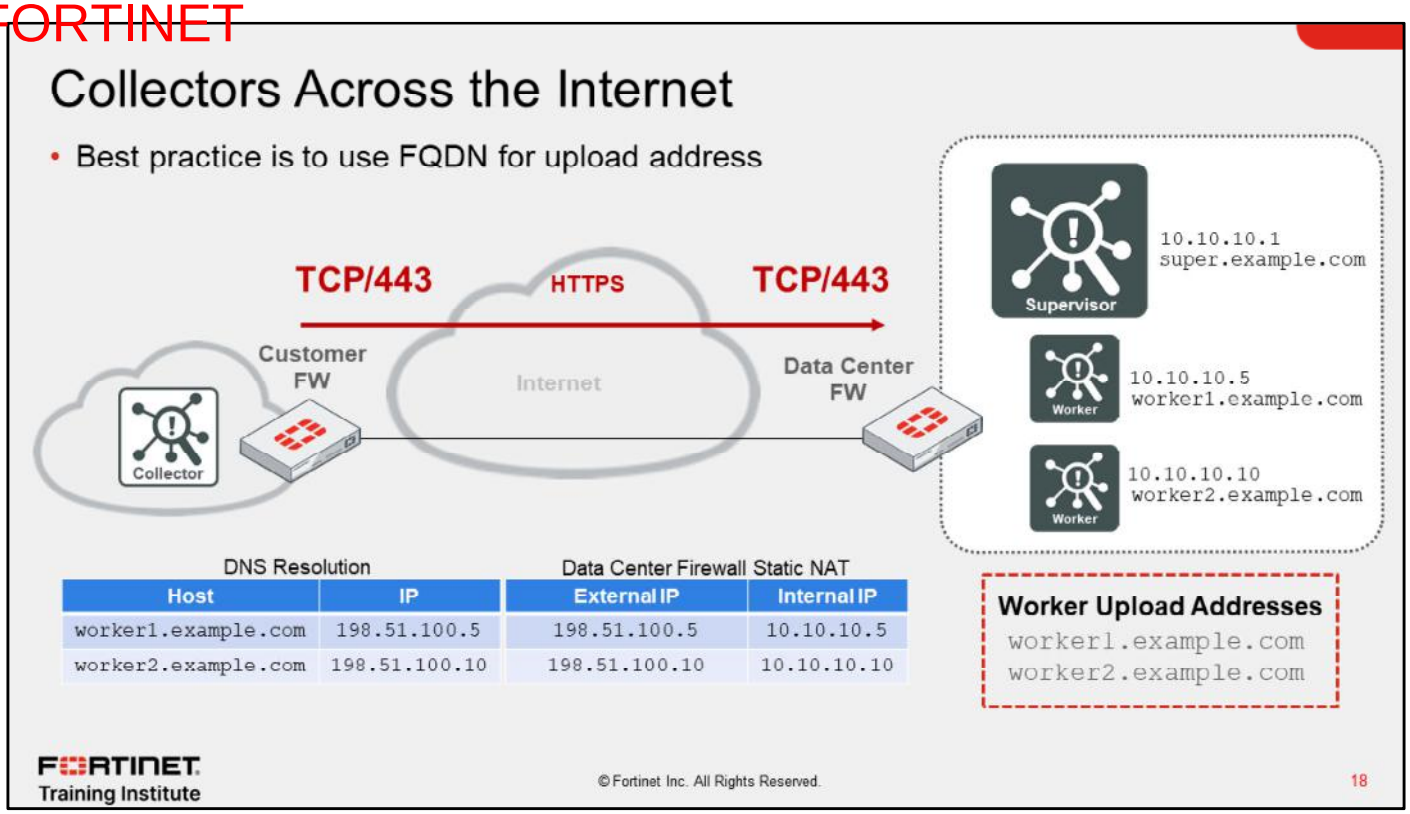
You can install workers as VA or hardware devices. After you deploy a VA worker for the first time, you must run the `configFSM.sh` script to configure the time zone and network configuration. To run the setup script, log in to the worker as root user and enter the commands shown on this slide. The setup script guides you through the network configuration where you configure the IP address, gateway, DNS, and hostname.

DO NOT REPRINT
© FORTINET



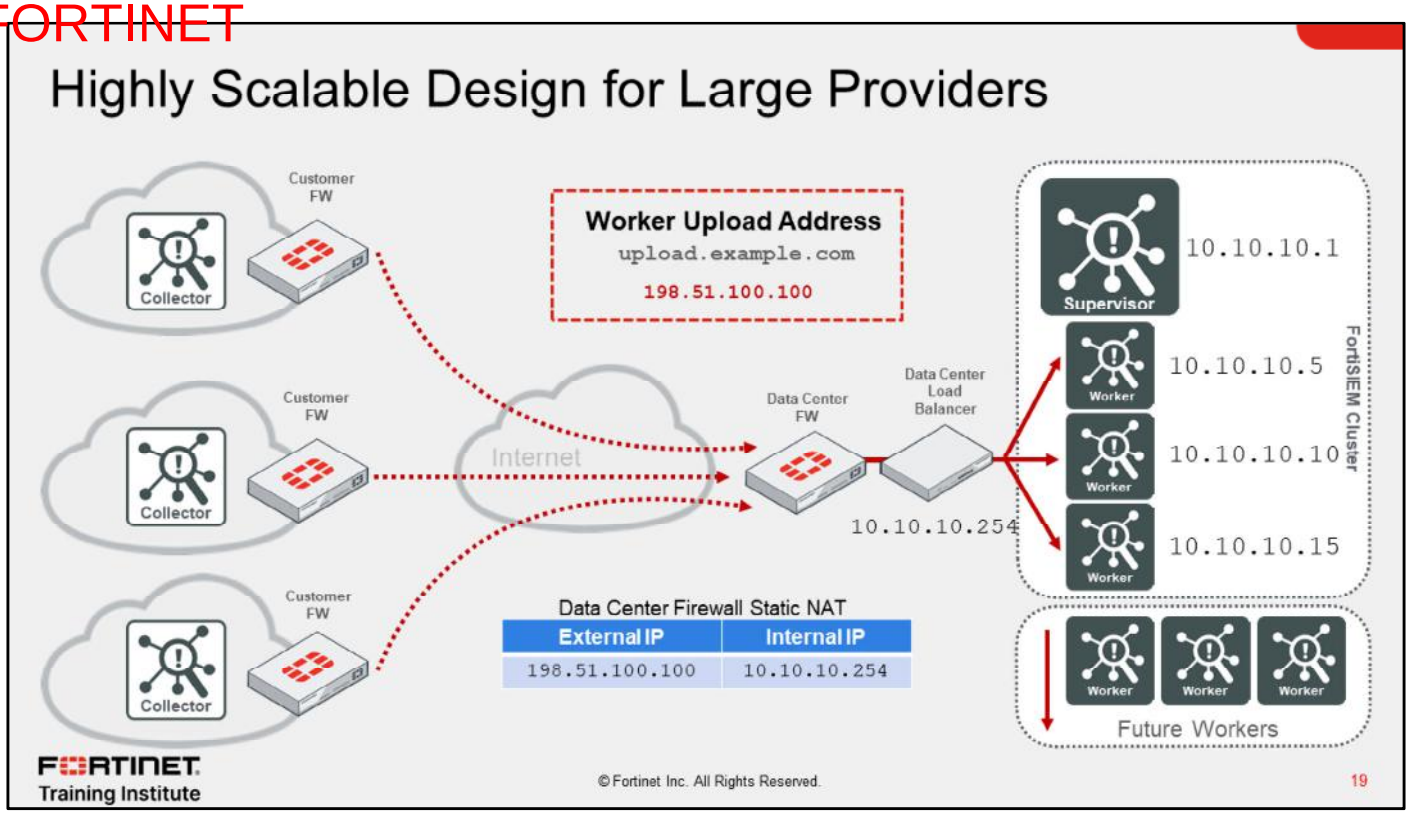
If you don't have a worker node, then you must define the IP address of the supervisor node as a worker address, so that you can add collectors to the cluster. For larger software-based deployments that involve multiple collectors, a large number of monitored devices, or devices with high event rates, it is highly recommended that you deploy one or more workers to distribute the workload of the supervisor node. Note that you can add workers to hardware-based appliances.

DO NOT REPRINT
© FORTINET



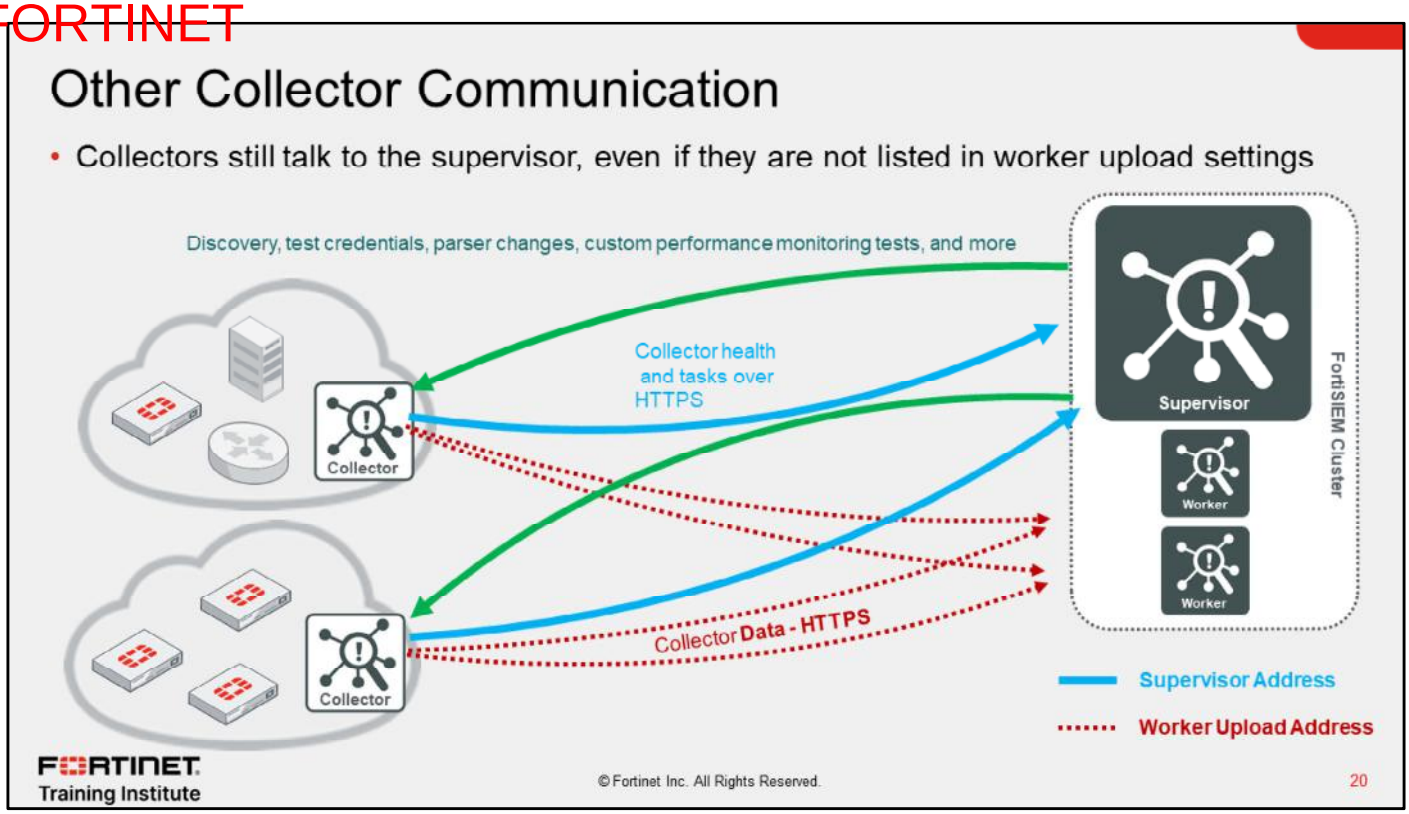
Each value in the worker upload address must listen on TCP port 443 for connections from collectors. The more workers you add to the cluster, the more addresses that must be published to the internet. On the data center firewall, create virtual IPs that map the external public IP to the worker private IP address. Collector communication is only one way, so the customer firewall policies must allow all outbound traffic from the collectors on port 443.

DO NOT REPRINT
© FORTINET



In a highly scalable design for large providers, you can place a load balancer behind the data center firewall. You can create a static NAT on the data center firewall to map the public IP address to the IP address of the load balancer. That public IP address is the worker upload address for all the customer collectors. The load balancer should evenly distribute the customer data to the workers in the FortiSIEM cluster.

DO NOT REPRINT
© FORTINET



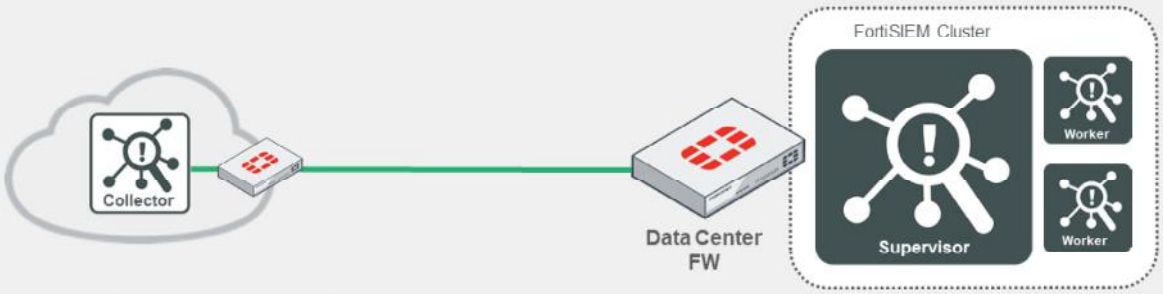
Even though collectors send data directly to the workers, they still communicate periodically with the supervisor node.

Collectors use the same communication channel that is used for sending the log data is also used to ask for any tasks, such as discovery, test credentials, parser changes, custom performance monitoring tests, and more, from the supervisor node. The supervisor node uses the same session to communicate with the collectors. You need to allow only an outbound connection in the customer firewall policy. The supervisor does not explicitly initiate any connections to the collector nodes.

DO NOT REPRINT
© FORTINET

Service Port Summary

- Summary of service provider firewall communications



From	To	Port	Description
Collector	Supervisor	TCP/443	Health/Tasks
Collector	All Nodes Defined Under Worker Upload Settings	TCP/443	Event Data
Super/Worker	Collector	None	

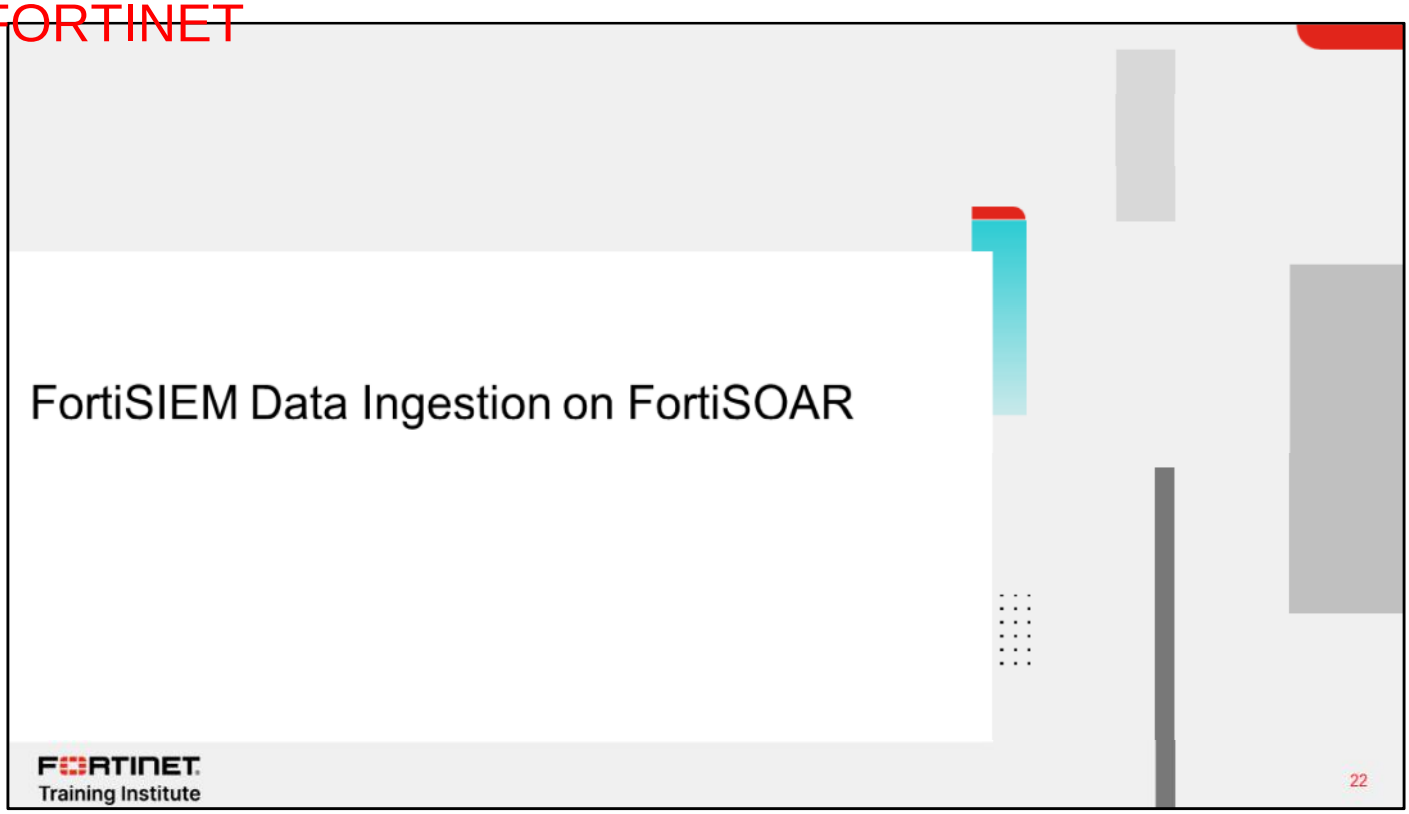
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

21

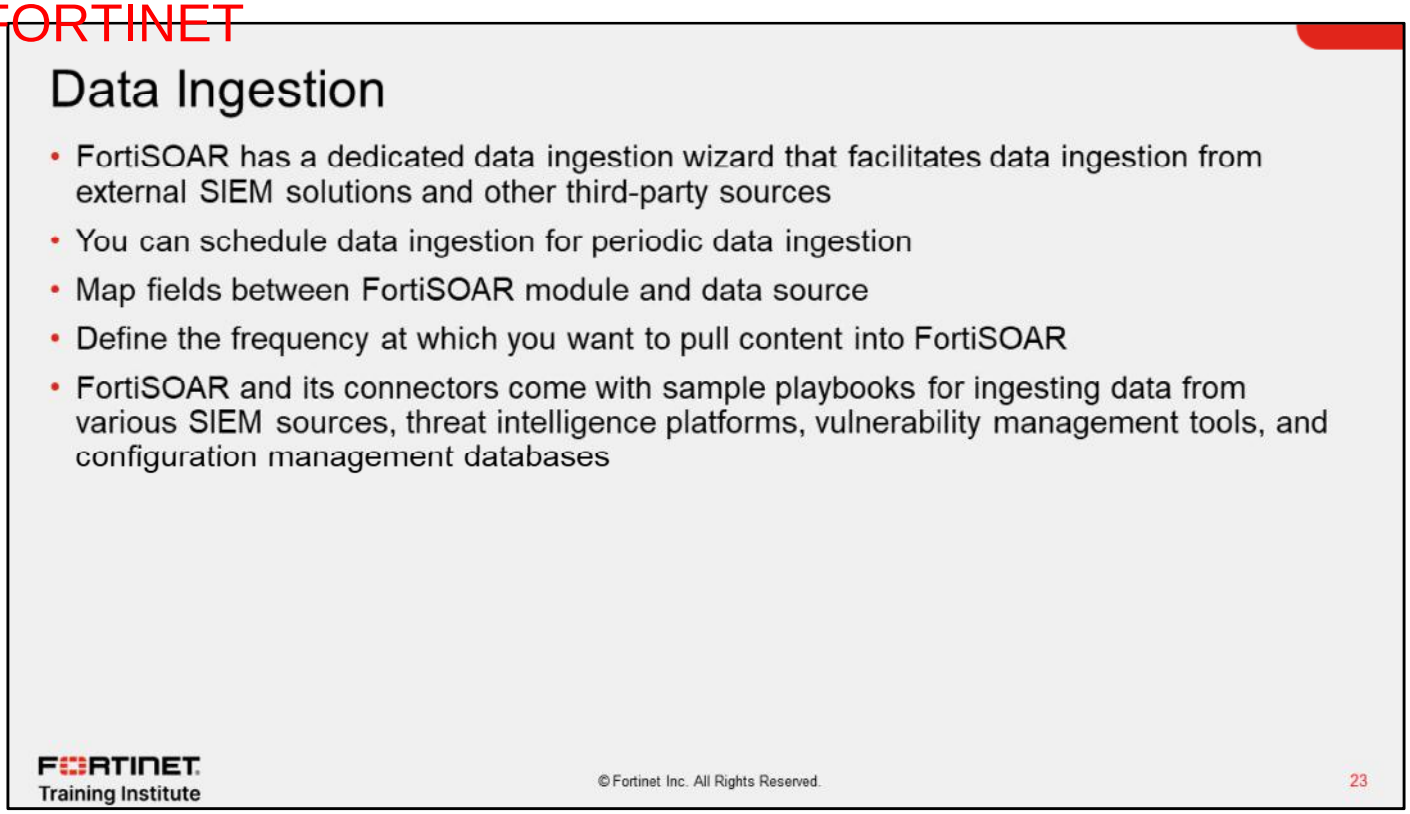
You must allow only TCP port 443 for the collector to communicate to the FortiSIEM cluster. Collector health information and tasks are sent only to the supervisor, and the event data is sent to all nodes that are defined in the worker upload settings. If the supervisor is also a worker, then the supervisor also receives event data.

DO NOT REPRINT
© FORTINET



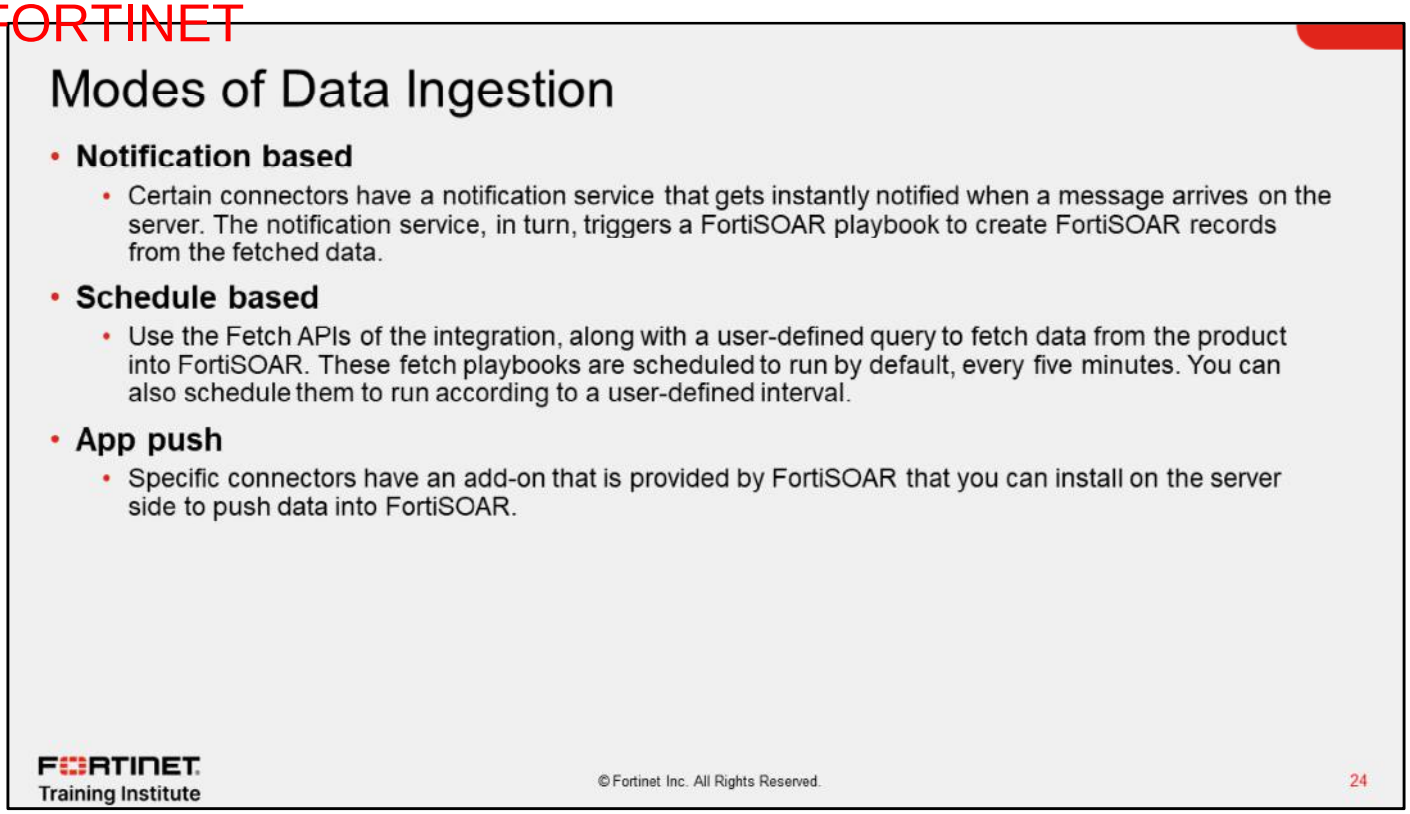
In this section, you will learn about the data ingestion mechanism on FortiSOAR, and various settings available for fine-tuning data ingestion and field mapping within a FortiSOAR module.

DO NOT REPRINT
© FORTINET



You can use the data ingestion wizard to simplify data ingestion configuration and mapping of fields between the two systems. The wizard fetches sample data from the source, mapping fields from the source data into a FortiSOAR module, and setting up an ingestion schedule if the connector is enabled for scheduling. Based on the inputs you provide in the wizard, the system dynamically generates the ingestion workflows without the user being exposed to the details of playbooks, and easing the process of configuring ingestion.

DO NOT REPRINT
© FORTINET



Data ingestion can run in one of three modes:
- Notification based
- Schedule based
- App push

When running in notification-based mode, connectors such as IMAP, Exchange, Syslog, and so on, have a notification service that is instantly notified when a message arrives on the server. In turn, the notification service triggers a FortiSOAR playbook to create FortiSOAR records from the fetched data.

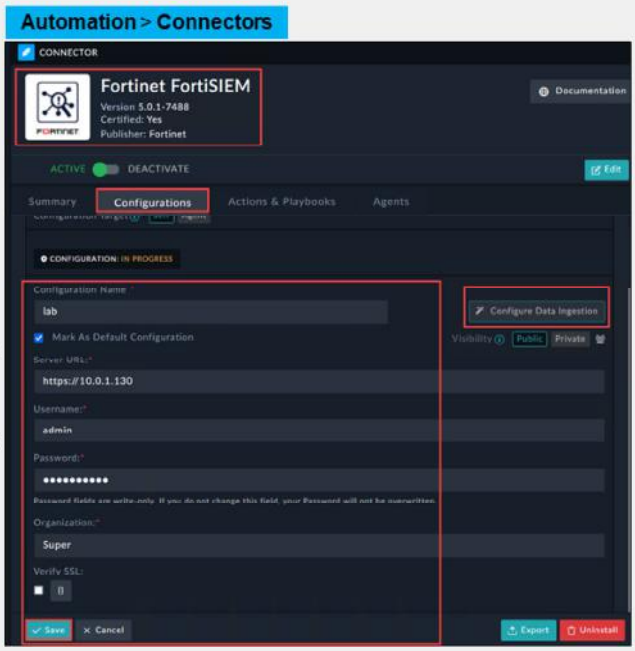
When running in schedule-based mode, connectors such as QRadar, Anomali, ServiceNow, and so on, use the fetch APIs of the integration. By default, these fetch playbooks are scheduled to run every five minutes, which the user can modify to run according to a user-defined interval. Note that most integrations that support a notification-based ingestion also support a schedule-based ingestion. However, you must configure only one of the two, otherwise both pull data from the same source, which might result in data loss due to conflicts.

When running in app push mode, connectors, such as Splunk, have a FortiSOAR add-on that you can install on the server side to push data from the server to FortiSOAR. You configure the add-on with a user password or appliance-based authentication in FortiSOAR and it triggers FortiSOAR playbooks to create the records in FortiSOAR.

DO NOT REPRINT
© FORTINET

Data Ingestion Tab

- A connector that fetches data from external sources has a tab called **Configure Data Ingestion**
- Sample ingestion playbooks were created keeping the following points in mind:
  - Each playbook that contributes to data ingestion contains the following tags: <connector_name>, <dataingestion>
  - **Fetch Playbook**: This playbook fetches sample data and is also used for the actual fetch operation once ingestion is configured
  - Data from external systems is presented to the users as source data on the **Field Mapping** screen in the data ingestion wizard



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

25

Any connector that supports data ingestion has a **Configure Data Ingestion** tab. The **Connectors** page has a section for data ingestion connectors that monitors the data ingestion and provides information on which connectors are configured for using the Data Ingestion Wizard, as well as other information, such as the status of a configuration, the schedule for ingestion, when data was last pulled using that configuration, and so on.

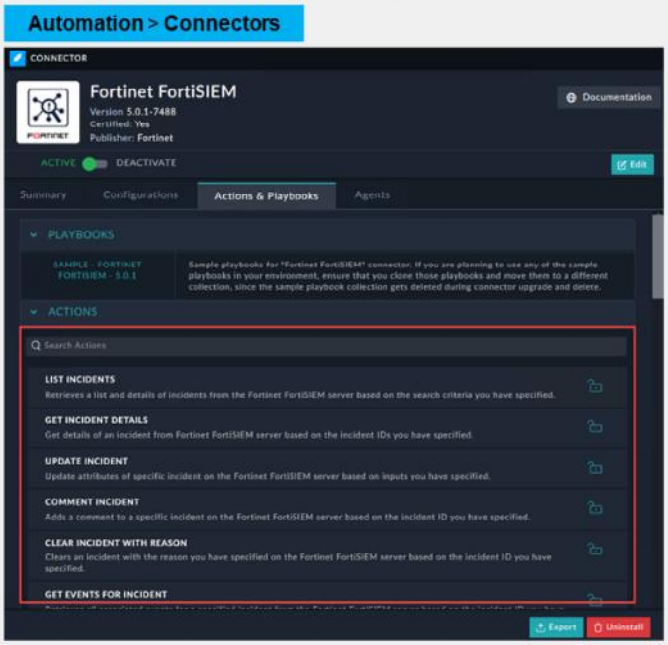
Sample ingestion playbooks were created keeping the following points in mind:

- Each playbook that contributes to data ingestion contains tags such as <connector_name> and <dataingestion>.
- **Fetch Playbook**: This playbook fetches sample data and is also used for the actual fetch operation once ingestion is configured.
- Data from external systems is presented to the users as source data on the **Field Mapping** screen in the data ingestion wizard.

DO NOT REPRINT
© FORTINET

Sample Playbook to Fetch Data—FortiSIEM

- Sample playbooks are used by the connector to fetch data but can also be used in a playbook environment



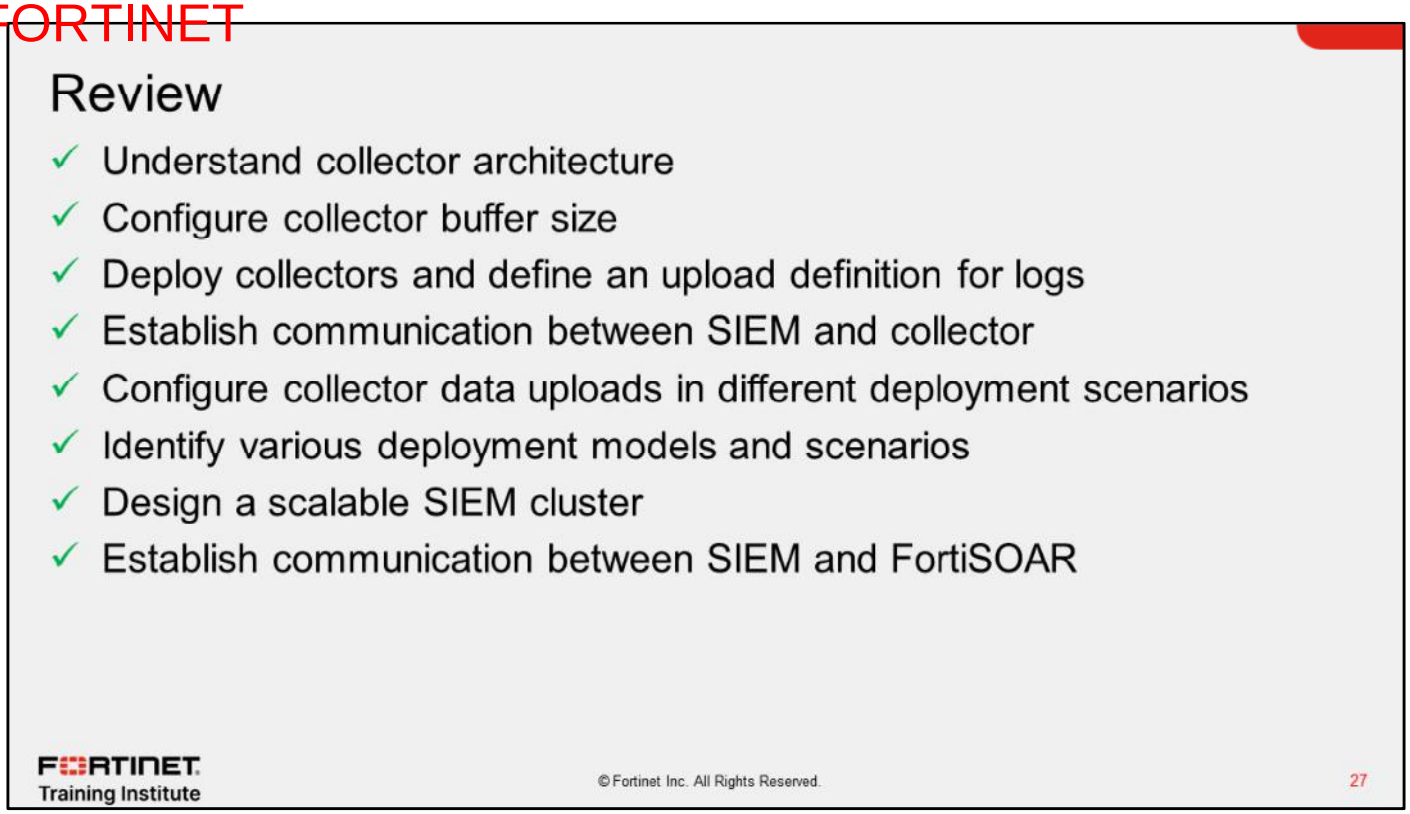
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

26

There are data ingestion connectors, such as **Fortinet FortiSIEM**, that include sample playbooks that can fetch data from a SIEM or other data aggregation systems. You can use these sample playbooks in other playbook environments. If you decide to use the sample playbooks in your environment, ensure that you clone those playbooks and move them to a different collection, since the sample playbook collection gets deleted during connector upgrade and delete.

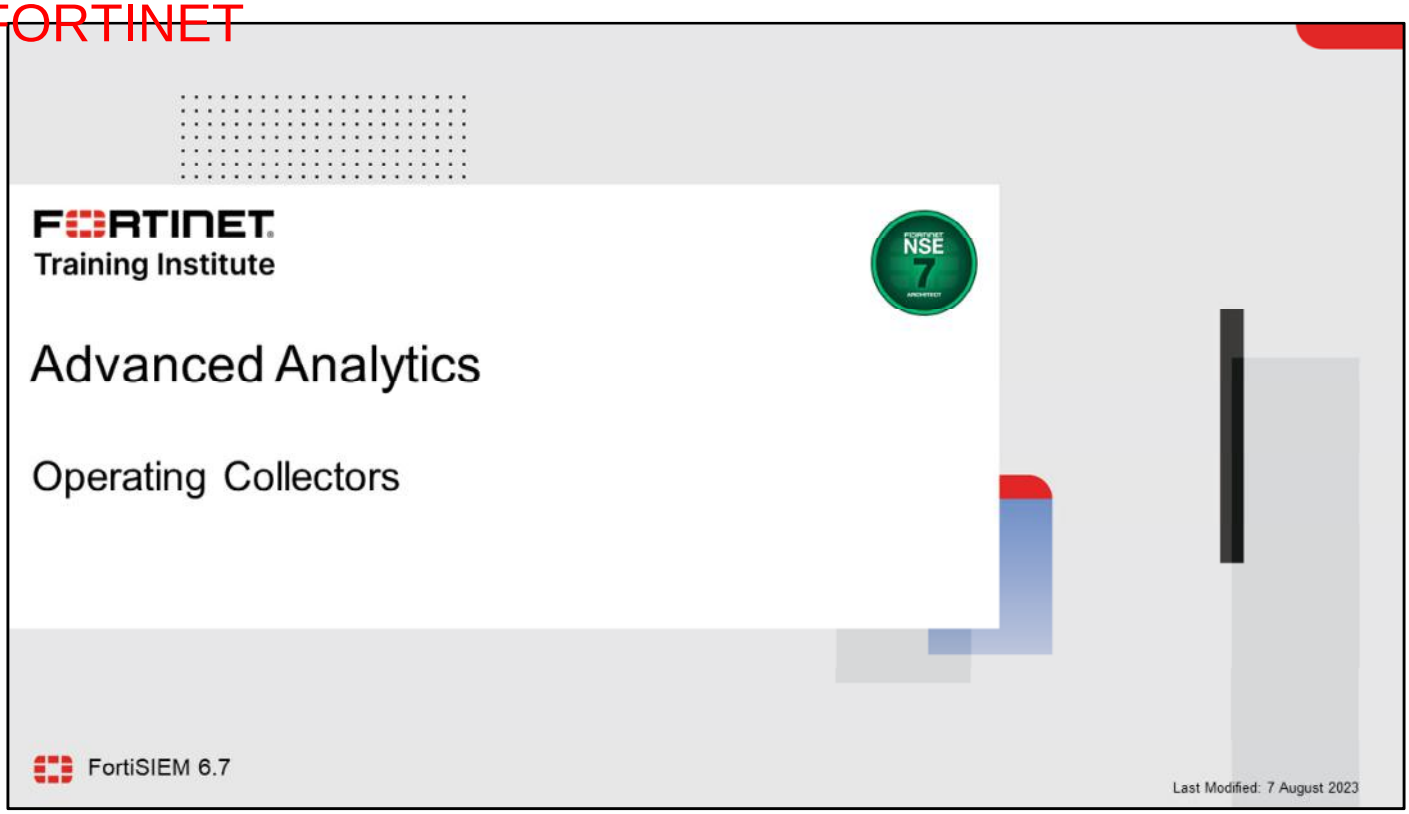
DO NOT REPRINT
© FORTINET



This slide shows the objectives that you covered in this lesson.

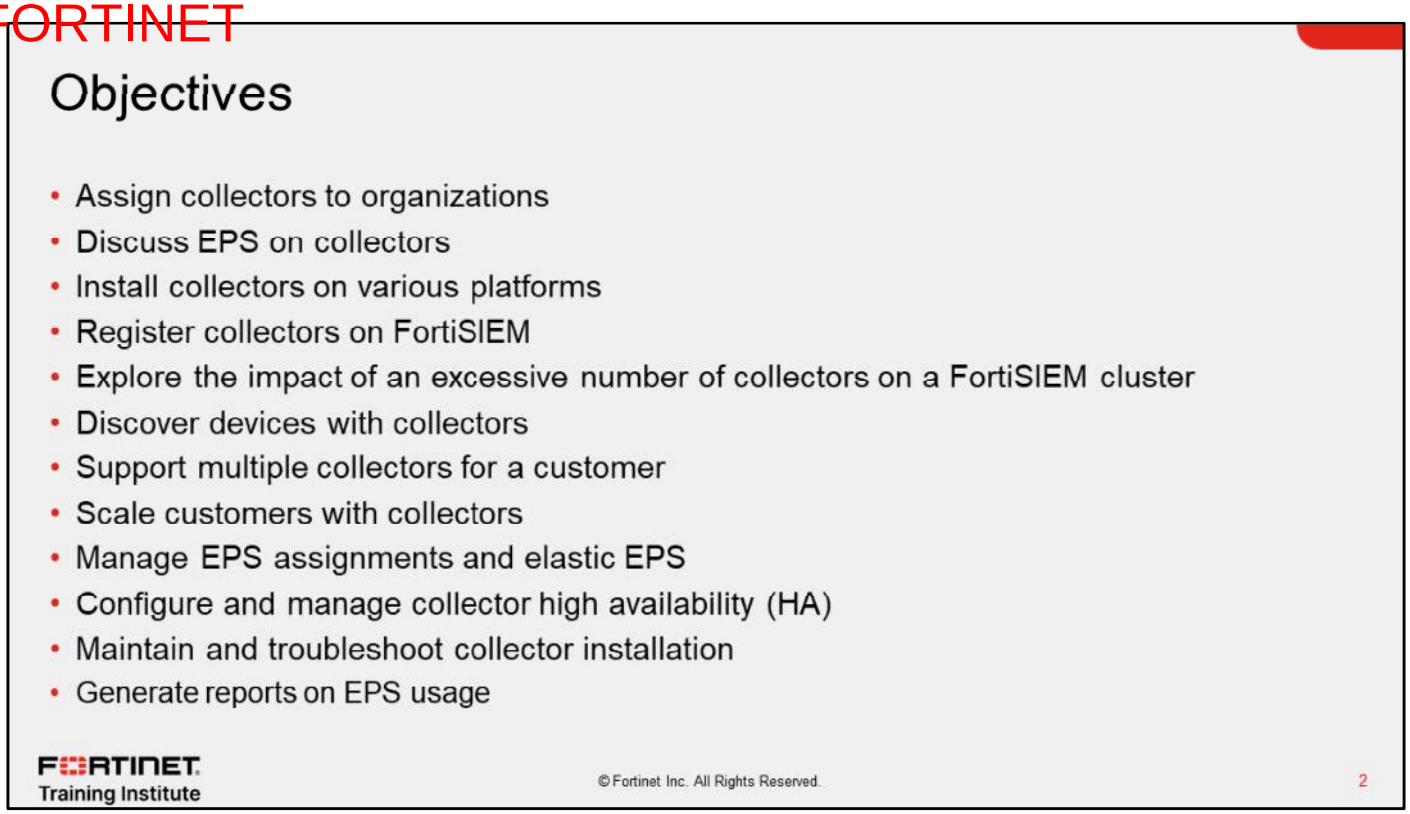
By mastering the objectives covered in this lesson, you learned about defining and deploying collectors. You also learned about collector operation, and scaling collectors as organizations grow.

DO NOT REPRINT
© FORTINET



In this lesson, you will learn about installing collectors on various platforms, registering collectors on FortiSIEM, and discovering devices with collectors. You will also learn about managing events per second (EPS) assignment for collectors.

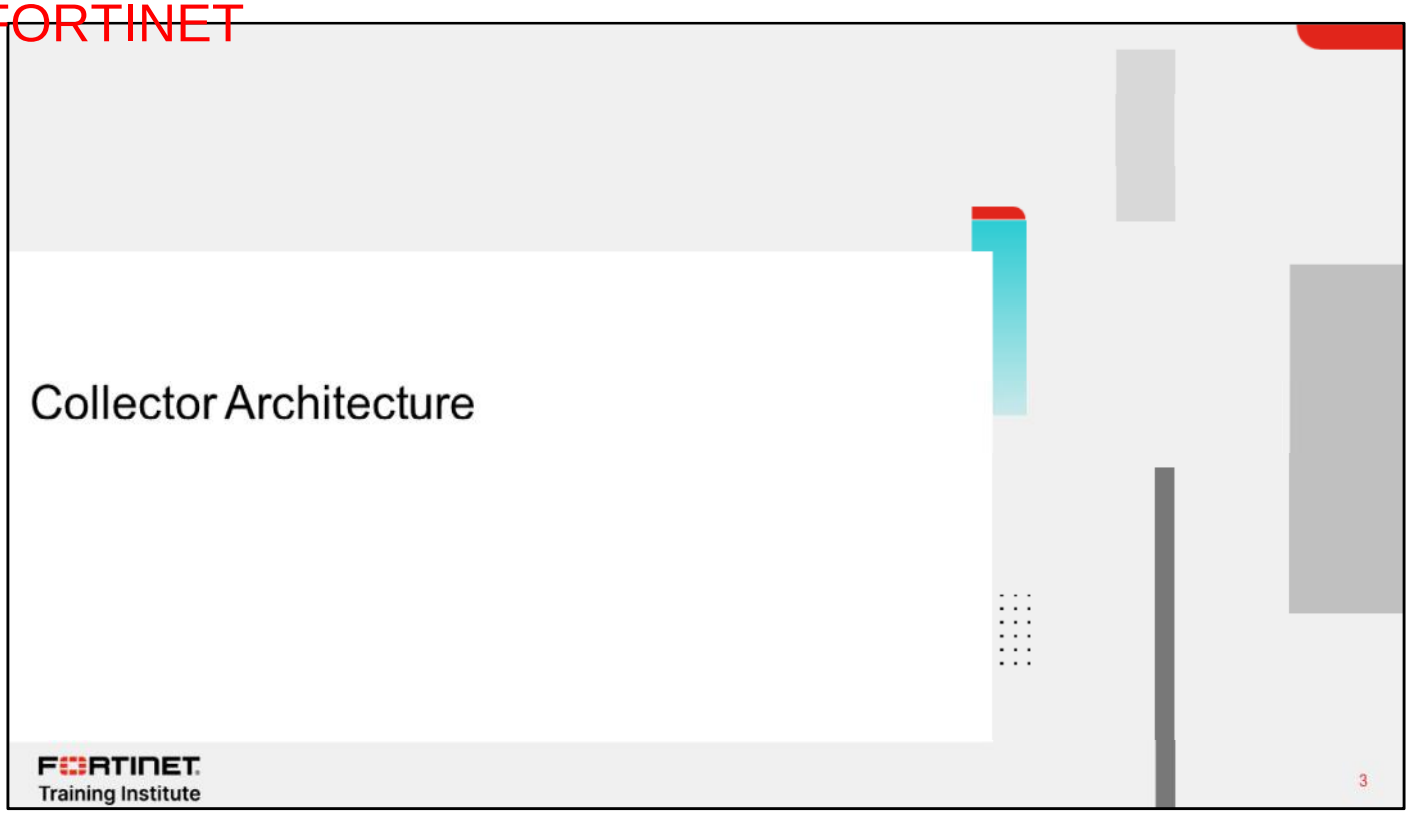
DO NOT REPRINT
© FORTINET



After completing this lesson, you should be able to achieve the objectives shown on this slide.

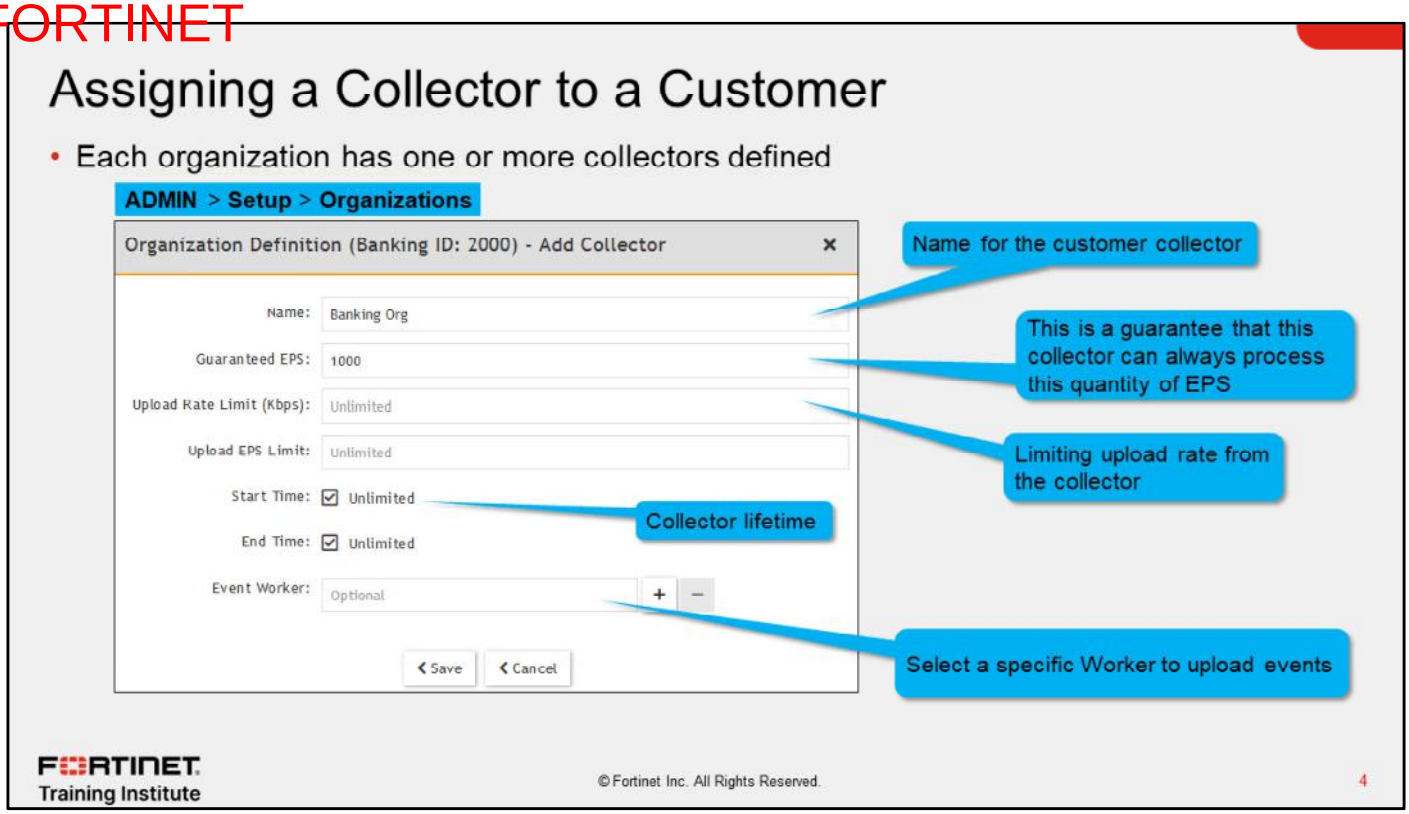
By demonstrating competence in EPS and collector installation, you should be able to identify EPS requirements for collectors, the impact of an excessive number of collectors on a FortiSIEM cluster, and troubleshoot collector installation.

DO NOT REPRINT
© FORTINET



In this section, you will learn about collector architecture, the minimum guaranteed EPS for collectors, and the overall super EPS.

DO NOT REPRINT
© FORTINET

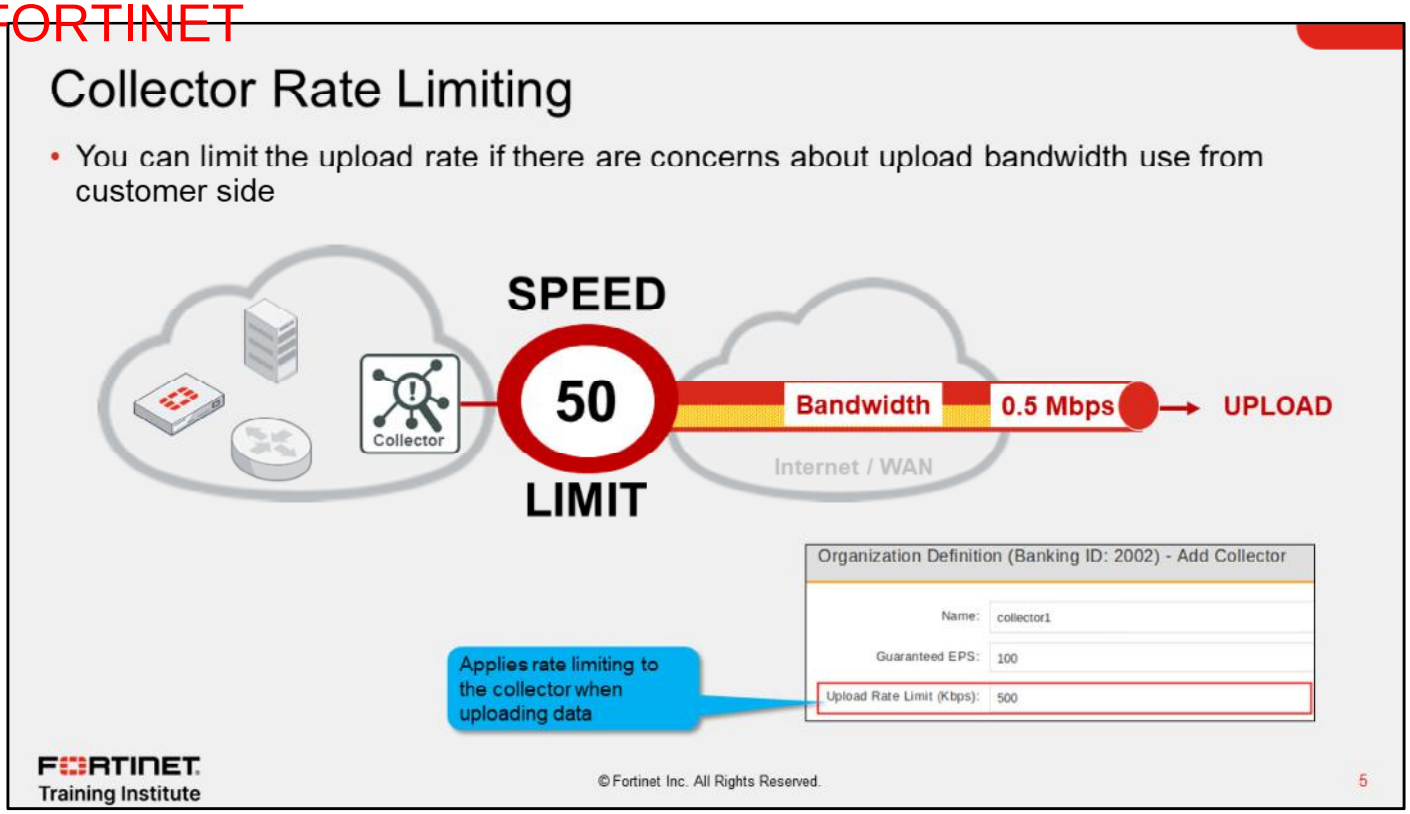


One of the ways of defining an organization is by associating it with a collector. The organization's collector monitors all devices belonging to that organization, as well as all events.

The **Guaranteed EPS** setting sets the threshold at which events are always accepted, as long as the event rate is below the configured value. FortiSIEM reallocates excess EPS, which is the licensed rate minus the sum of guaranteed EPS across all collectors. It is based on need, but the allocation never goes below the configured guaranteed EPS value.

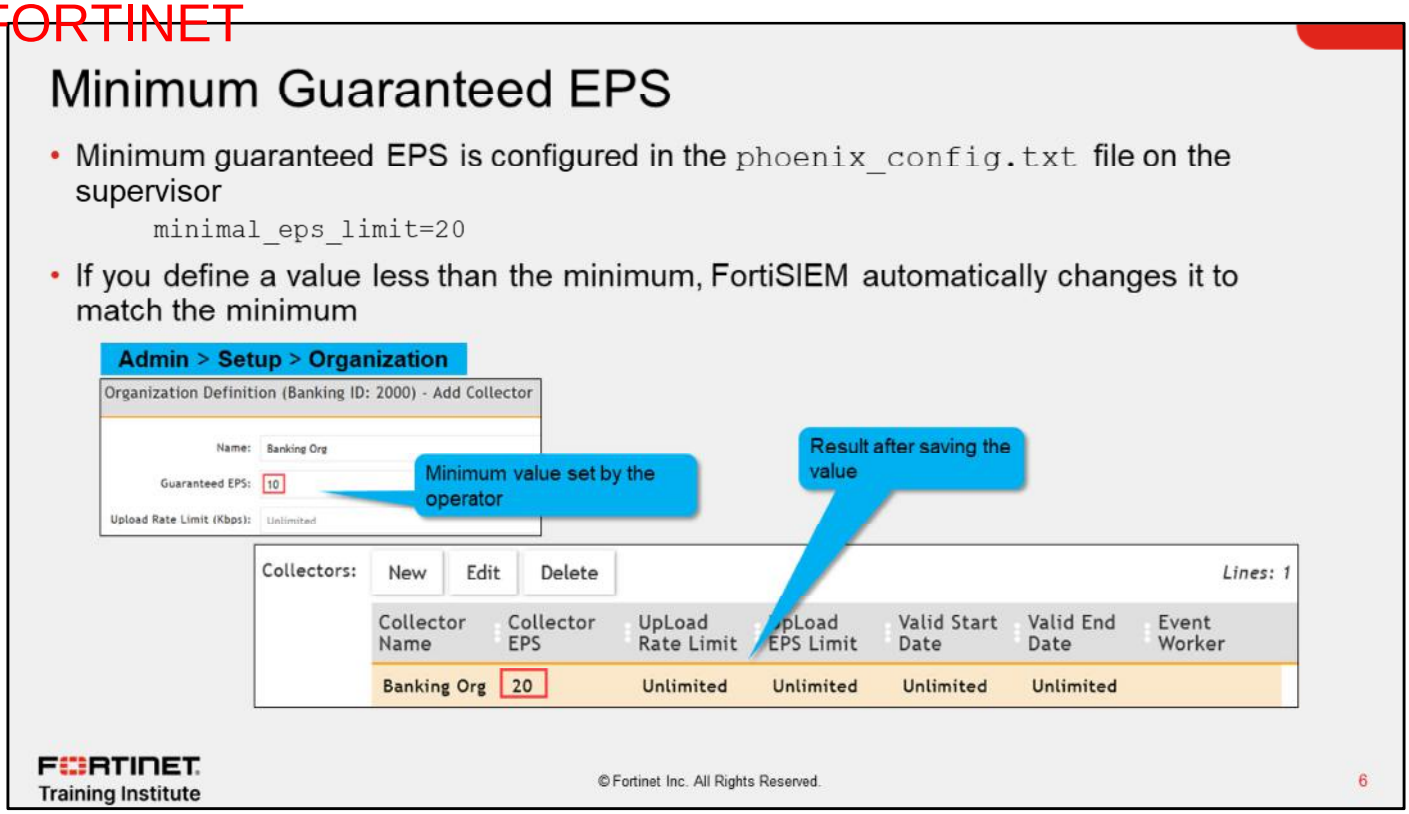
The start and end time settings specify the dates during which the collector is active.

DO NOT REPRINT
© FORTINET



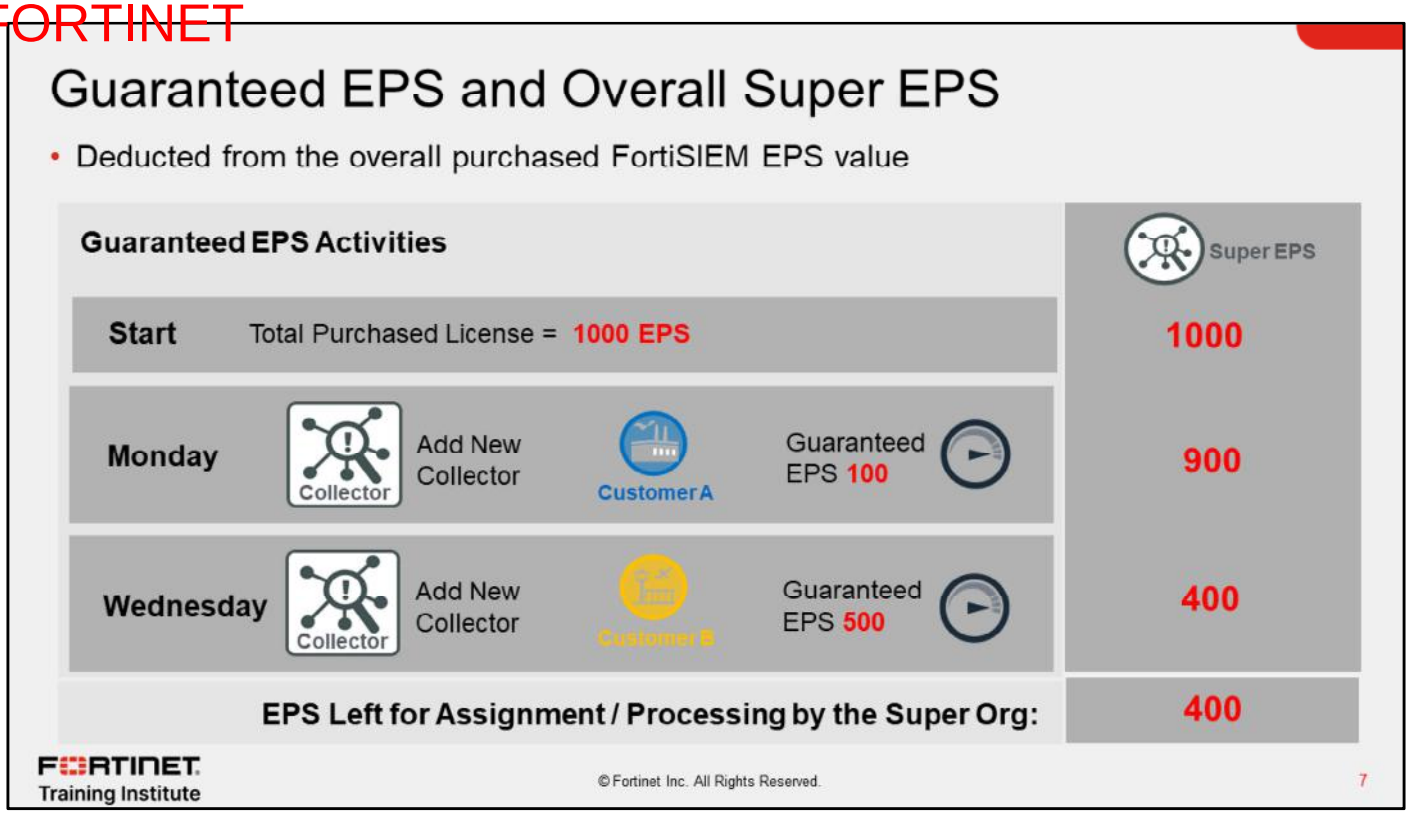
Sometimes concerns may be made over the amount of bandwidth used or available for collector uploads. You can configure the maximum upload bandwidth value to address those concerns. The default value is unlimited.

DO NOT REPRINT
© FORTINET



FortiSIEM comes preconfigured with a minimum guaranteed EPS. It is set to 20 in the `phoenix_config.txt` file on the supervisor. So, if you configure an EPS lower than the minimum value, FortiSIEM automatically changes the value to the match the minimum value when you try to save the changes.

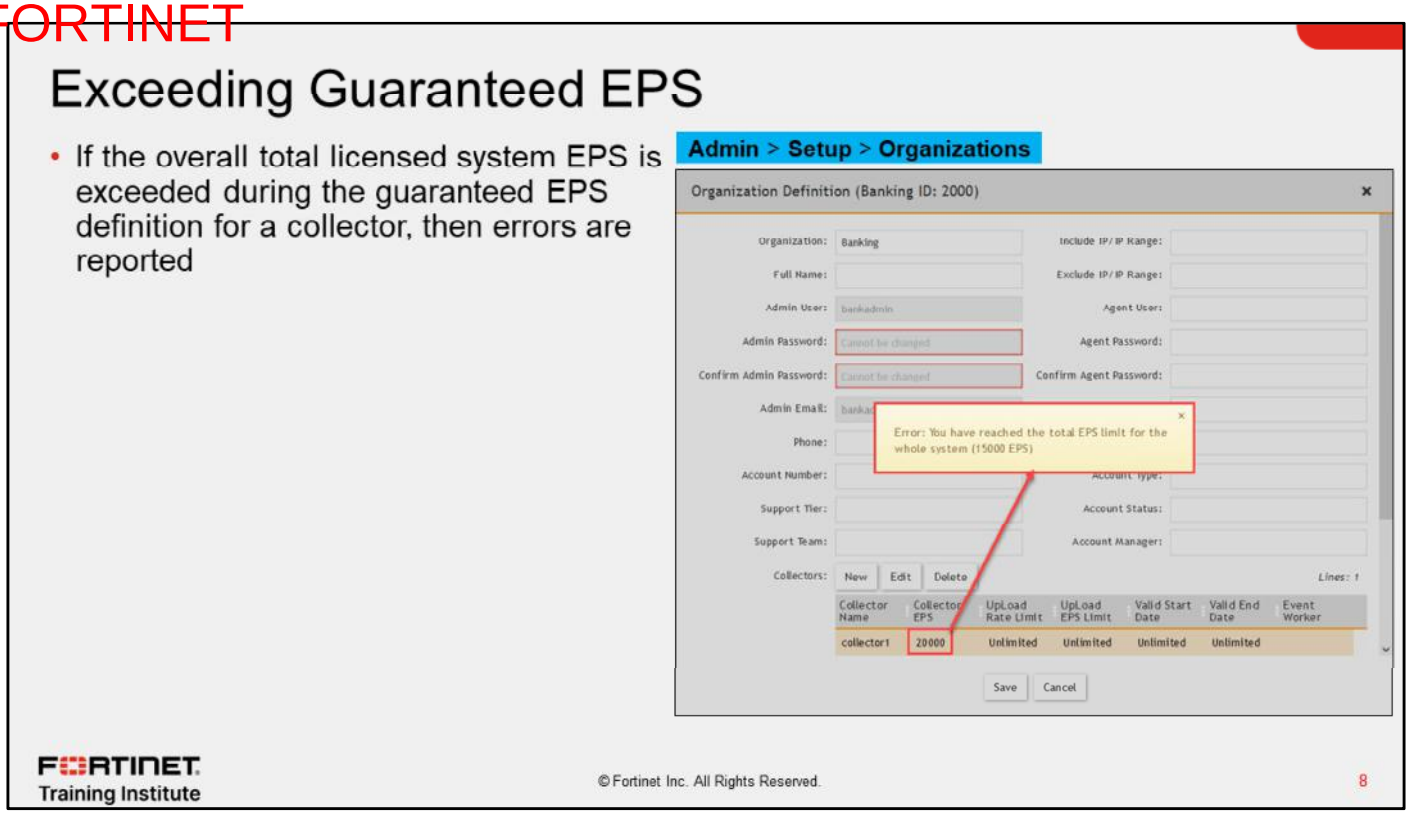
DO NOT REPRINT
© FORTINET



The purchased EPS licenses are treated like a pool, which you have to distribute across all collectors. The guaranteed EPS value for each collector is deducted from the overall purchased FortiSIEM EPS value.

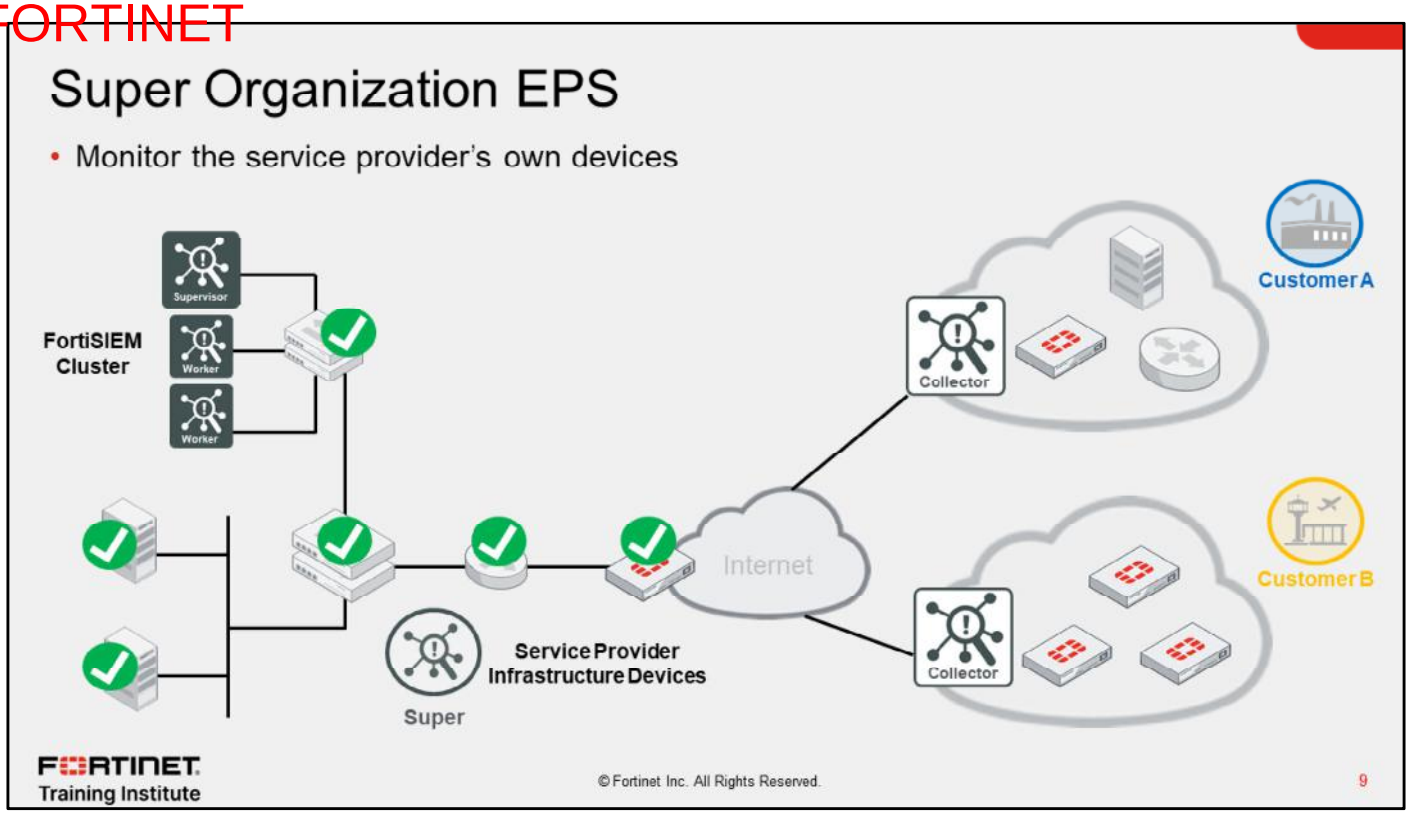
For example, if you've purchased a 1000-EPS license and configured a new collector with 100 EPS for customer A, 900 EPS remain in the pool. If you add another collector with 500 EPS for customer B, 400 EPS remain in the pool. You can use the remaining EPS pool for future customer collectors, or use it for your own organization to monitor and manage local assets.

DO NOT REPRINT
© FORTINET



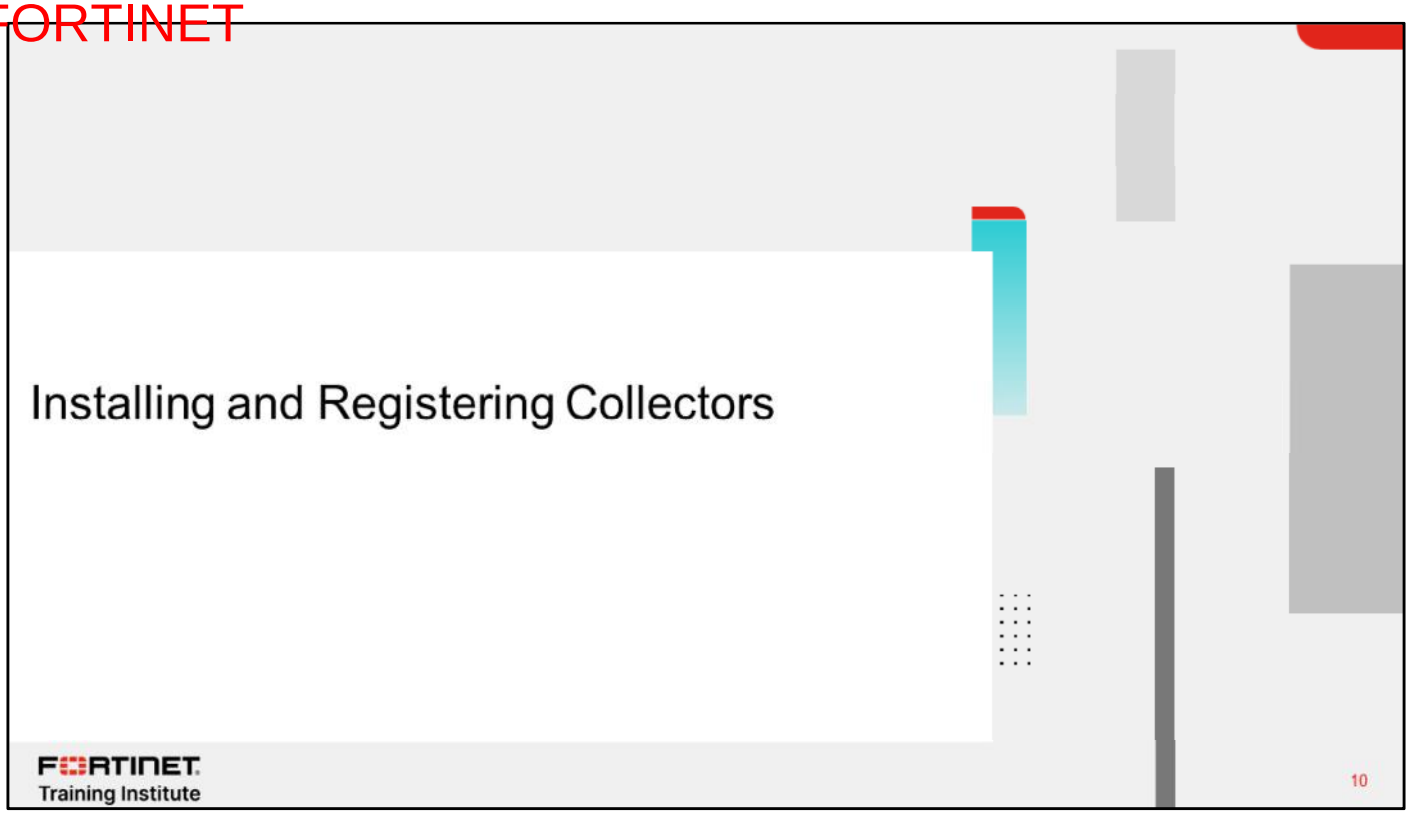
If the overall total licensed system EPS is exceeded during the guaranteed EPS definition for a collector, then errors are reported. For example, if you purchased a 15,000-EPS license and you try to allocate more than that value to a single collector, then you will get an error stating that you have reached the total EPS limit for the whole system.

DO NOT REPRINT
© FORTINET



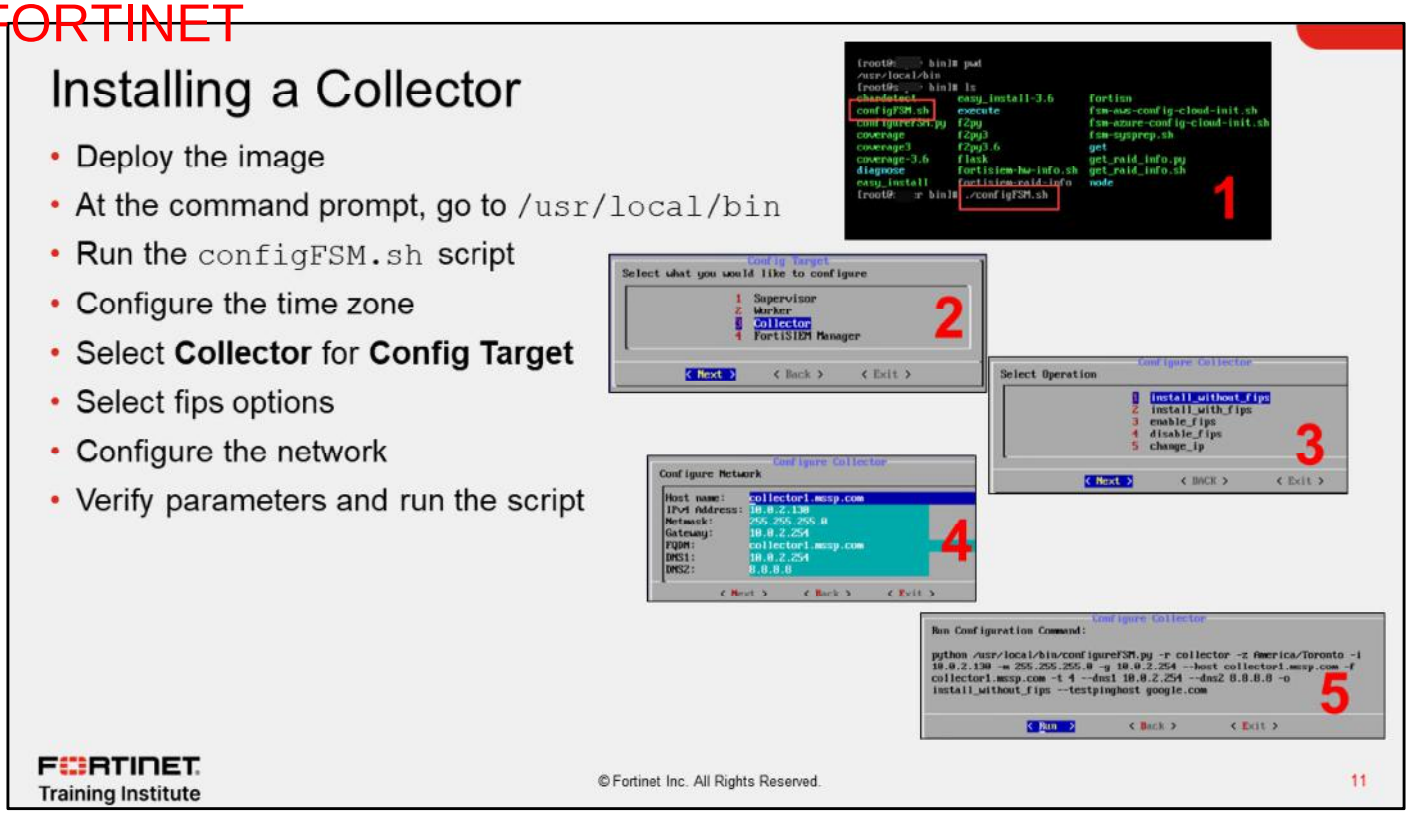
A super organization needs its own EPS to manage its own assets. A service provider may have critical assets within its own organization that must be monitored. For that reason, it's always a good idea to reserve some EPS for the service provider's own needs.

DO NOT REPRINT
© FORTINET



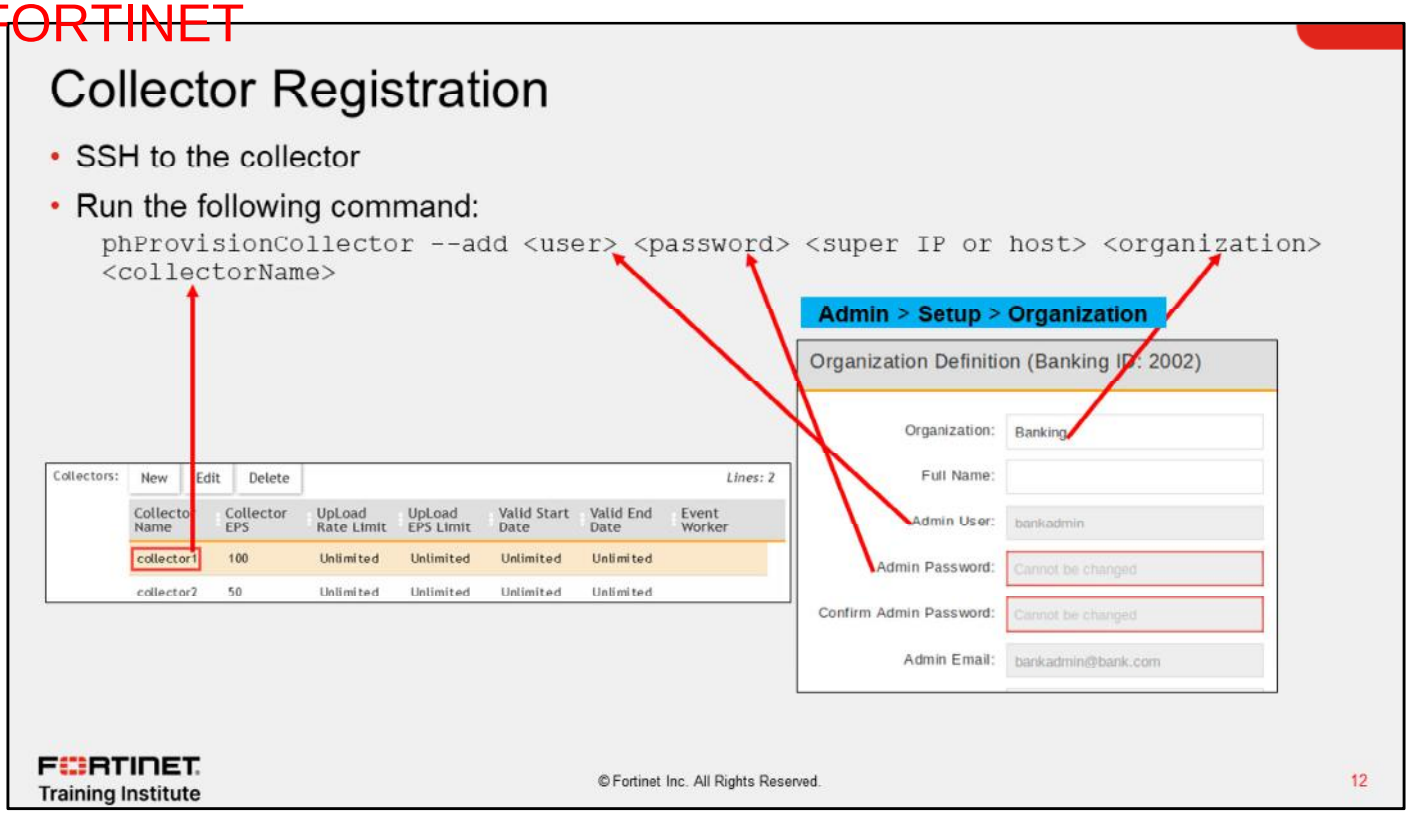
In this section, you will learn about installing a collector and then registering it on the supervisor.

DO NOT REPRINT
© FORTINET



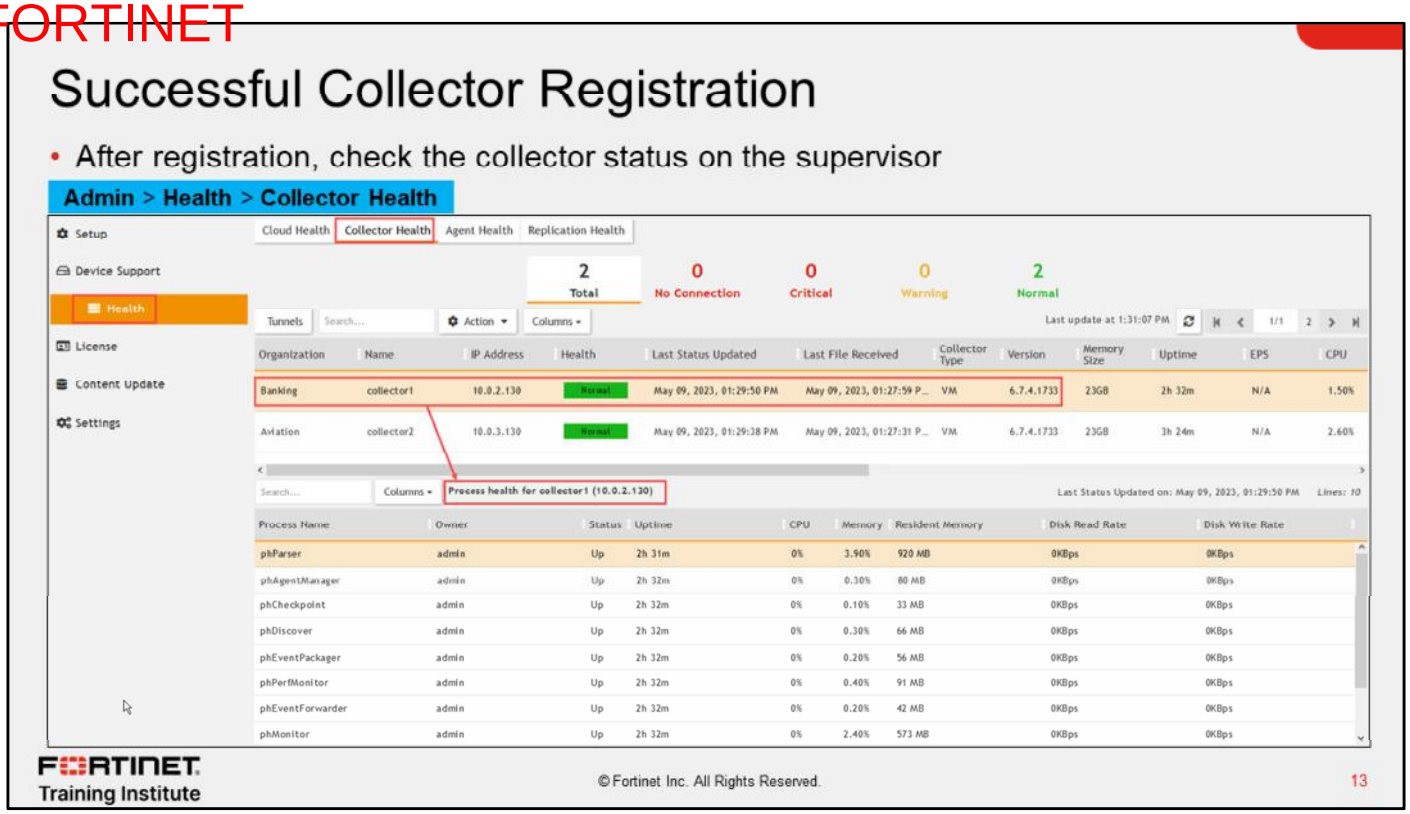
You can install collectors as VA or hardware appliances. After you deploy a VA collector for the first time, you must run the `configFSM.sh` script to configure the time zone and the network configuration. To run the setup script, log in to the collector as root user, and then enter the commands shown on this slide. The setup script guides you through the network configuration where you configure the IP address, gateway, DNS, and hostname.

DO NOT REPRINT
© FORTINET



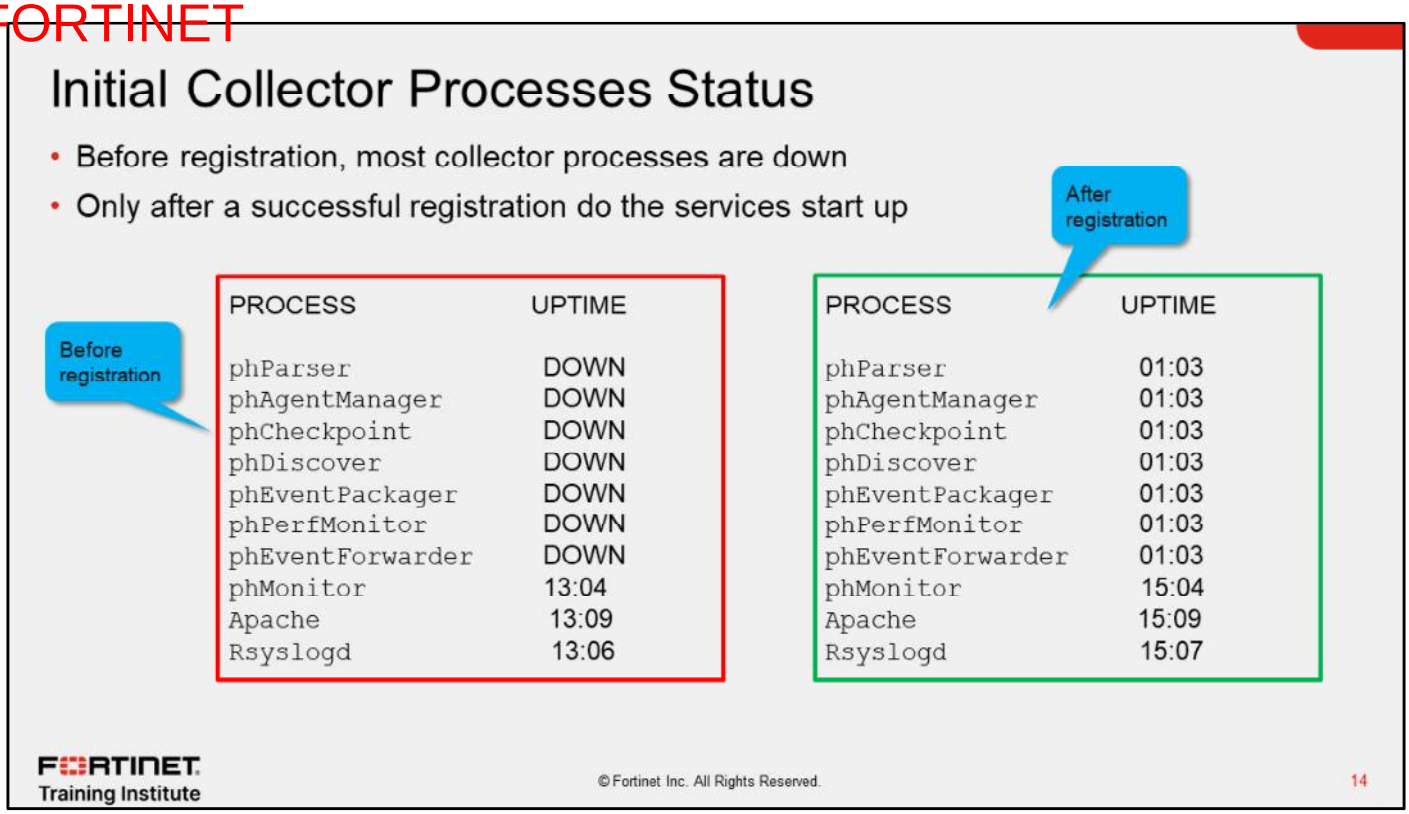
After you complete the collector installation, you must register the collector on the supervisor. Before registration, you must collect information from the organization definition setting of the supervisor, such as organization name, administrator user, and administrator password. After you have that information, use the commands shown on this slide to register the collector on the supervisor.

DO NOT REPRINT
© FORTINET



After you register the collector, check the status on the supervisor GUI. The **Process health** bottom pane displays the processes running on the collector and their status, up time, CPU, memory, and more.

DO NOT REPRINT
© FORTINET



Before a collector is registered, most processes are down. The services start up only after the collector is registered successfully on the supervisor.

DO NOT REPRINT
© FORTINET

Collector License File

- After registration, the collector pulls a license file from the supervisor
- Found under `/etc/opsd/`

```
[root@OrgA_Collector ~]# ls -la /etc/opsd/

drwxrwxrwx    2 admin admin  4096 Mar 28 12:27 .
drwxr-xr-x. 100 root  root  12288 Mar 28 14:06 ..
-rw-rw-r--    1 admin admin  1121 Mar 28  2019 .fortisiem4x0
```

- Removing this file takes the collector processes *down*, but not instantly

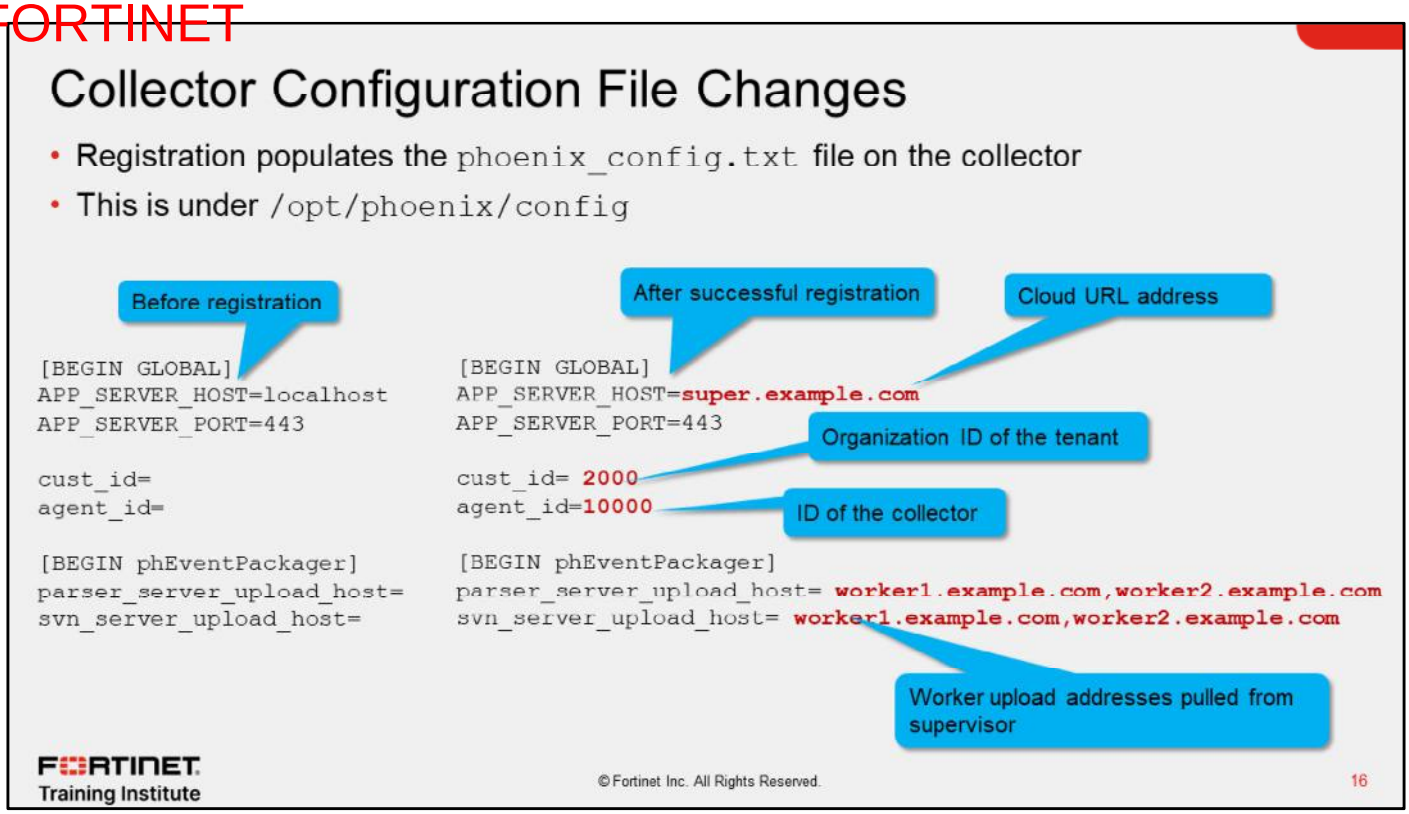
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

15

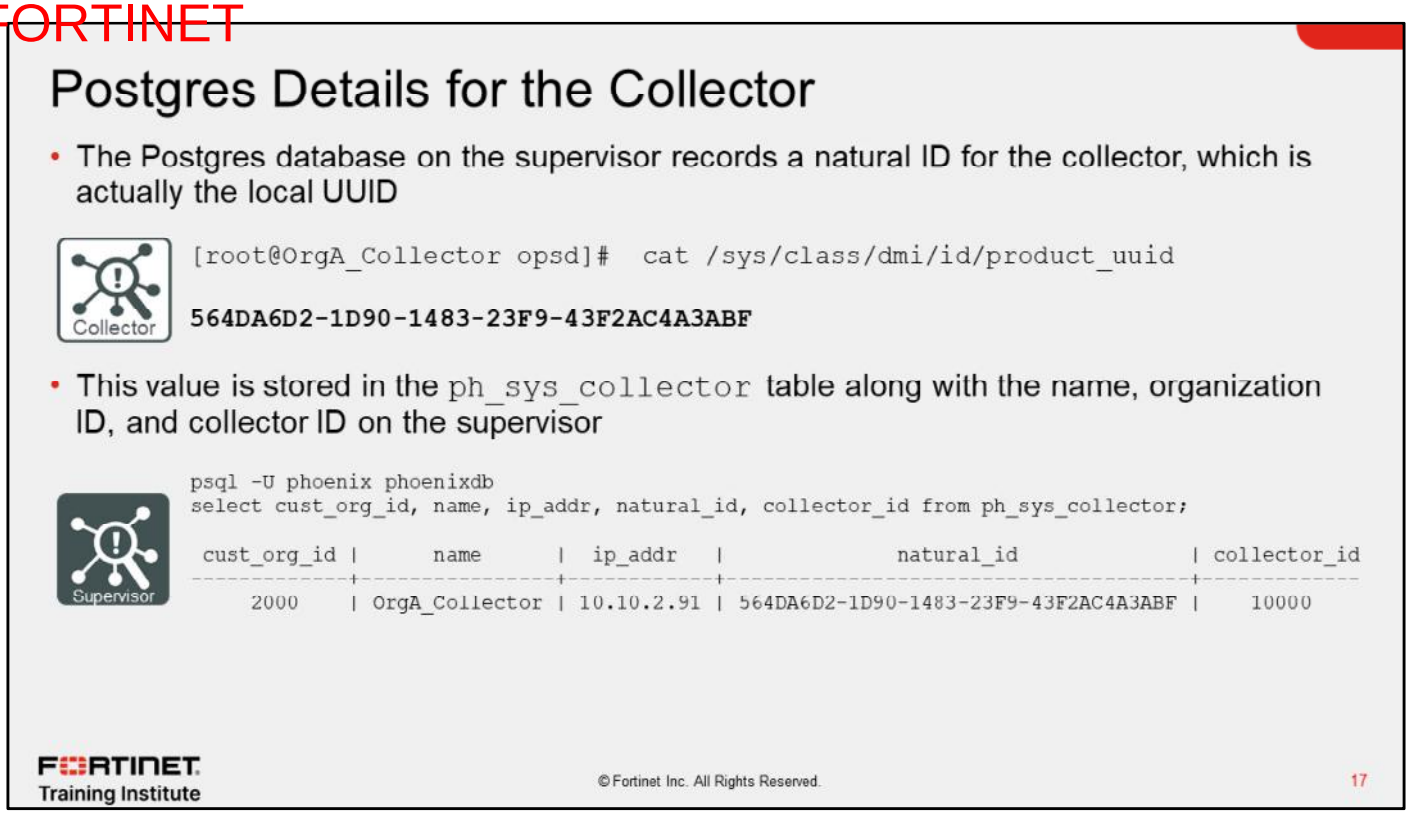
Once a collector is registered to a supervisor, it pulls a license file from the supervisor. The file is saved to the directory location shown on this slide. You should not remove this file because doing so can take down the collector processes.

DO NOT REPRINT
© FORTINET



Registration populates the collector configuration file with information, such as the cloud URL address of the supervisor, customer ID, agent ID, and worker upload address.

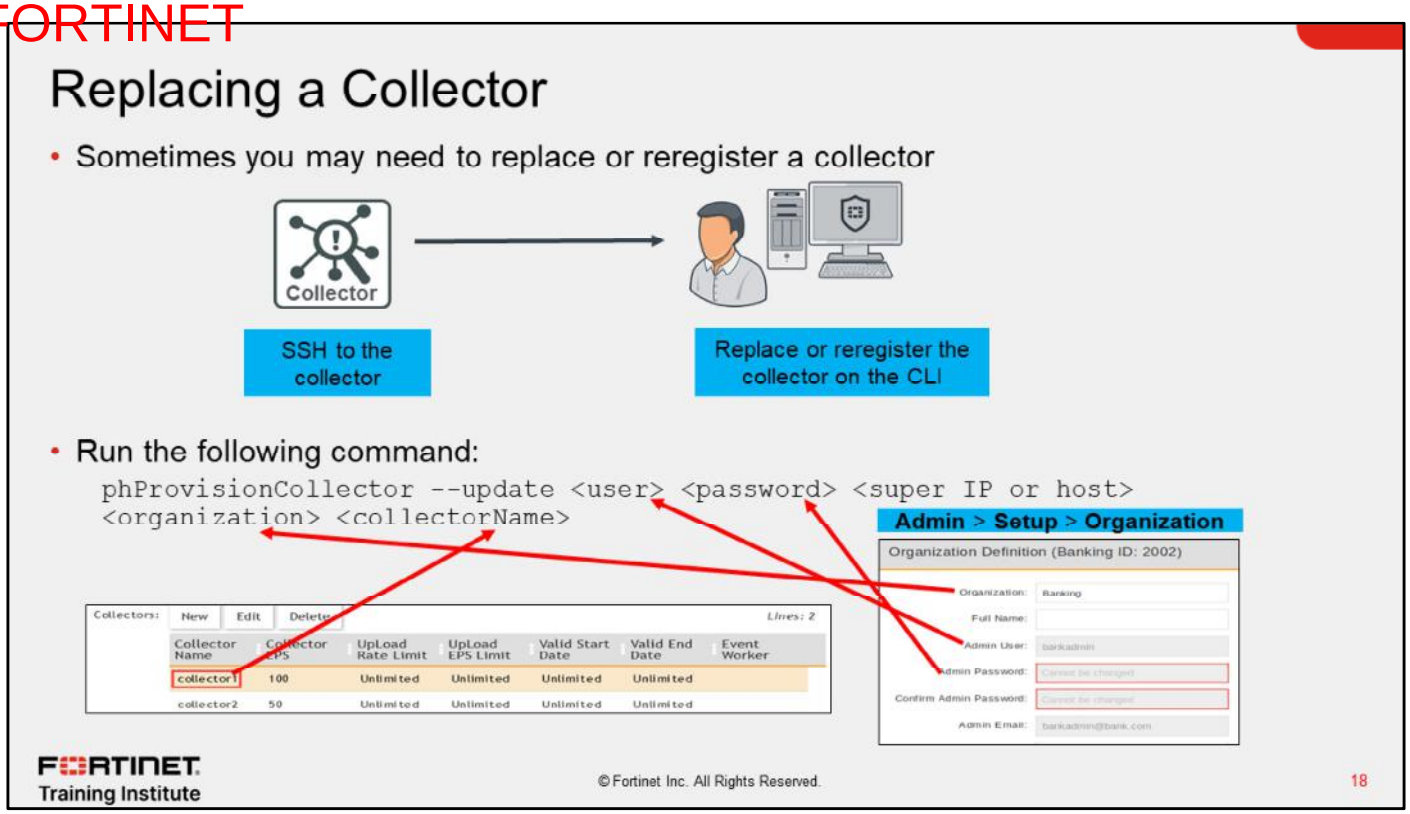
DO NOT REPRINT
© FORTINET



The Postgres database on the supervisor records a natural ID for the collector, which is actually the local UUID. The local UUID on the collector is recorded in the `product_uuid` file, the location for which is shown on this slide. The supervisor stores this value in a Postgres table along with the name, organization ID, and collector ID.

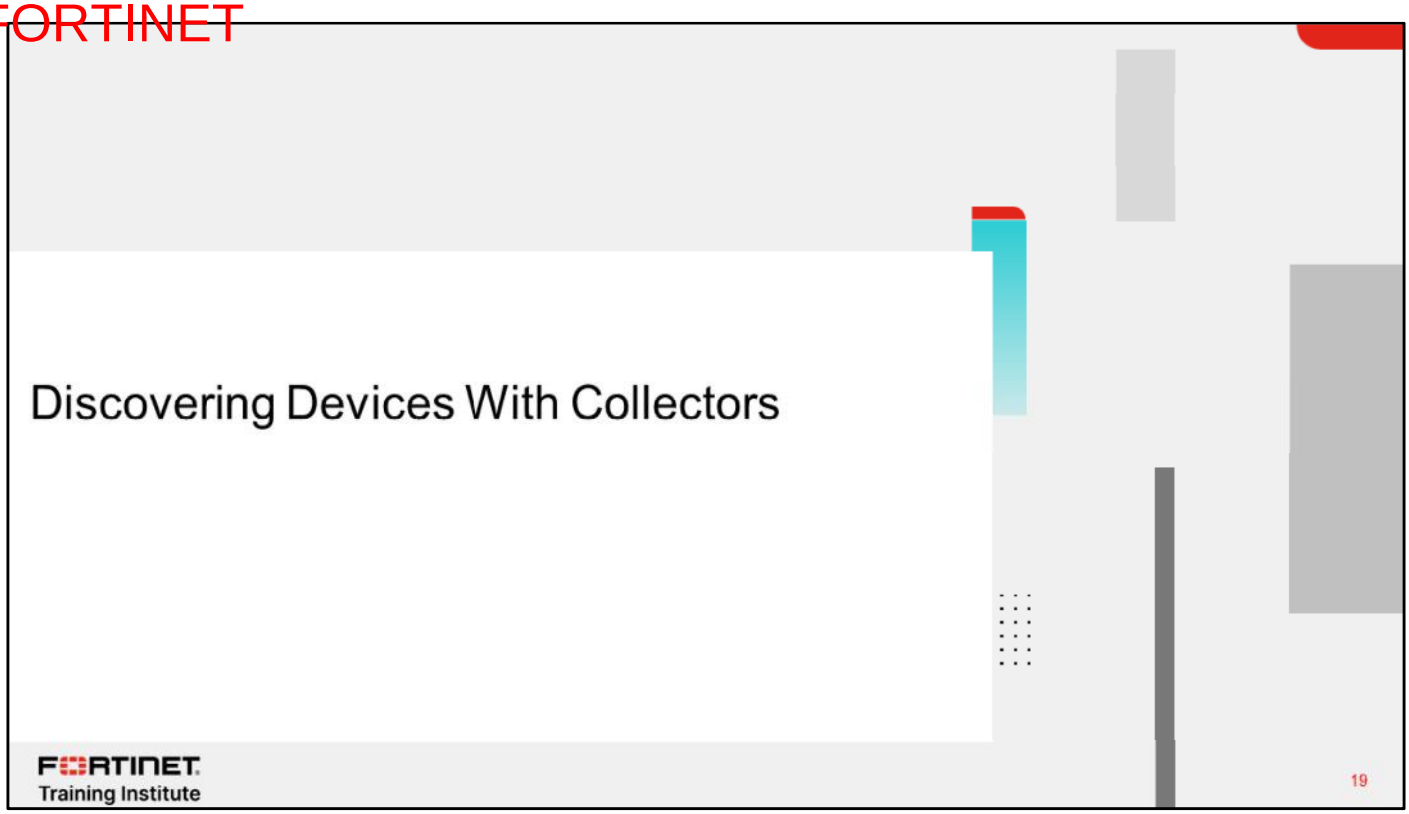
You can view the Postgres table by running the command shown on this slide. The output shows a table listing all of the collectors registered to the supervisor. If your collector is listed but it won't connect, then there was a problem during registration, which you should investigate.

DO NOT REPRINT
© FORTINET



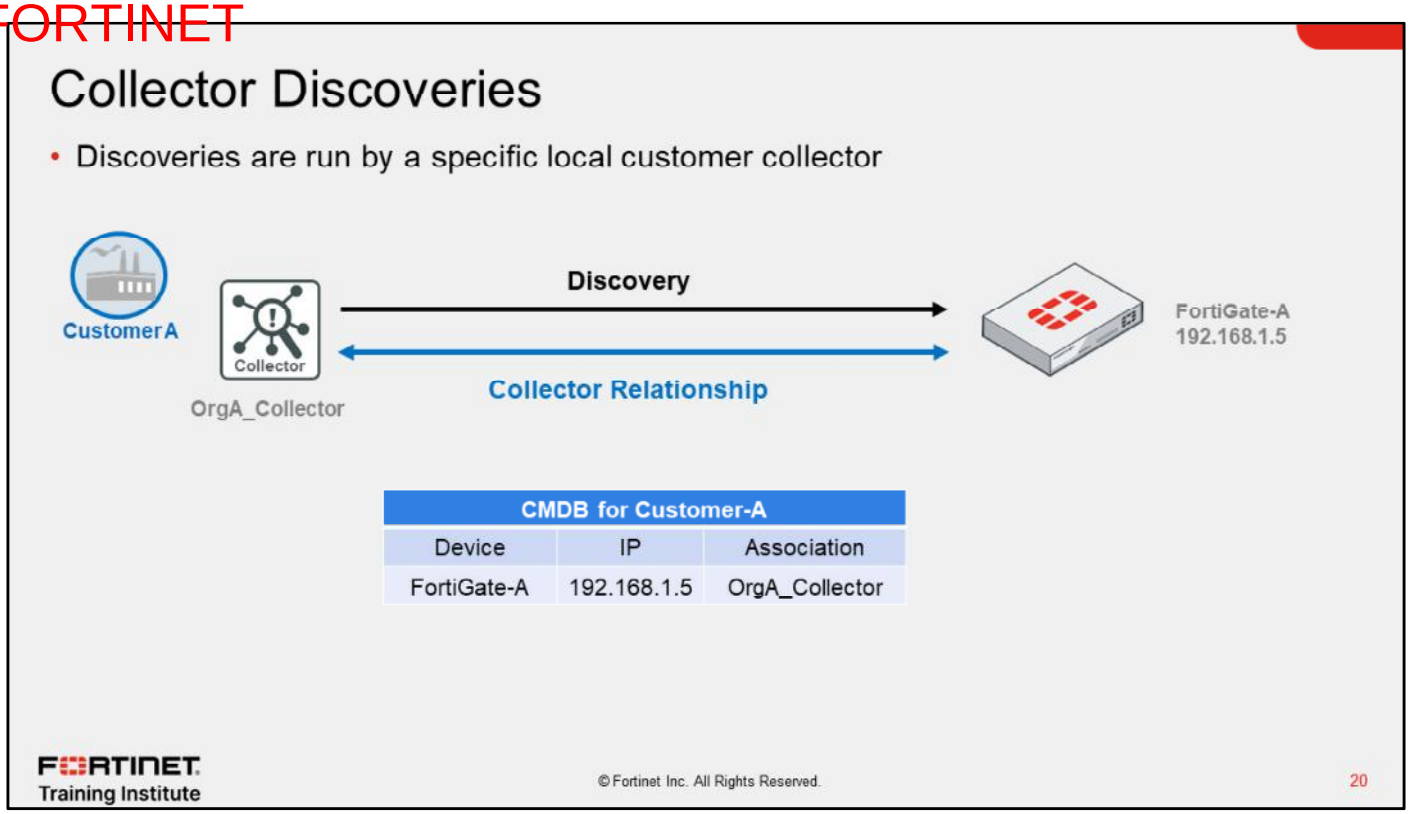
You can replace or reregister a collector by using command shown on the slide.

DO NOT REPRINT
© FORTINET



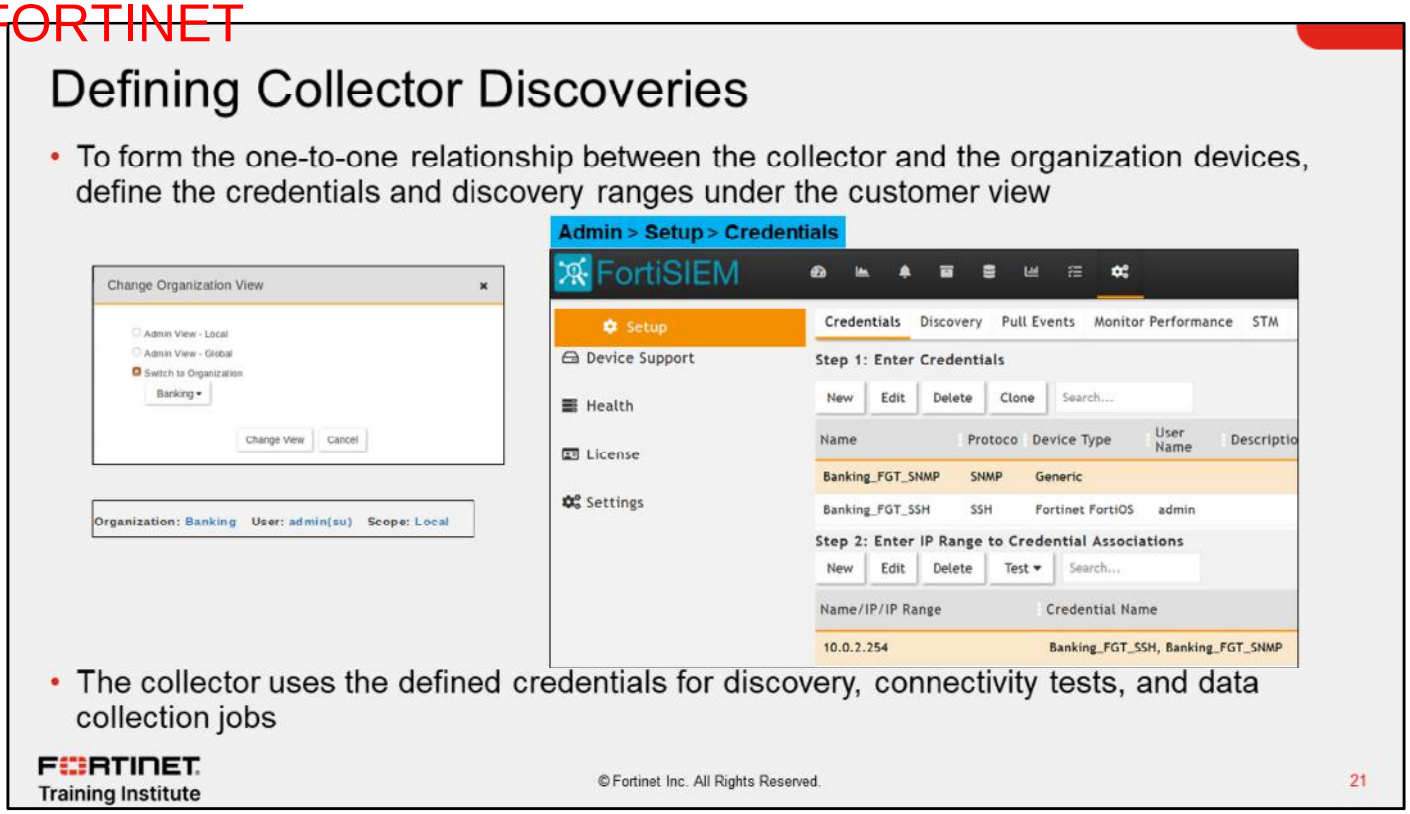
In this section, you will learn about discovering devices with the collector.

DO NOT REPRINT
© FORTINET



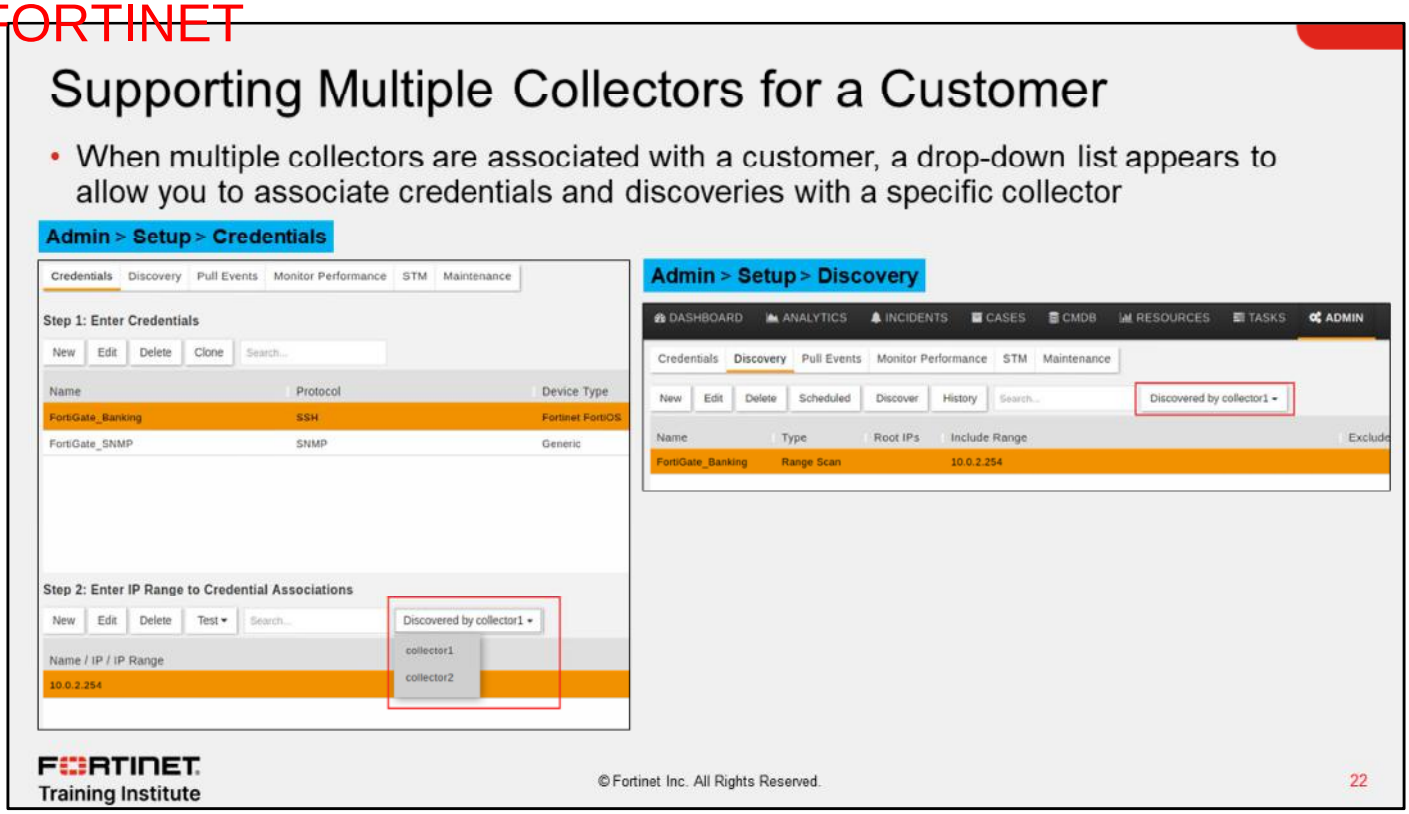
The customer's specific, on-premises collectors perform device discoveries on customer networks. This forms a one-to-one relationship between devices and the collector for all further tasks, such as performance and availability jobs, or pull event jobs, such as WMI event log collection.

DO NOT REPRINT
© FORTINET



To form the one-to-one relationship between devices and the collector, you must define the device credentials for that specific organization and associate them with the applicable IP ranges. Collectors for that organization use the credentials that you defined for discovery, connectivity tests, and data collection jobs.

DO NOT REPRINT
© FORTINET

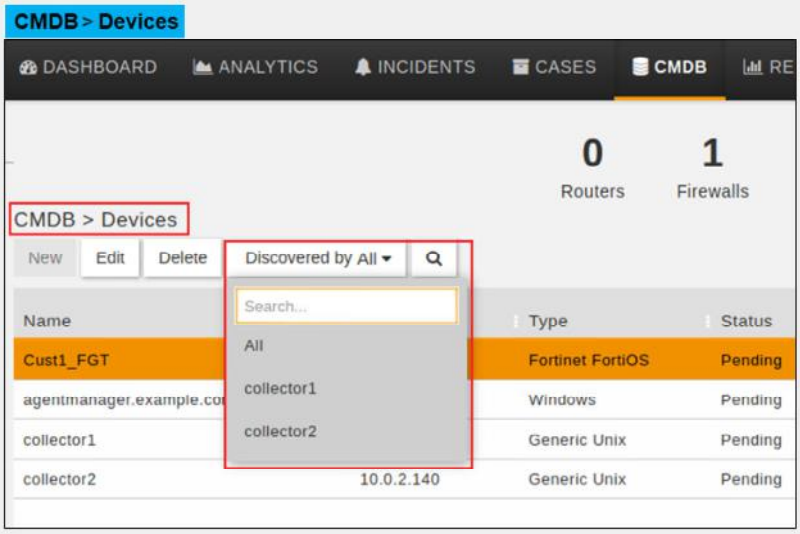


If a customer has multiple collectors, you can associate credentials and discovery targets with specific collectors. Only the collector that you associate the credentials with can discover devices for the IP range so that devices do not get duplicated in the CMDB. After you discover the devices, you can view which device was discovered by which collector.

DO NOT REPRINT
© FORTINET

Filtering the CMDB by Collector

- CMDB provides a filter to view which devices are associated with each customer collector



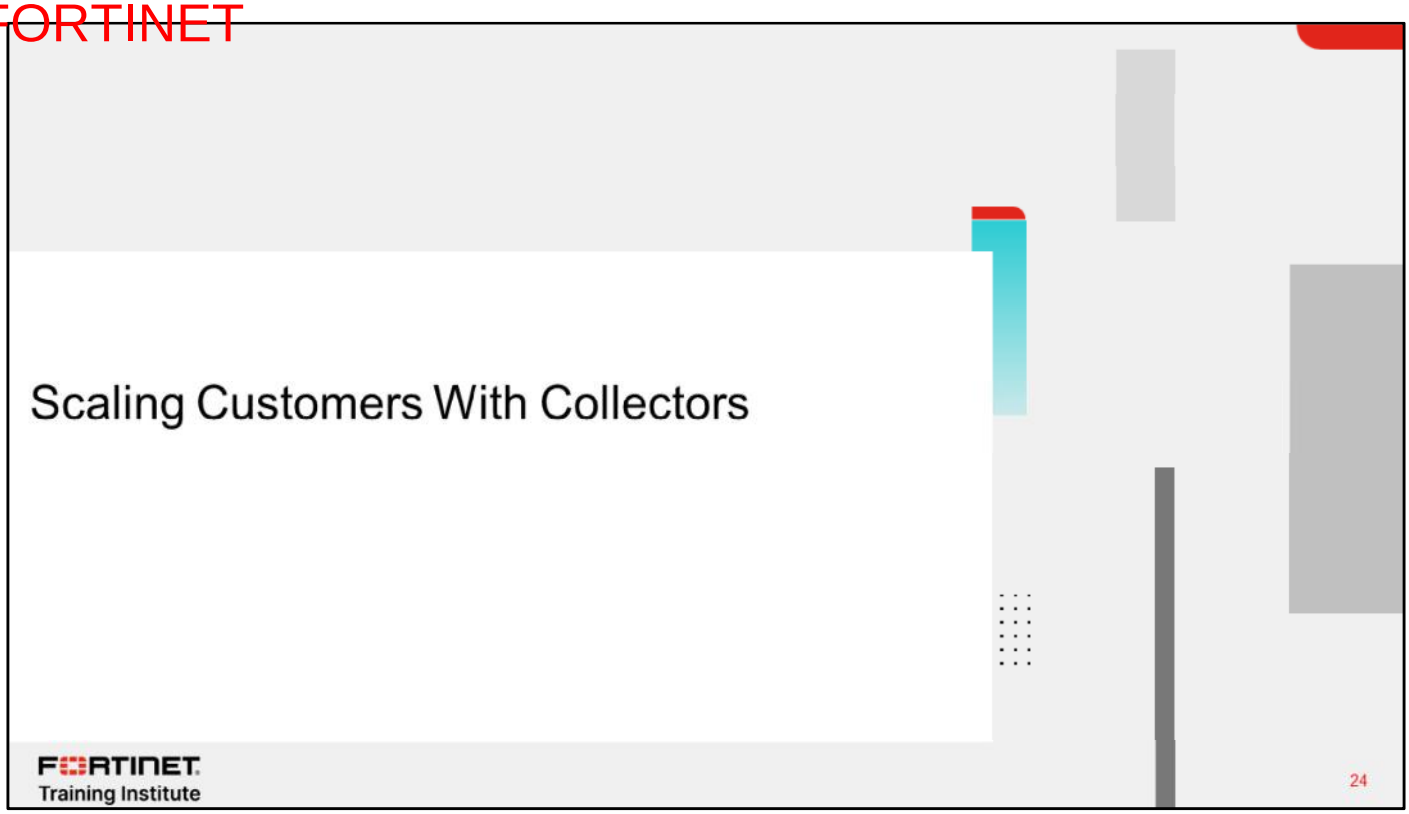
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

23

The CMDB provides a filter to view which devices are associated with each customer collector. If you select **Discovered by All**, then all the devices discovered by all the collectors for that organization are displayed. You can also select the individual collector alone, which shows only the devices discovered by that specific collector.

DO NOT REPRINT
© FORTINET

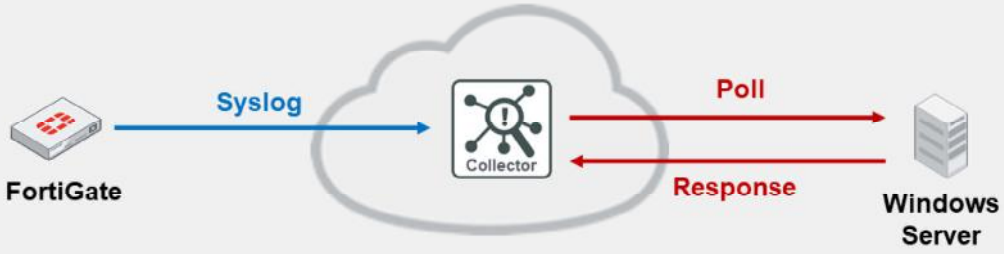


In this section, you will learn about scaling customers with collectors.

DO NOT REPRINT
© FORTINET

Scaling Collectors at Customer Sites—Log Collection

- Two things to take into consideration:
  - Direct log collection
    - Devices reporting syslog messages to the collector
  - Performance monitoring and SIEM collection jobs
    - SNMP, WMI collection jobs of CPU, memory, interface stats, and more
    - WMI collection of Windows event logs



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

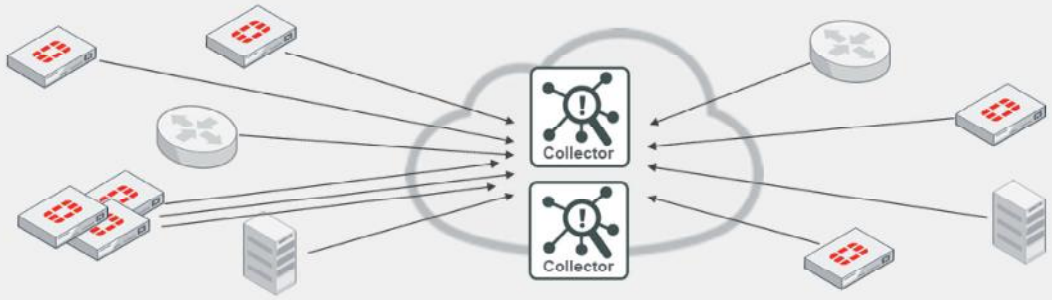
25

If you need to scale collectors at a customer location, you must take two things into consideration. The first is direct log collection from devices that report syslog messages to collectors. The second is performance monitoring jobs that are performed by the collector by polling devices using SNMP or the WMI protocol.

DO NOT REPRINT
© FORTINET

Scaling Collectors at Customer Sites—Log Collection (Contd)

- As the incoming EPS received by collectors increases, you must deploy more collectors, or increase resources



* **Resource** estimates

Resources	EPS	Resources	EPS
2 vCPU, 4GB Ram	500	8v CPU, 12 GB Ram	6-7k
4v vCPU, 8GB Ram	5k	16 vCPU, 16GB Ram	10k

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

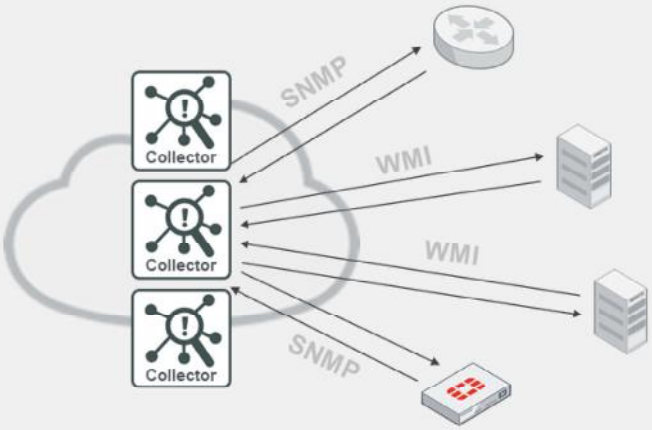
26

As the incoming EPS received by collectors increases, you must deploy more collectors, or increase the resources on the existing collectors. For example, for an incoming EPS of 10,000, the collector needs at least 16 vCPU and 16 GB of RAM.

DO NOT REPRINT
© FORTINET

Scaling Collectors at Customer Sites—Polling Jobs

- As performance jobs and SIEM collection jobs increase, you should deploy more collectors
  - No more than 300 devices per collector for SNMP
  - No more than 100 devices per collector for WMI event log collection
  - SNMPv3 increases the load on collectors, so reduce values to accommodate



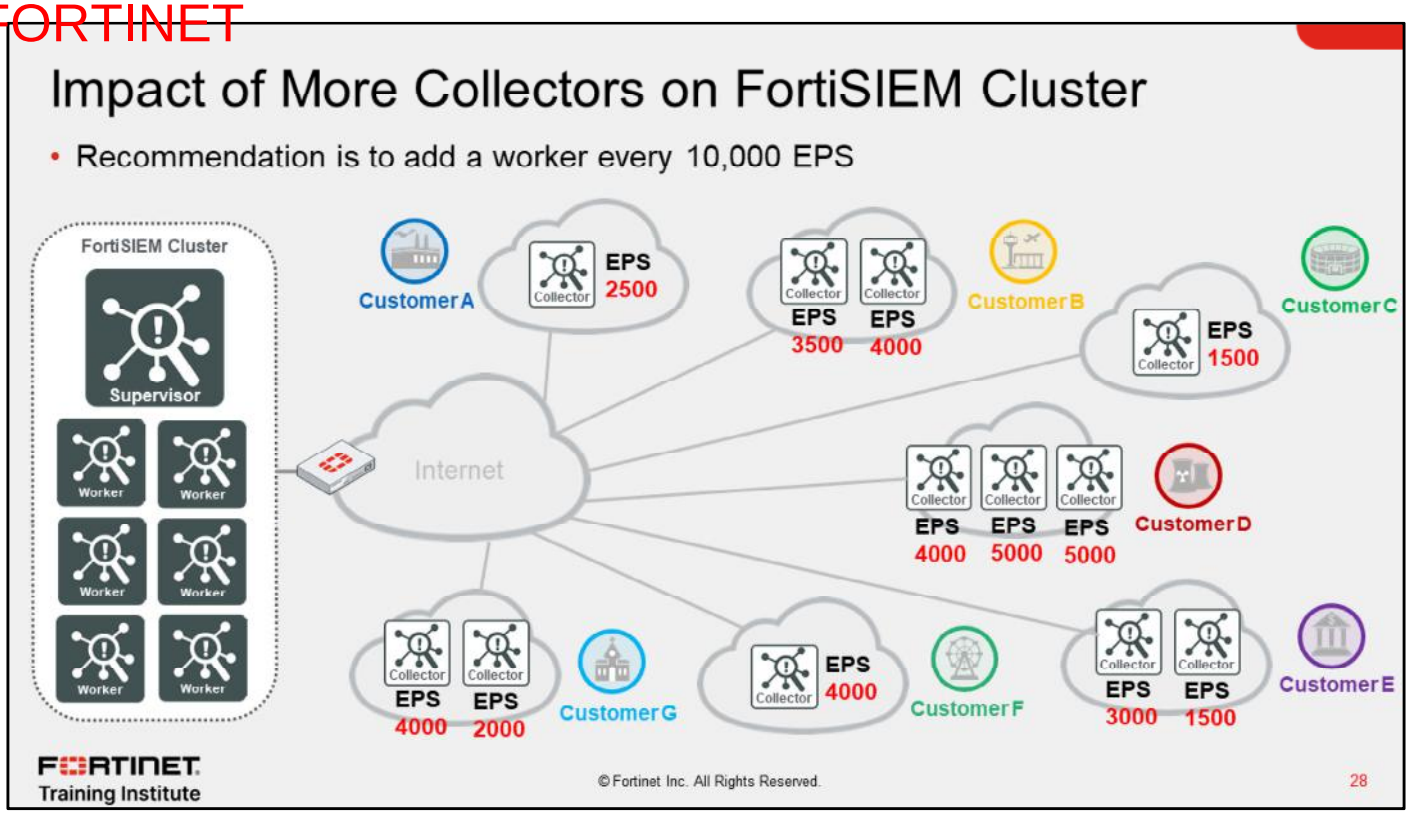
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

27

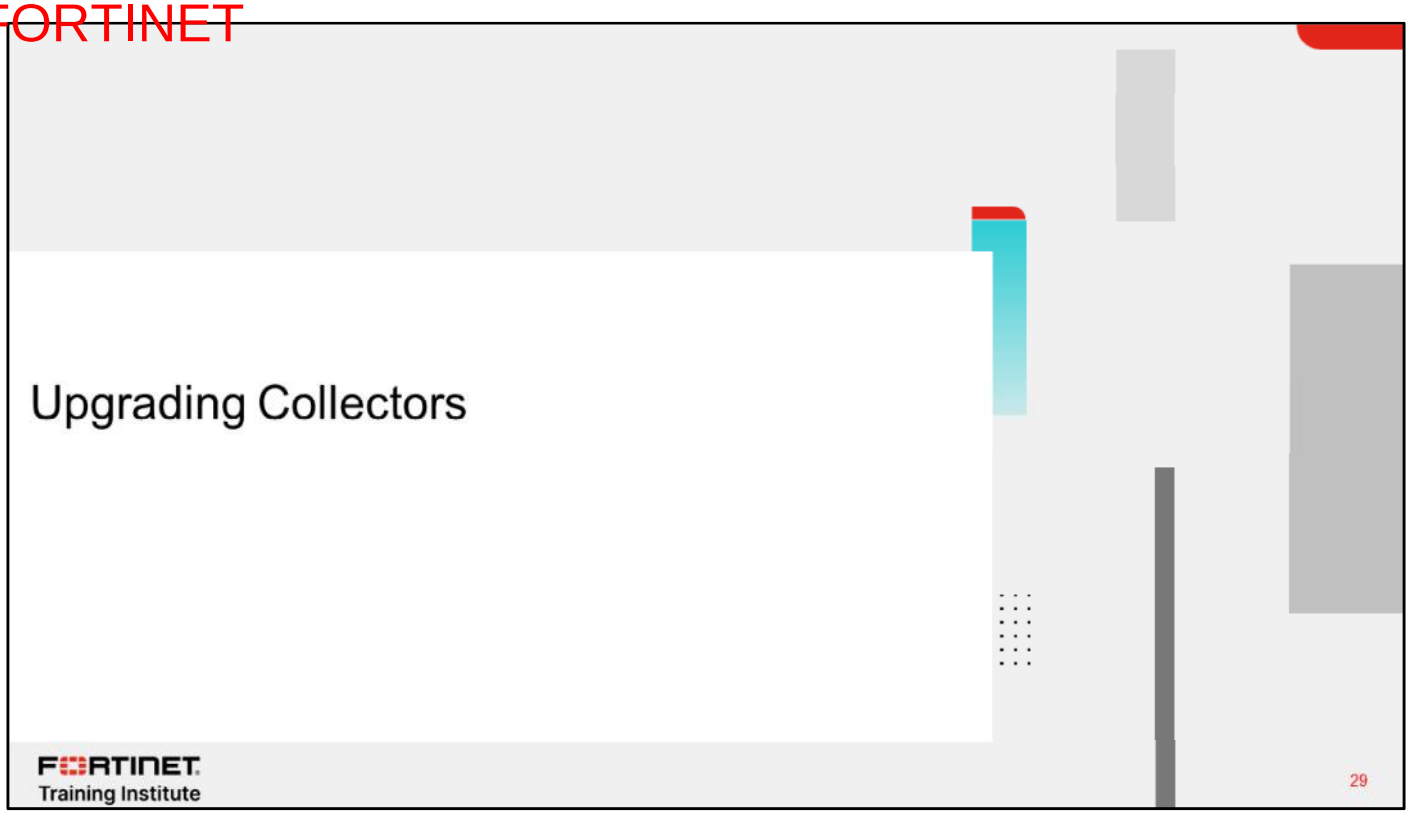
For polling jobs, the collector uses either SNMP or the WMI protocol. Unlike syslog, where the devices sent logs to the collector, the SNMP and WMI protocols require the collector to poll the devices periodically for performance and security metrics. The fact that the collector must actively poll devices requires additional processing overhead. You should not have more than 300 devices per collector for SNMP polling. If you're using SNMPv3, then that device limit is even lower because SNMPv3 is more resource intensive. You should not have more than 100 devices per collector when using WMI event log collection. You must add more collectors if the number of devices that you are monitoring keeps growing.

DO NOT REPRINT
© FORTINET



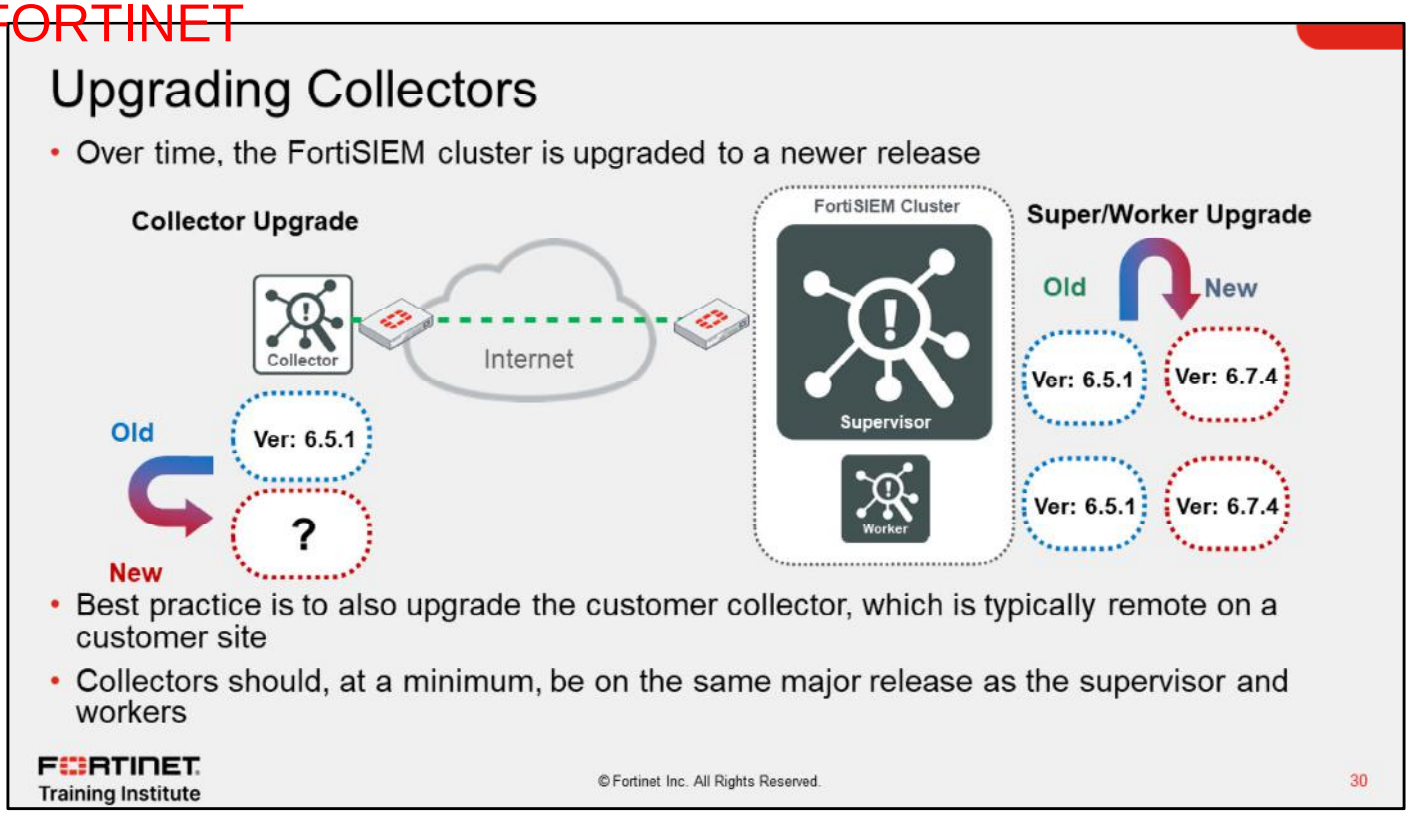
As you keep adding collectors to organizations to scale their growth, you also must take into account the FortiSIEM cluster processing load. With more EPS to process, the cluster needs additional processing power. For that, you should add more workers. Ideally, you should add a new worker for every 10,000 increase in EPS.

DO NOT REPRINT
© FORTINET



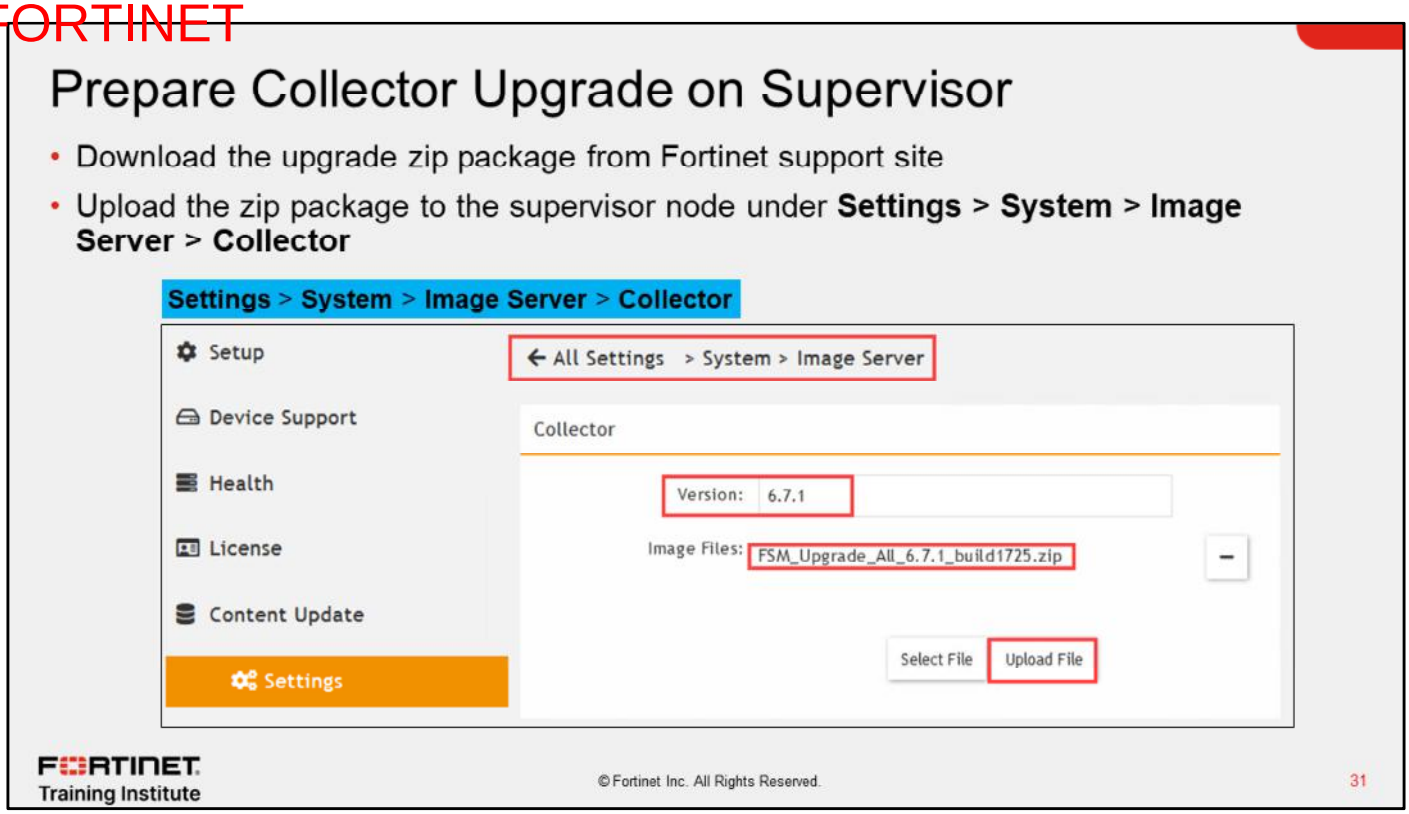
In this section, you will learn about upgrading collectors.

DO NOT REPRINT
© FORTINET



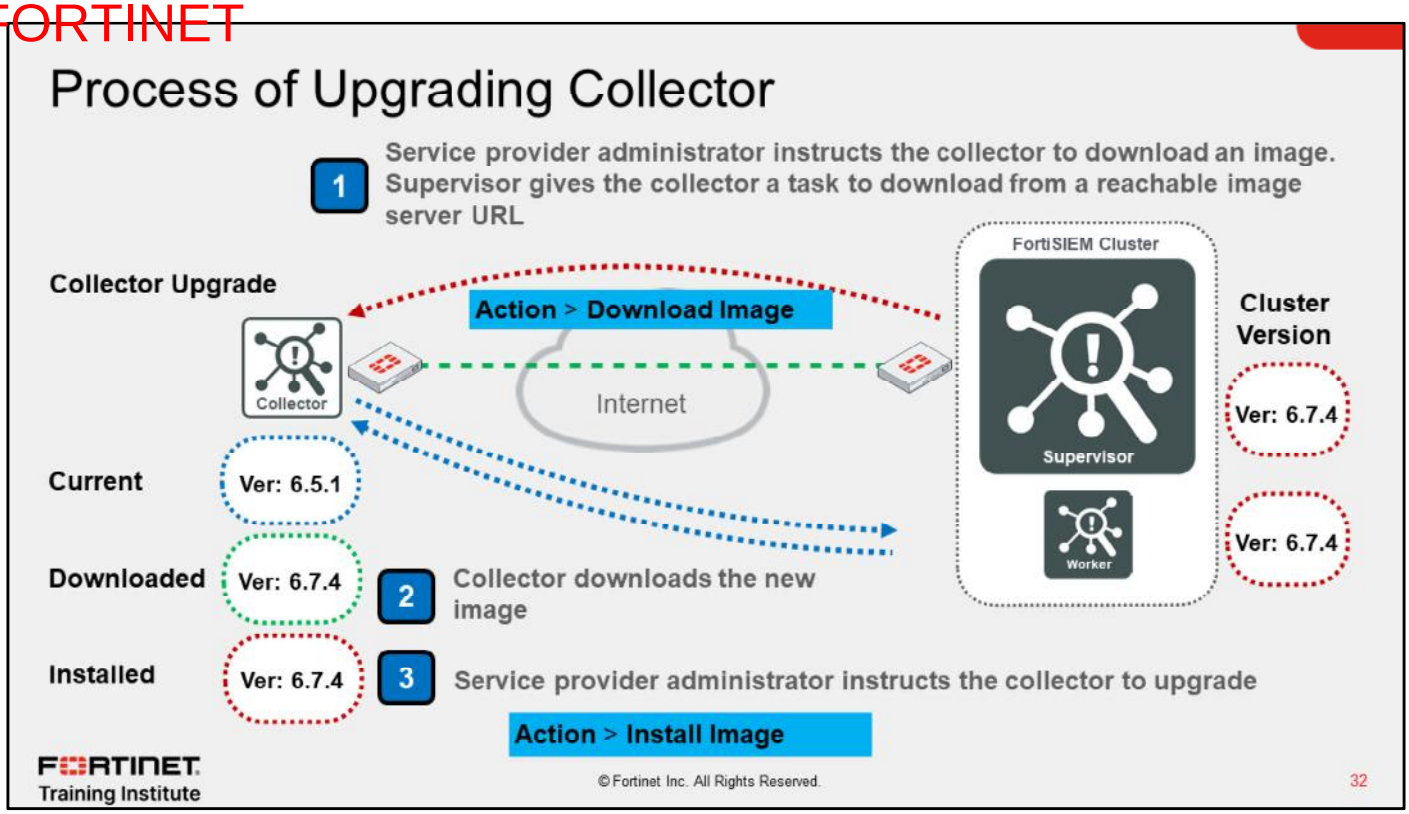
Collectors should, at a minimum, be on the same major release as the supervisor and workers. For example, if the supervisor and worker cluster is on version 6.7.4, the collectors could be on version 6.6.0. If the supervisor and worker cluster is on version 6.7.4, collectors should not be on version 4.10 or issues may arise. Major releases sometimes introduce new functionality that collectors on older releases may not support. Always follow the upgrade process mentioned in the *Upgrade Guide* found on `docs.fortinet.com`.

DO NOT REPRINT
© FORTINET



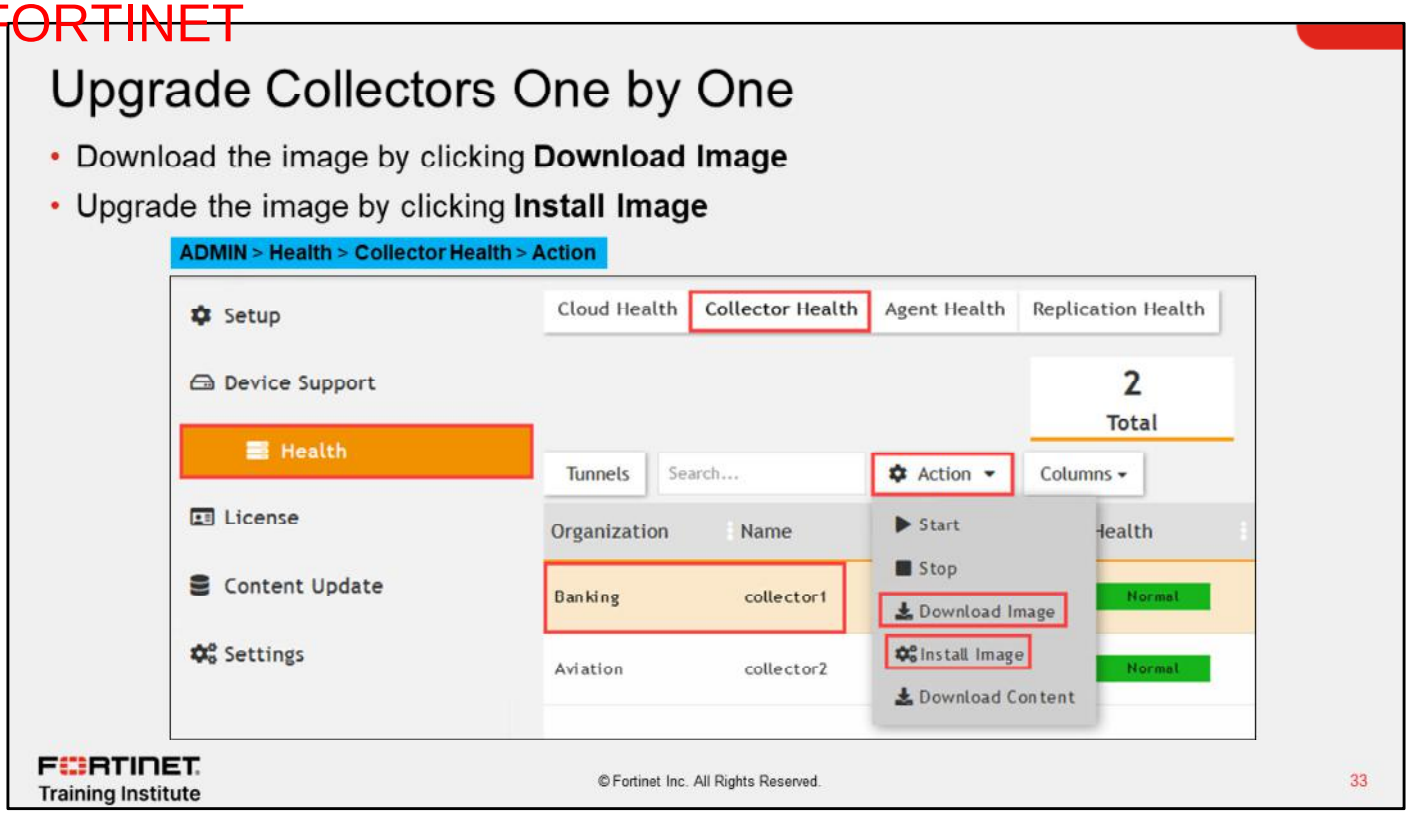
Before you can begin to upgrade collectors for organizations, you must prepare the supervisor so that the collectors can download the upgrade image from the supervisor and then upgrade. Download the upgrade zip package from the Fortinet support site, and then upload it to the supervisor node under the **Settings** > **System** > **Image Server > Collector**.

DO NOT REPRINT
© FORTINET



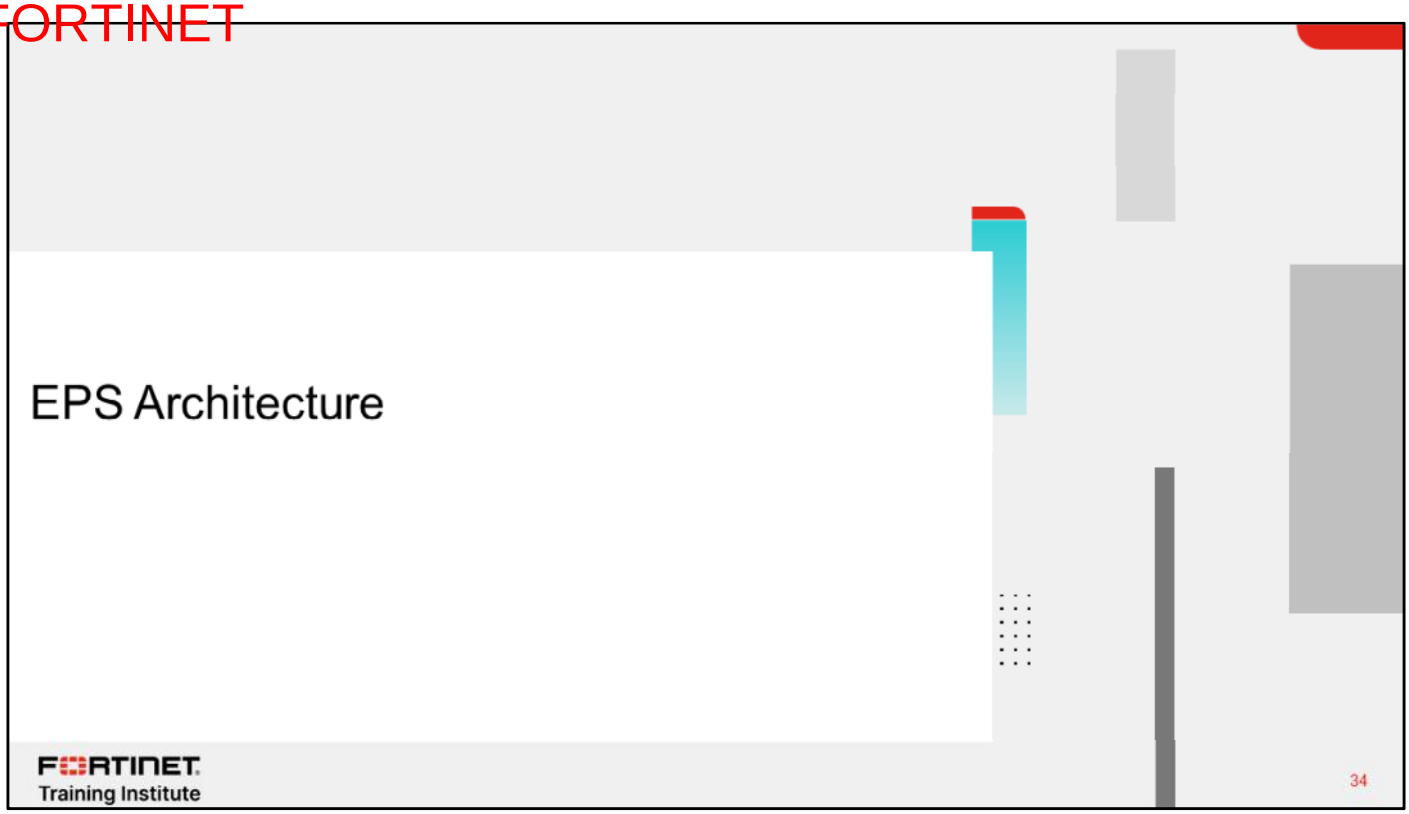
The service provider administrator who is logged in to the supervisor instructs the collector to download the collector upgrade image. The supervisor gives the collector a task to download the image from a reachable image server URL, which, in this case, is the supervisor. The collector downloads the new image. The service provider administrator then instructs the collector to upgrade itself.

DO NOT REPRINT
© FORTINET



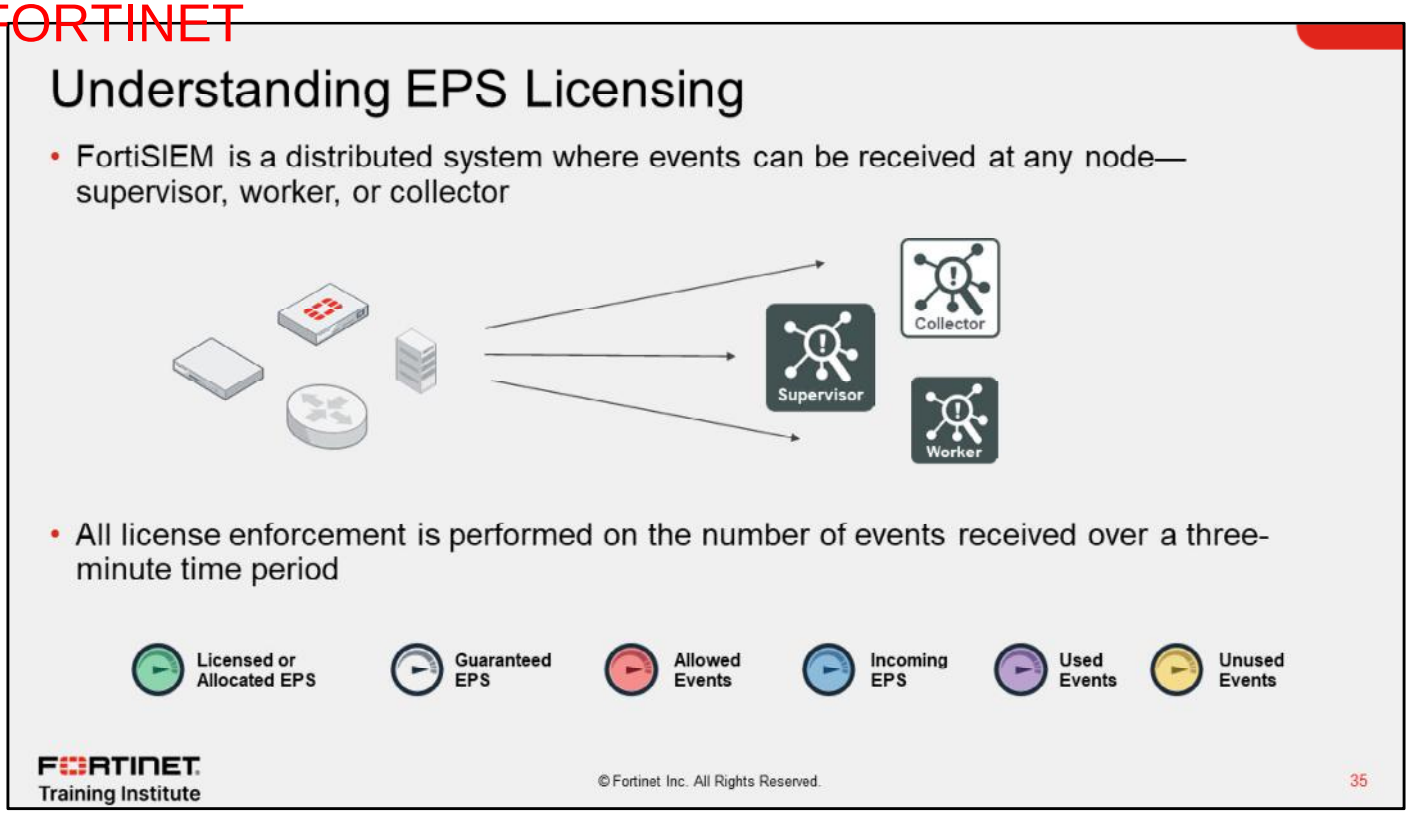
Log in to the FortiSIEM GUI and navigate to the **Collector Health** page. Select the collector that you want to upgrade and click **Action**. In the drop-down list, select **Download Image**, which instructs the collector to download the collector upgrade image. The service provider administrator then instructs the collector to upgrade itself by clicking **Install Image**.

DO NOT REPRINT
© FORTINET



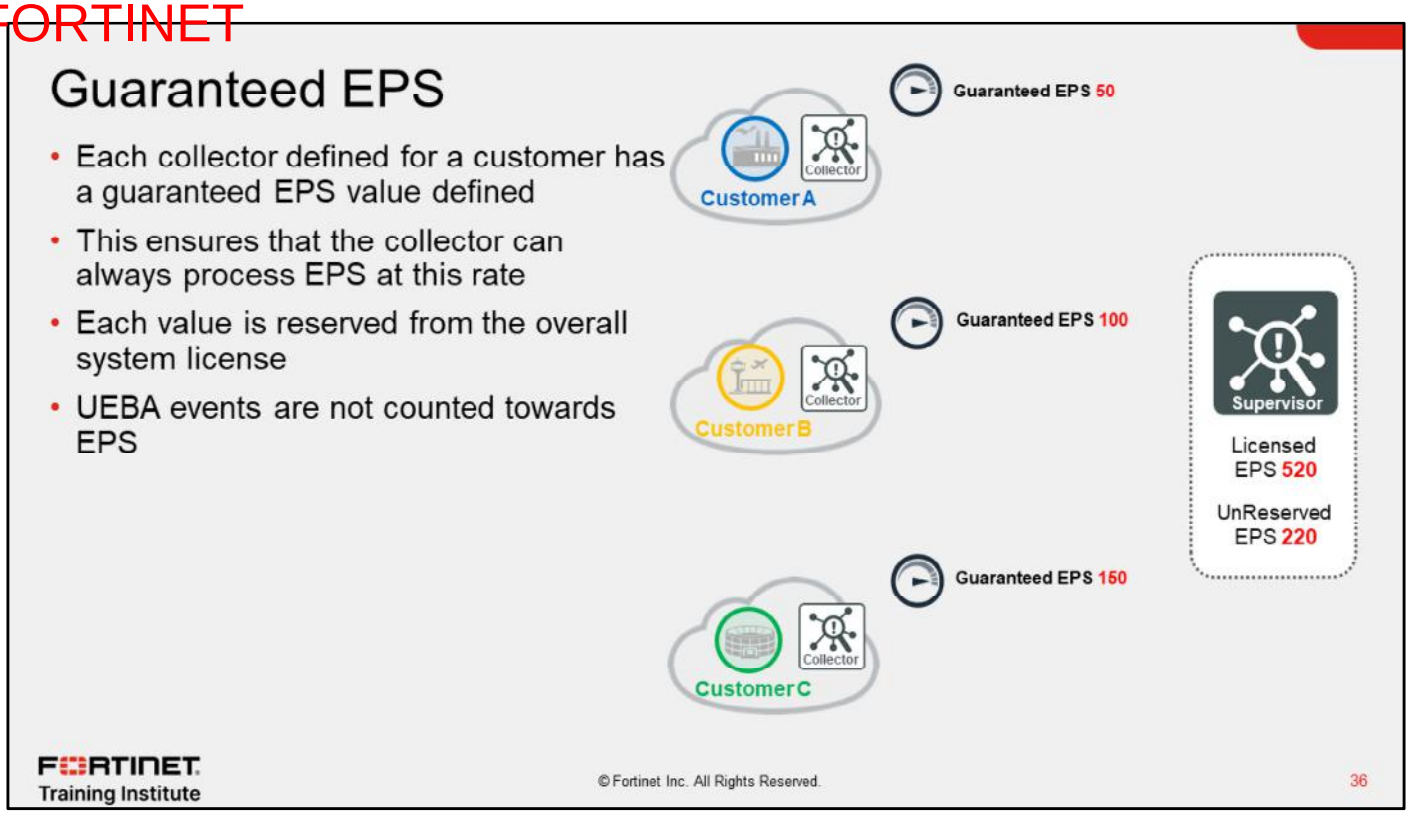
In this section, you will learn about EPS architecture.

DO NOT REPRINT
© FORTINET



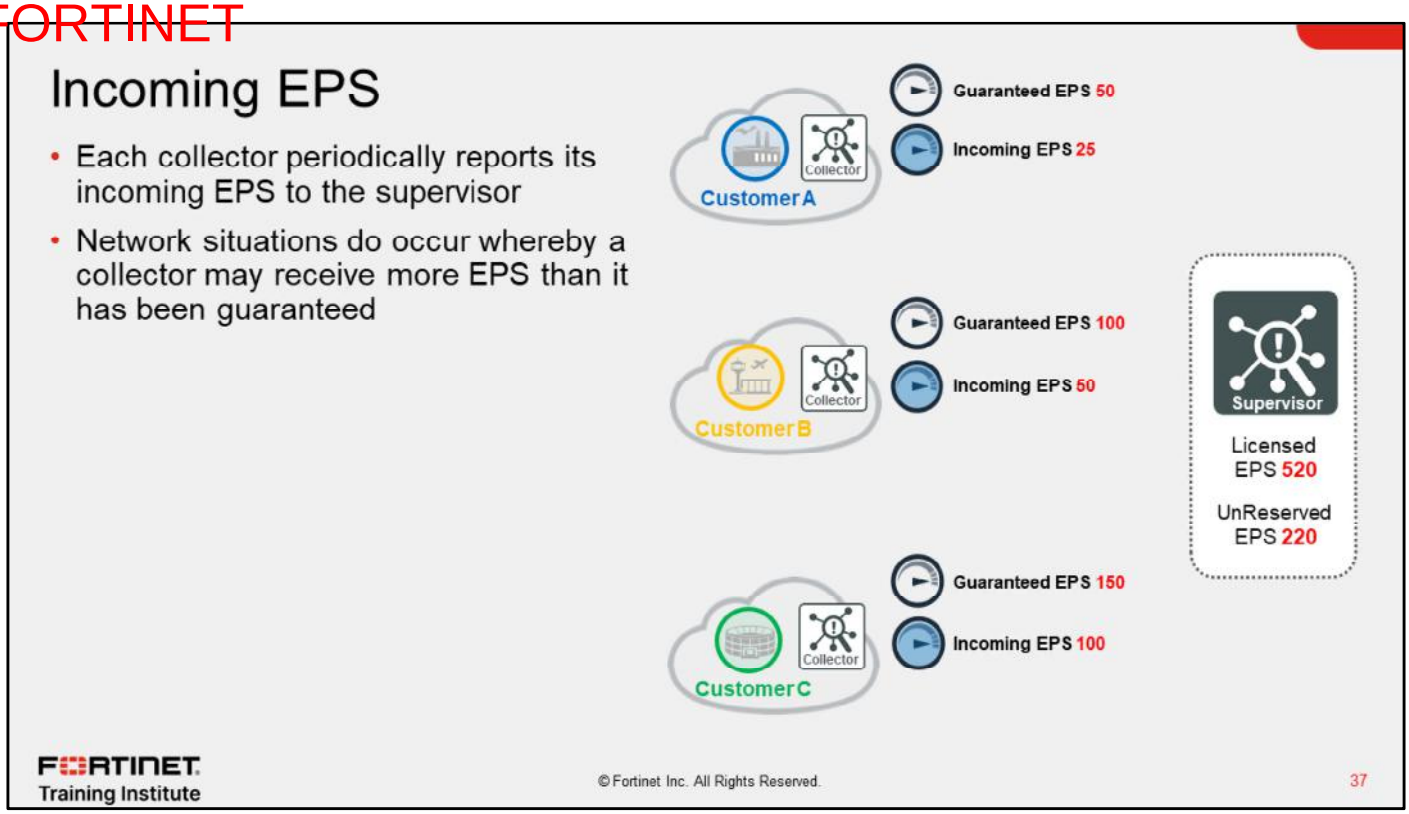
FortiSIEM is a distributed system where events can be received at any node, such as a supervisor, worker, or collector. One of the key requirements for FortiSIEM licensing is the aggregate EPS that FortiSIEM receives. FortiSIEM calculates EPS over a three-minute period. The total number of events received over a three-minute period is divided by 180. So, sometimes EPS is referred to as events within a three-minute time period.

DO NOT REPRINT
© FORTINET



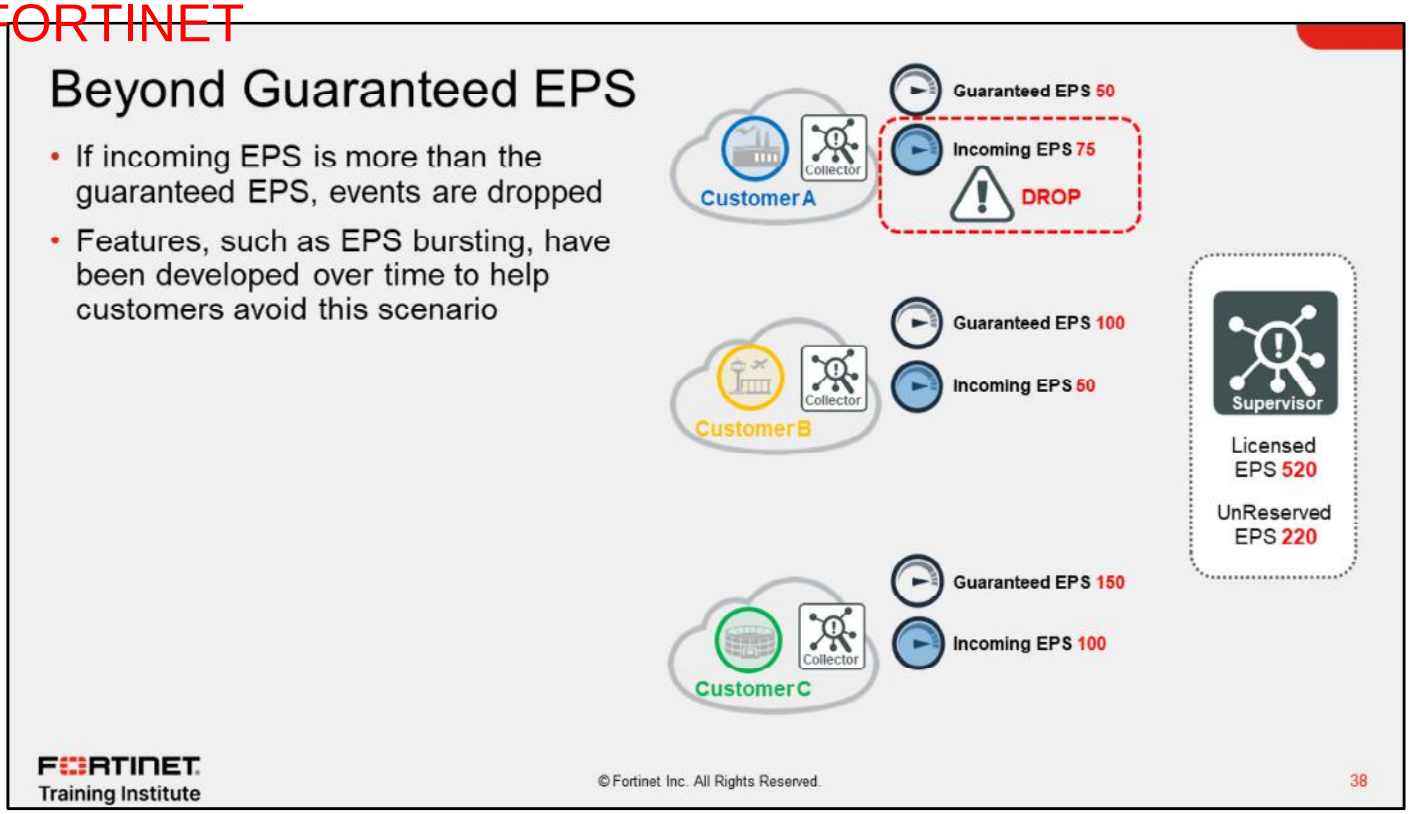
You define the guaranteed EPS during the collector configuration process while setting up an organization. FortiSIEM ensures that the collector can always send EPS at this rate. This is a constant that never changes during the operation of the collector, unless you modify the collector definition. UEBA events are not counted toward EPS.

DO NOT REPRINT
© FORTINET



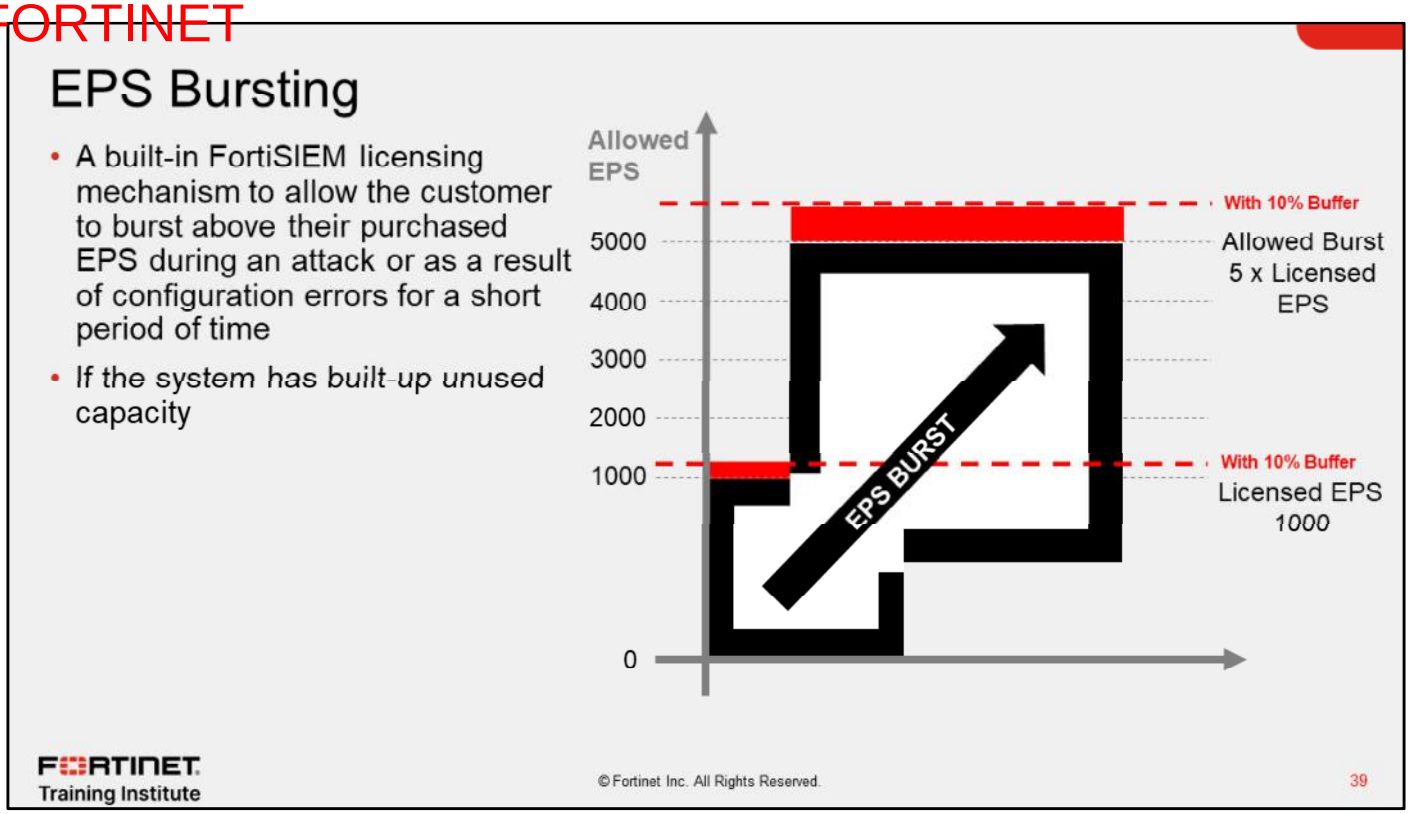
Incoming EPS is the EPS that the collector sees, and this value changes continuously. Every collector periodically reports the incoming EPS to the supervisor. You can view this metric for the collector on the **Collector Health** page.

DO NOT REPRINT
© FORTINET



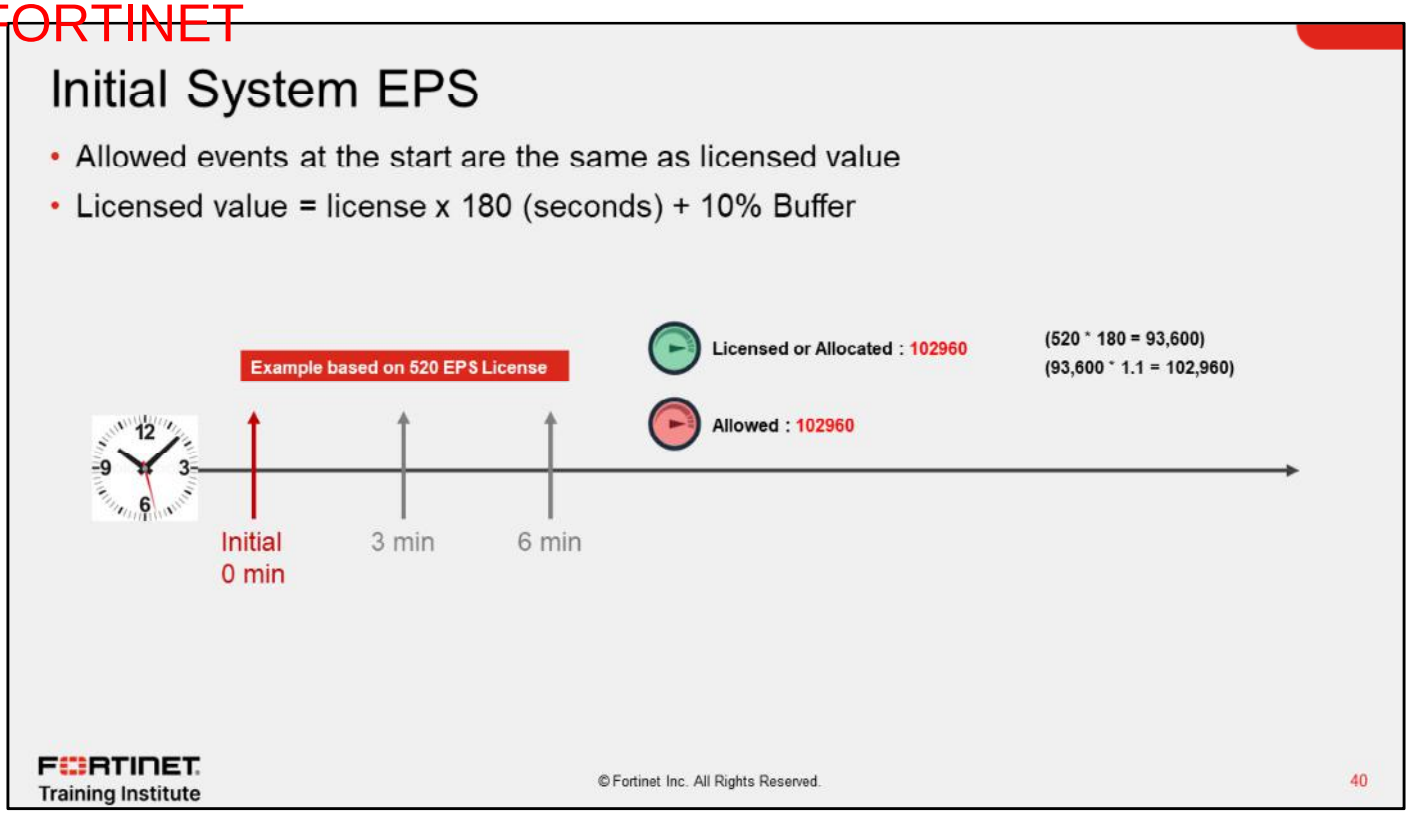
Usually, if the incoming EPS is greater than the guaranteed EPS, then the collector begins to drop events until the incoming EPS is less than or equal to the guaranteed EPS. However, features such EPS bursting have been developed over time to help service providers avoid this scenario.

DO NOT REPRINT
© FORTINET



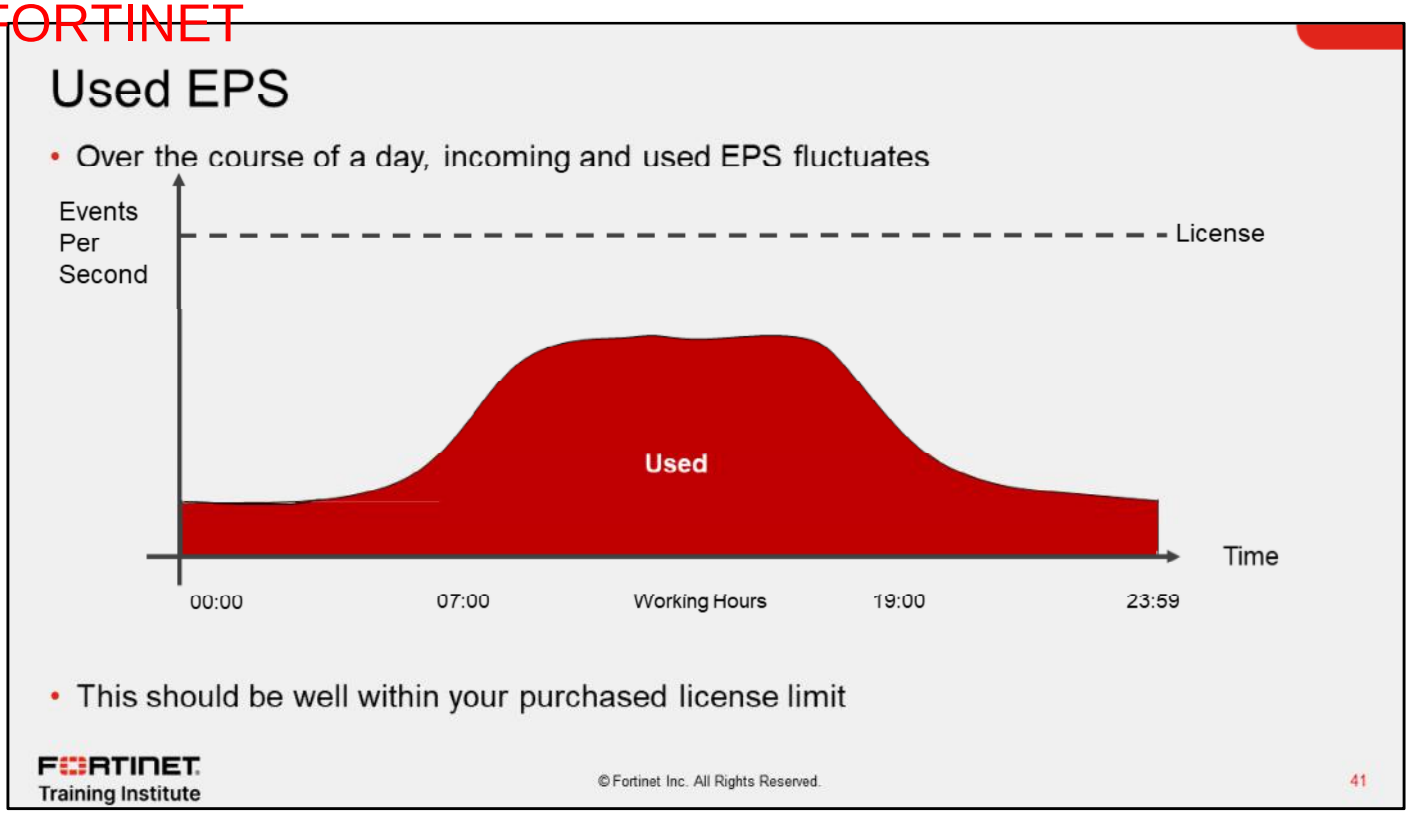
Before firmware version 4.10, FortiSIEM dropped events that exceeded 110% of the licensed EPS. Starting with the 4.10 release, FortiSIEM adds a reservoir of events to reduce the possibility of events being dropped. This allows greater EPS bursts without event loss. FortiSIEM lets you burst up to five times the licensed EPS using the currently accumulated unused EPS. To benefit from this EPS bursting, you must have enough computational power and storage to process the extra EPS. You must provision the system with five times the licensed EPS to handle potential event surges.

DO NOT REPRINT
© FORTINET



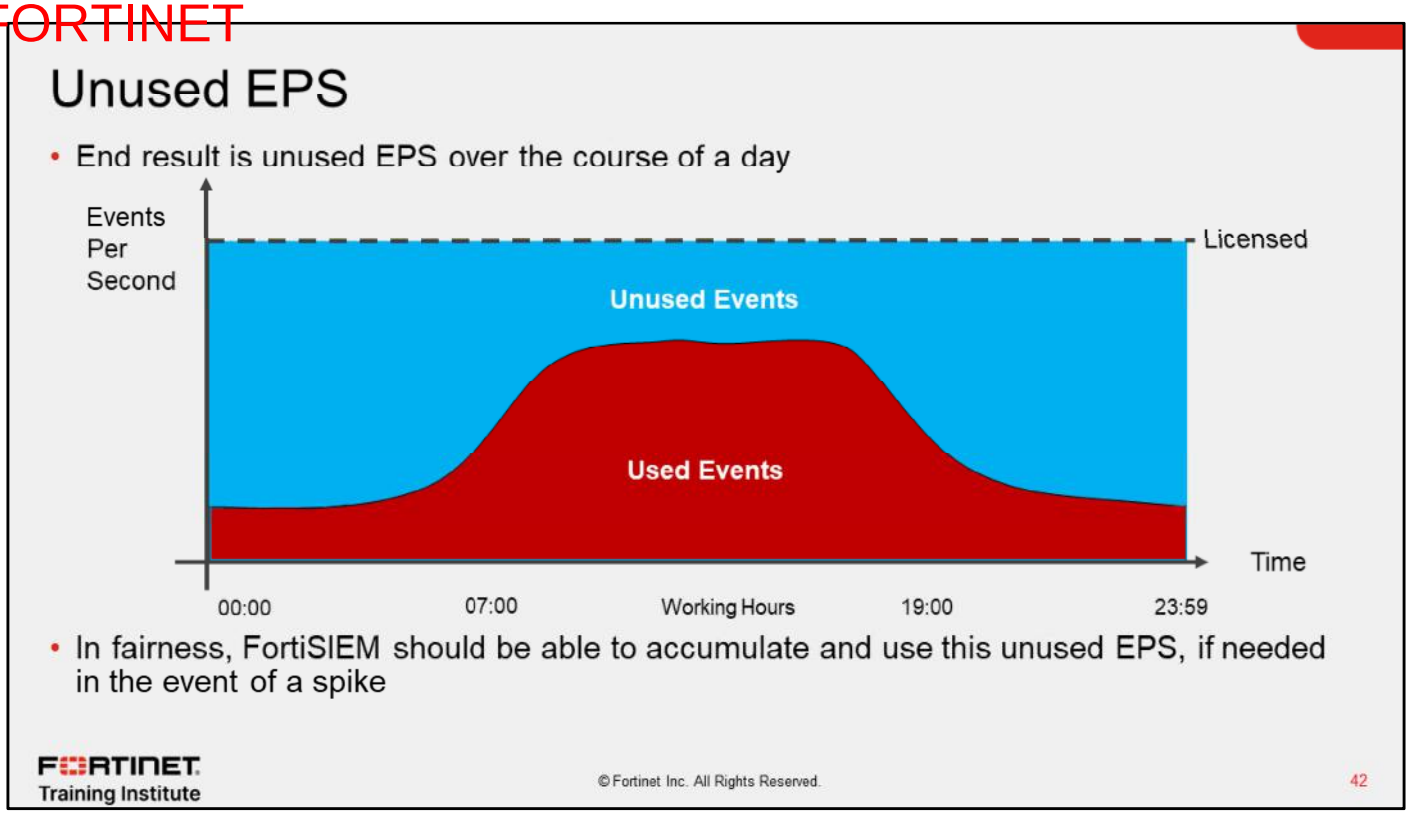
When a node is first deployed, the allowed events is the same as the licensed EPS value for a three-minute duration. This example is based on a 520 EPS license, so the allocated EPS for a three-minute duration is 10,2960, which includes the 10% buffer.

DO NOT REPRINT
© FORTINET



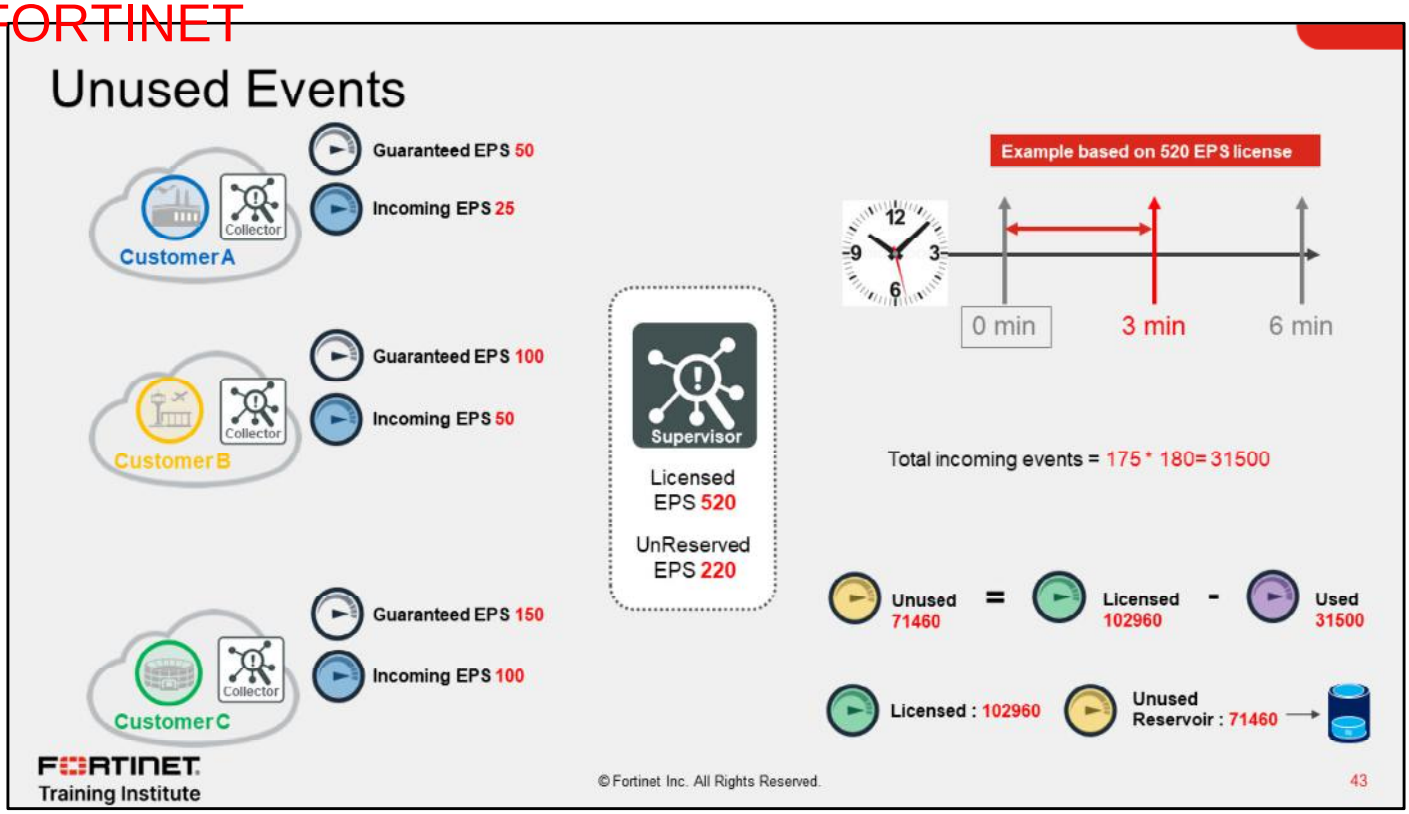
Over the course of the day, the incoming and used EPS fluctuates, but it should be within your purchased license limit during normal operation.

DO NOT REPRINT
© FORTINET



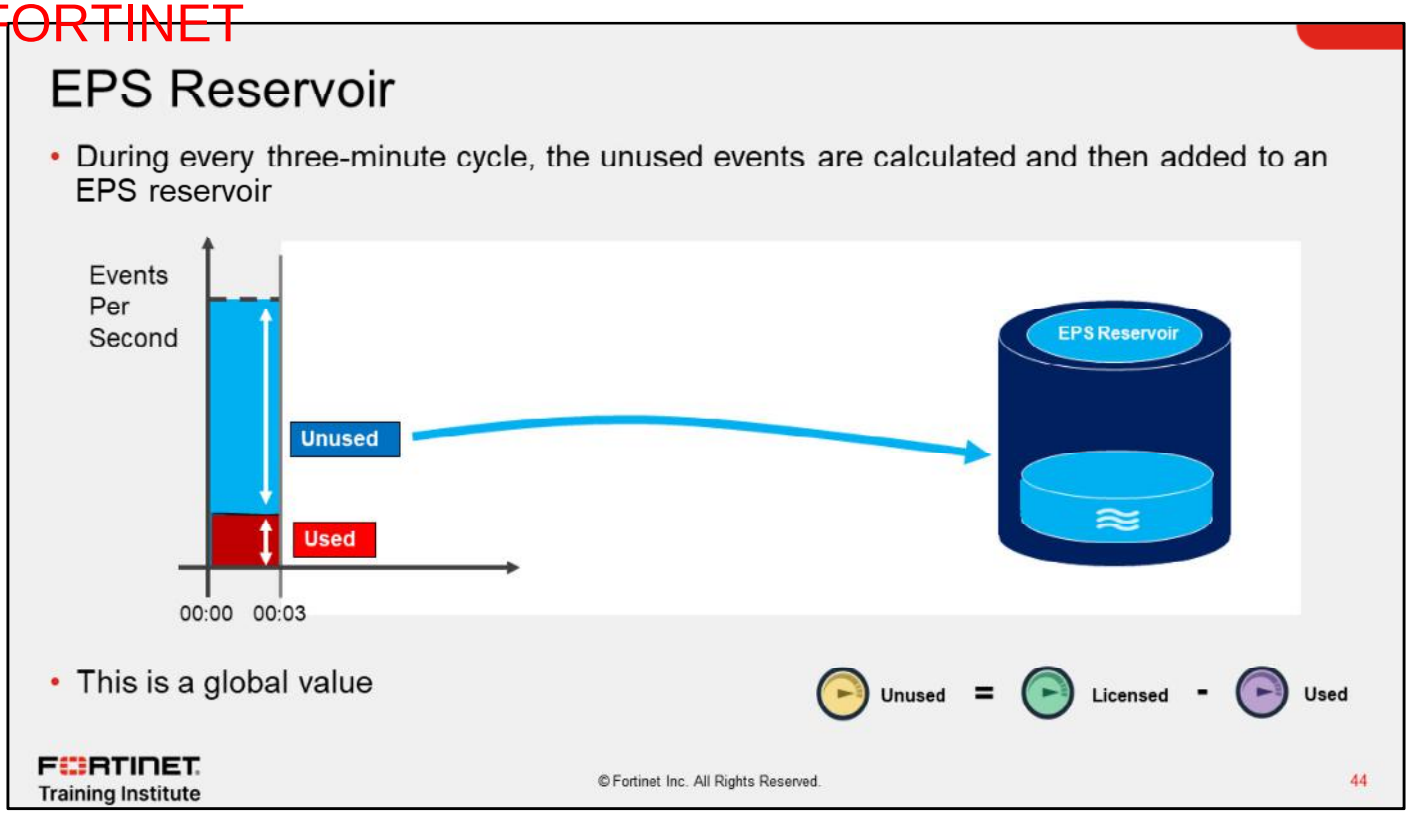
FortiSIEM keeps track of unused EPS as the sum of positive differences of allocated EPS and incoming EPS over all nodes. FortiSIEM can use unused EPS for bursting during attacks or other event surge periods, beyond licensed EPS.

DO NOT REPRINT
© FORTINET



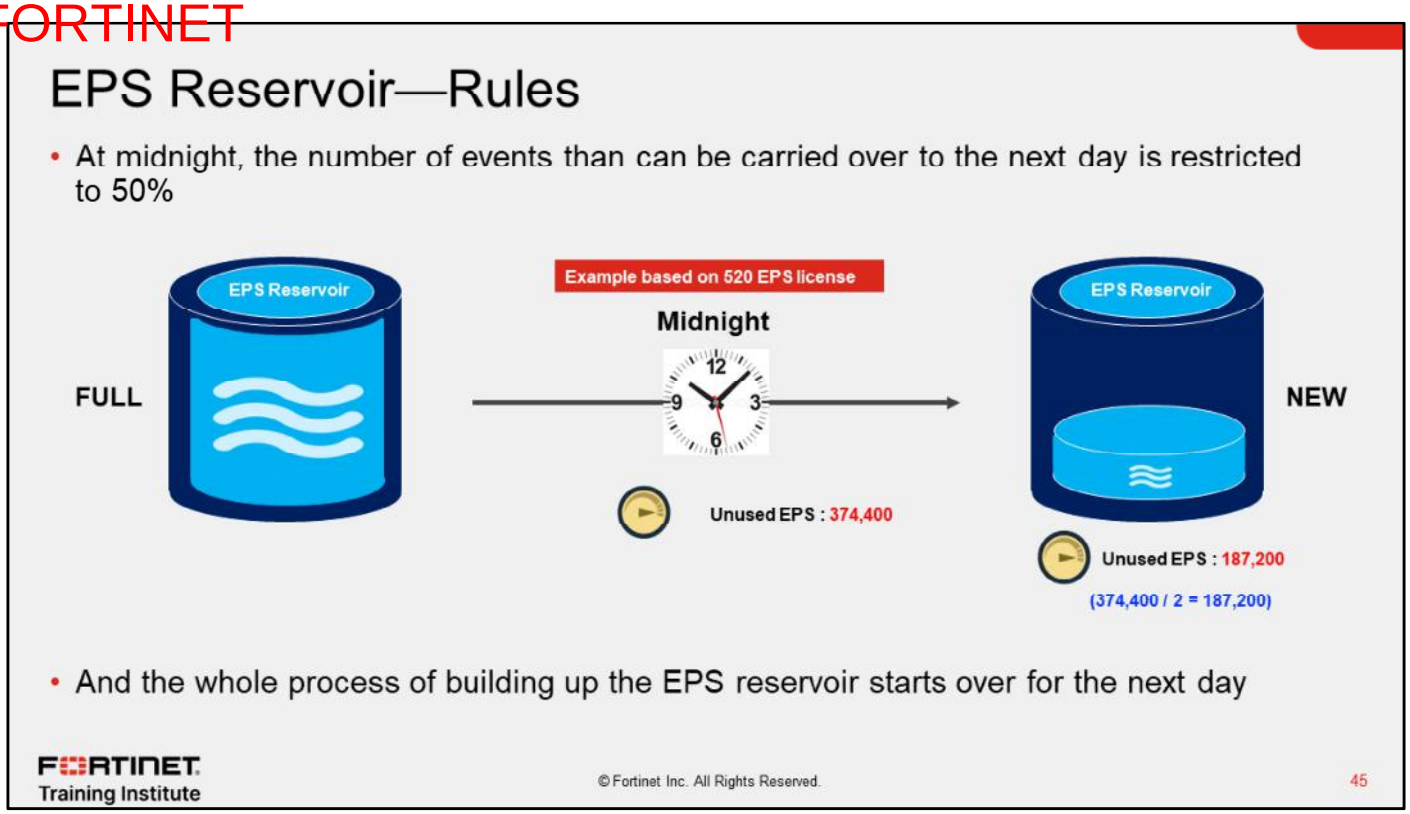
At the end of every three-minute interval, incoming EPS is calculated at each event entry node—collector, worker, or supervisor—and the value is sent to the central decision-making engine on the supervisor node. The supervisor calculates the unused events by subtracting the total incoming events from the licensed EPS. In the example shown on this slide, the total incoming EPS from the three collectors is 175. Over a three-minute duration, that is 31,500 events. The total unused events is 71,460, which is added to the EPS reservoir.

DO NOT REPRINT
© FORTINET



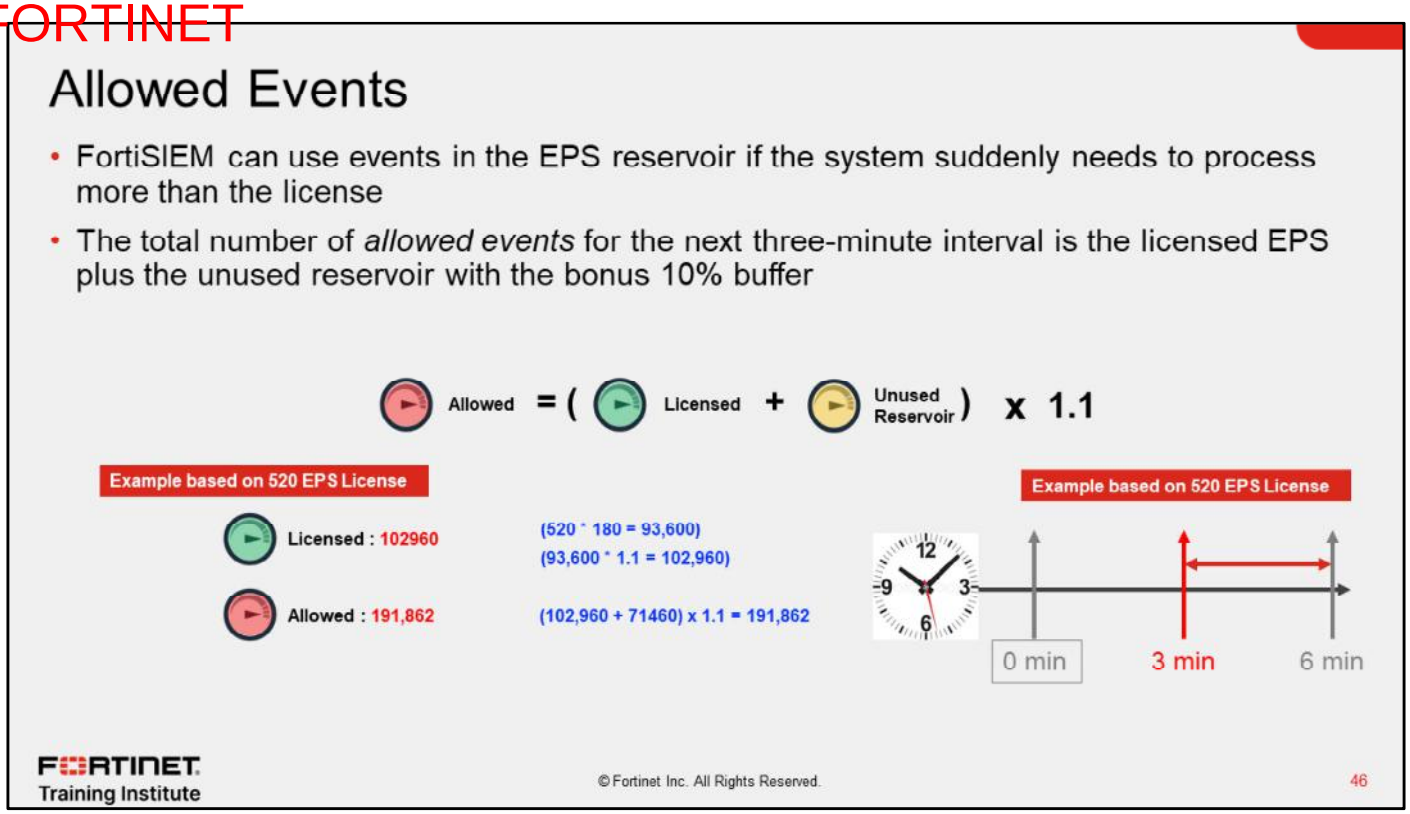
FortiSIEM creates a reservoir of events to use during event bursts. The supervisor receives the incoming EPS from all the nodes every three minutes, and that incoming EPS is the used EPS available to every node. The unused EPS is calculated by subtracting the used EPS from the licensed EPS, which is a global value that can be shared by any node in case there is a burst.

DO NOT REPRINT
© FORTINET



At midnight, the number of events than can be carried over to the next day is restricted to 50% and the process of building the EPS reservoir starts over for the next day.

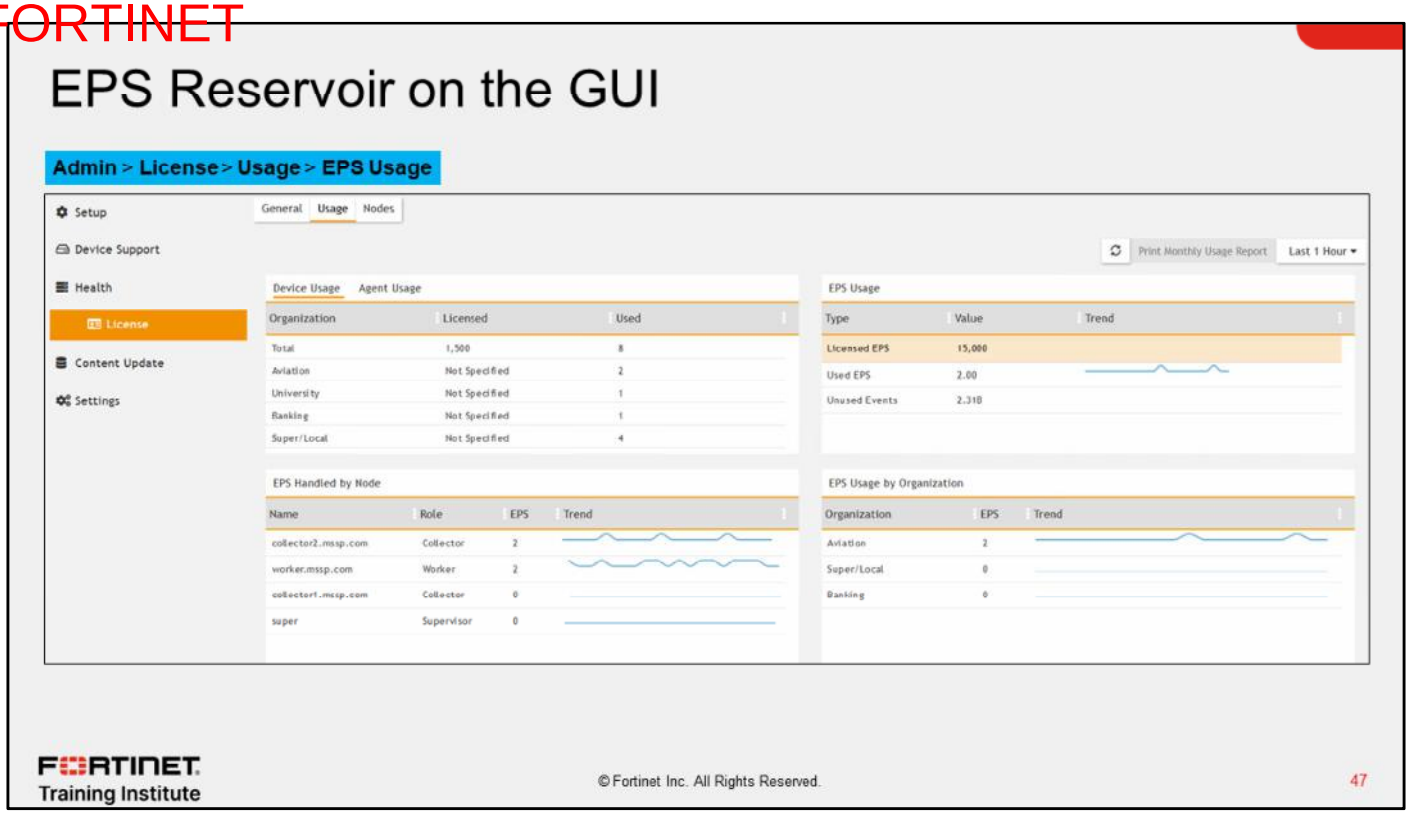
DO NOT REPRINT
© FORTINET



FortiSIEM can use events in the EPS reservoir if the system suddenly needs to process more than the license.

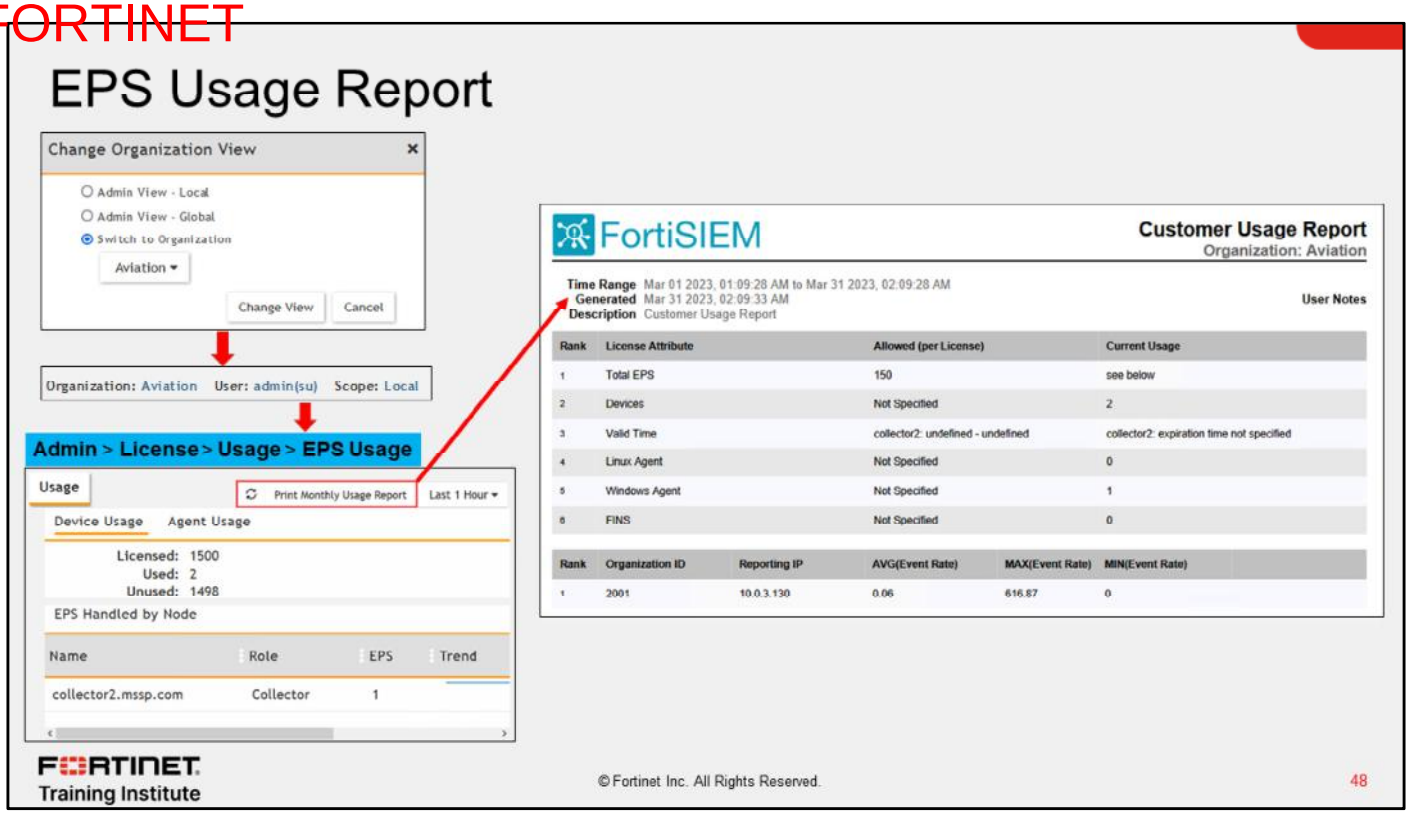
The supervisor then computes the allocated EPS for the next three-minute interval for every node, and communicates those values to every node. The total number of allowed events for the next three-minute interval is the licensed EPS plus the unused reservoir value, with the bonus 10% buffer. Each node reports the incoming EPS for the current interval to the supervisor node.

DO NOT REPRINT
© FORTINET



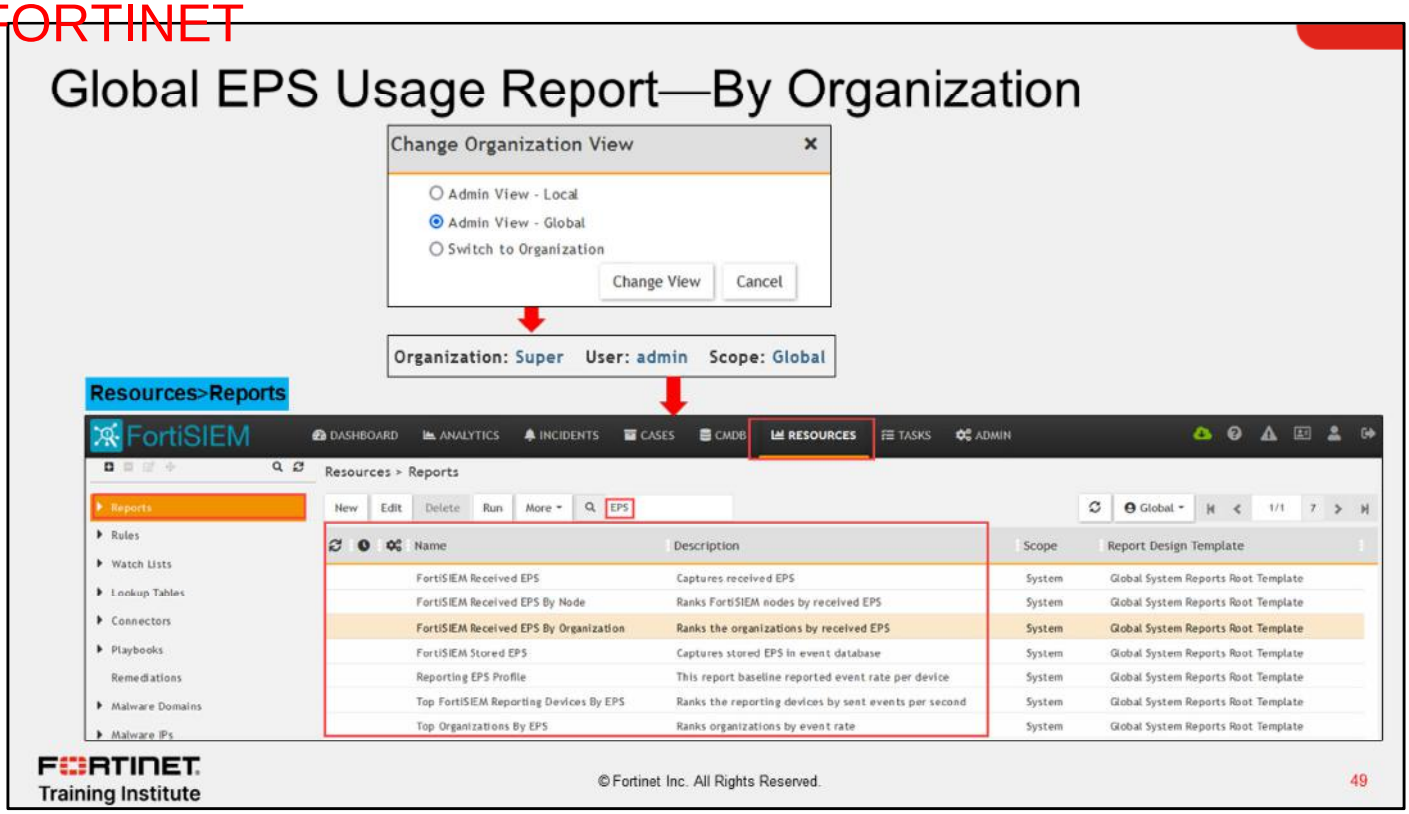
You can check the **Licensed EPS**, **Used EPS**, and **Unused Events** on the **EPS Usage** page on the FortiSIEM GUI.

DO NOT REPRINT
© FORTINET



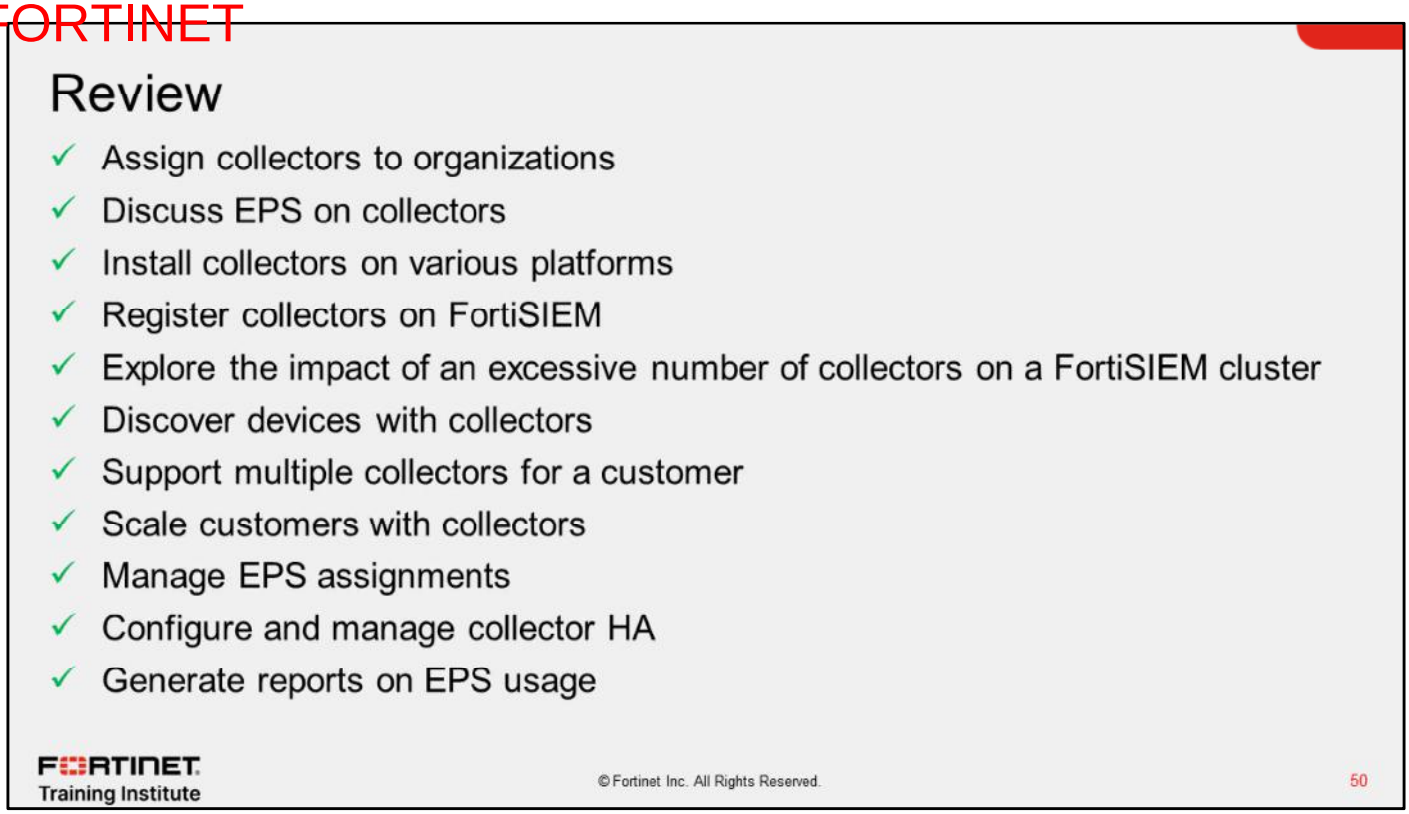
You can create a custom EPS usage report for each organization.

DO NOT REPRINT
© FORTINET



You can create a global EPS usage report for all organizations by using predefined system EPS reports.

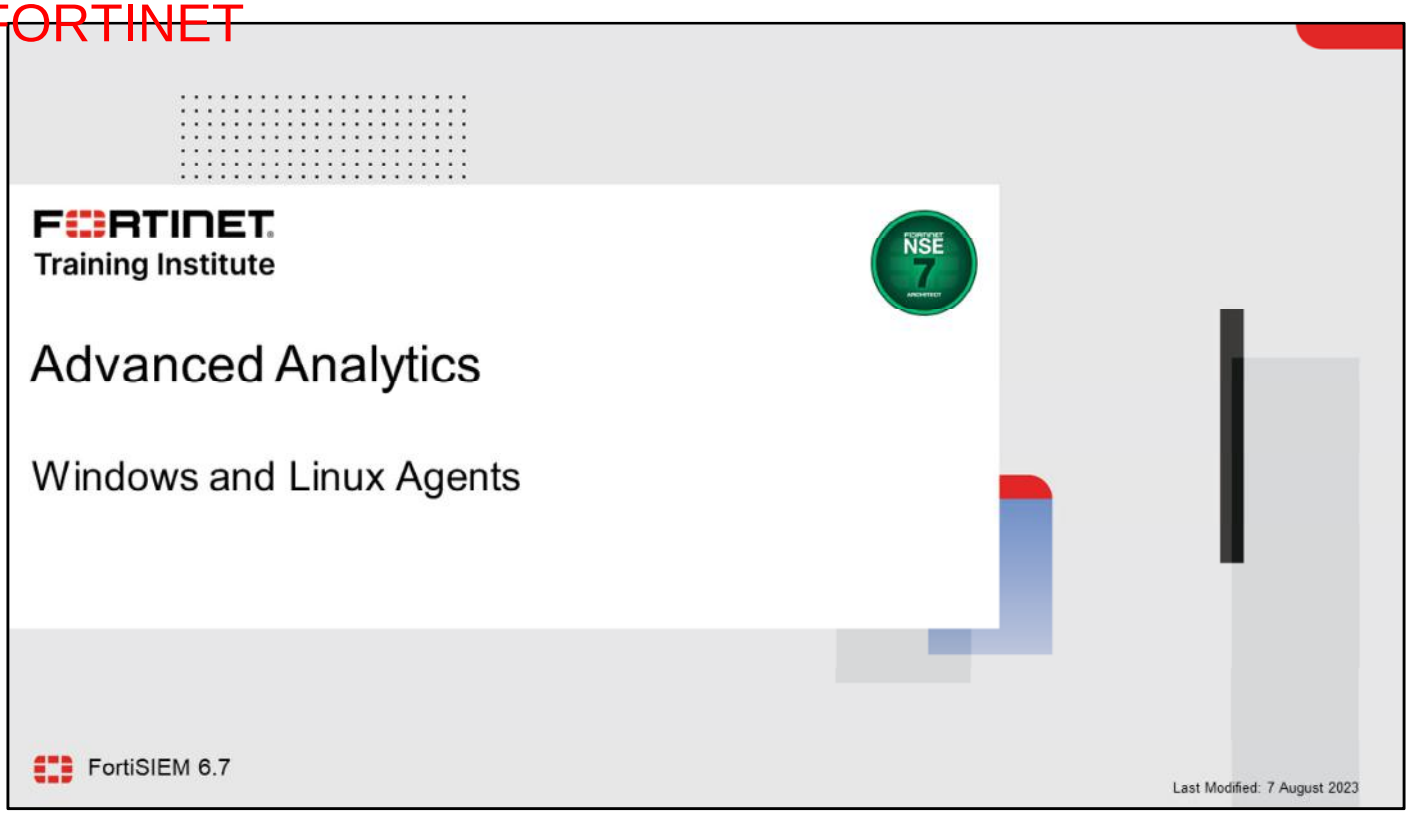
DO NOT REPRINT
© FORTINET



This slide shows the objectives that you covered in this lesson.

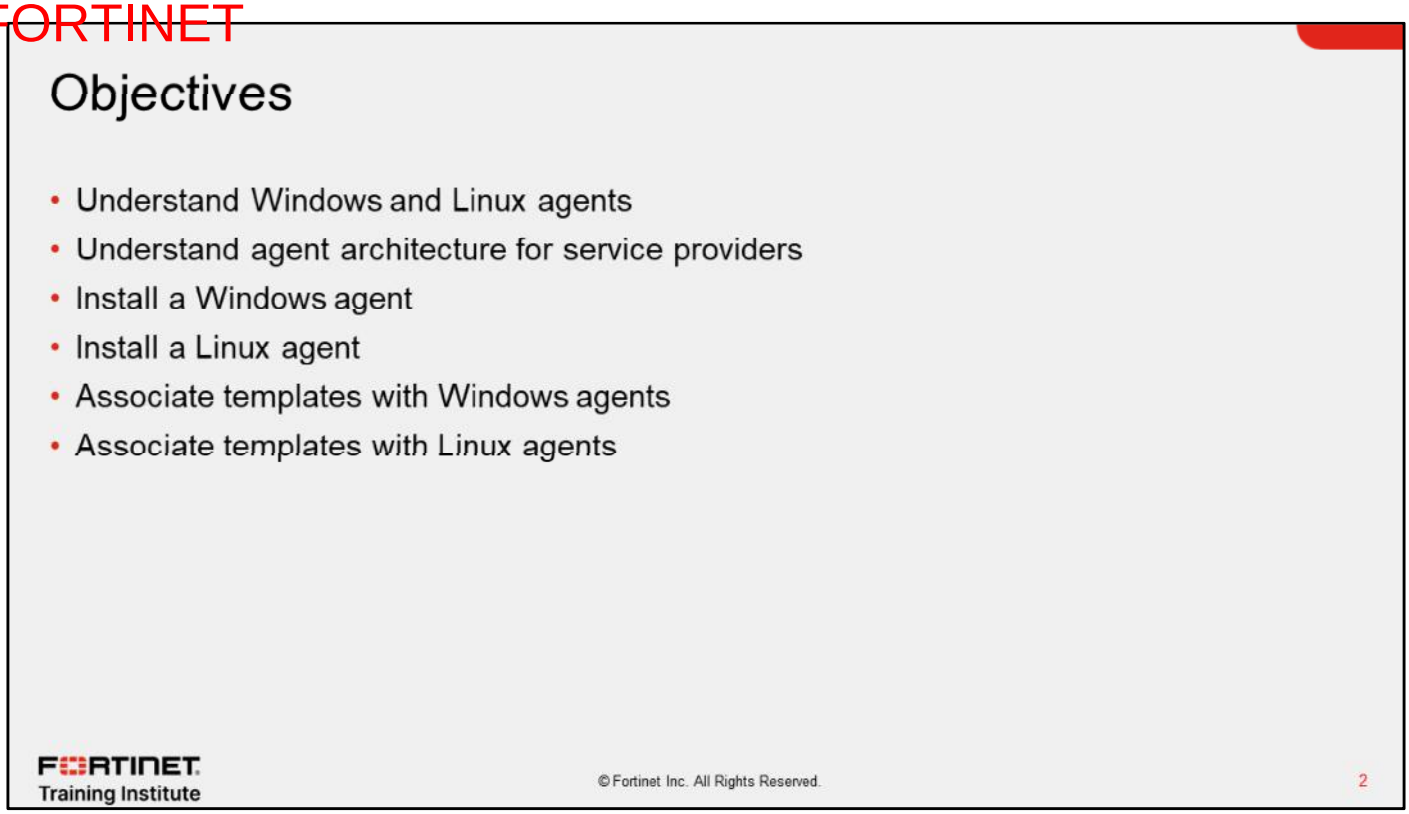
By mastering the objectives covered in this lesson, you learned about installing collectors on various platforms, registering collectors on FortiSIEM, and discovering devices with collectors. You also learned about managing events per second (EPS) assignment for collectors.

DO NOT REPRINT
© FORTINET



In this lesson, you will learn about installing and managing FortiSIEM Windows and Linux agents.

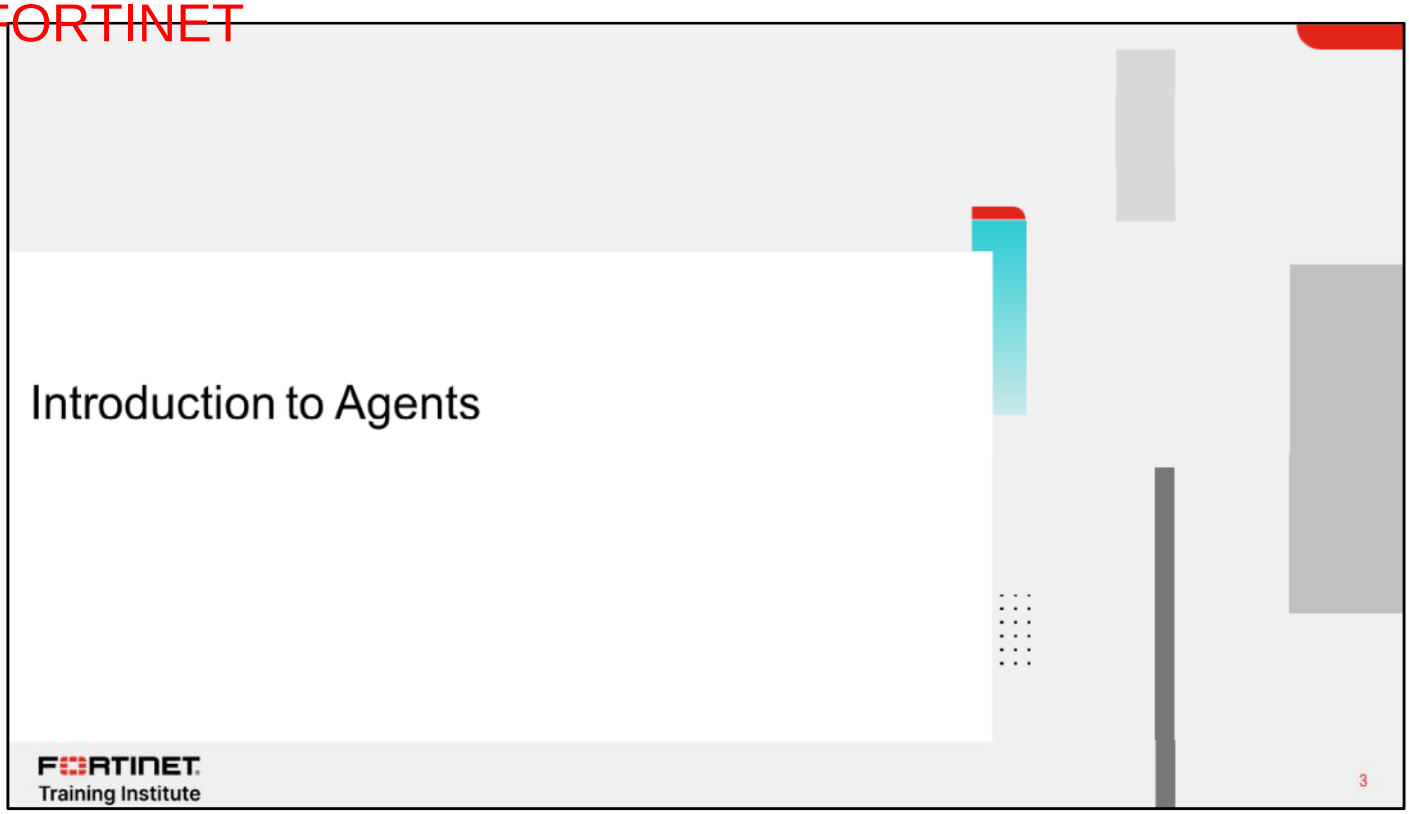
DO NOT REPRINT
© FORTINET



After completing this lesson, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in Windows and Linux agents, you will be able to install and manage agents in a multi-tenant environment.

DO NOT REPRINT
© FORTINET

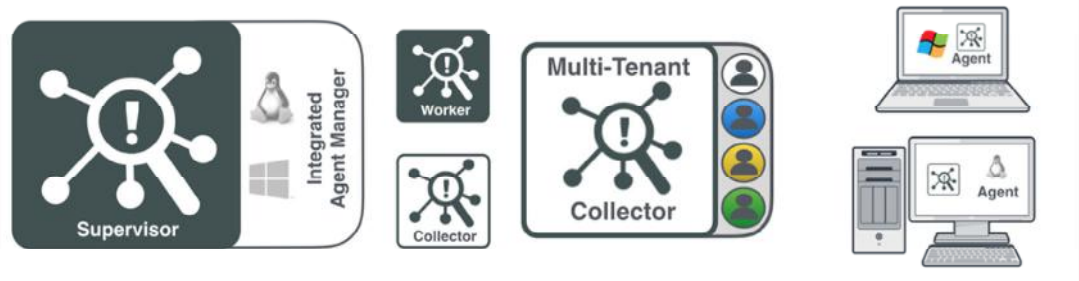


In this section, you will learn about various FortiSIEM agent components and their features.

DO NOT REPRINT
© FORTINET

FortiSIEM Agent Components

- Supervisor node has an integrated agent manager component that manages FortiSIEM Windows and Linux agents



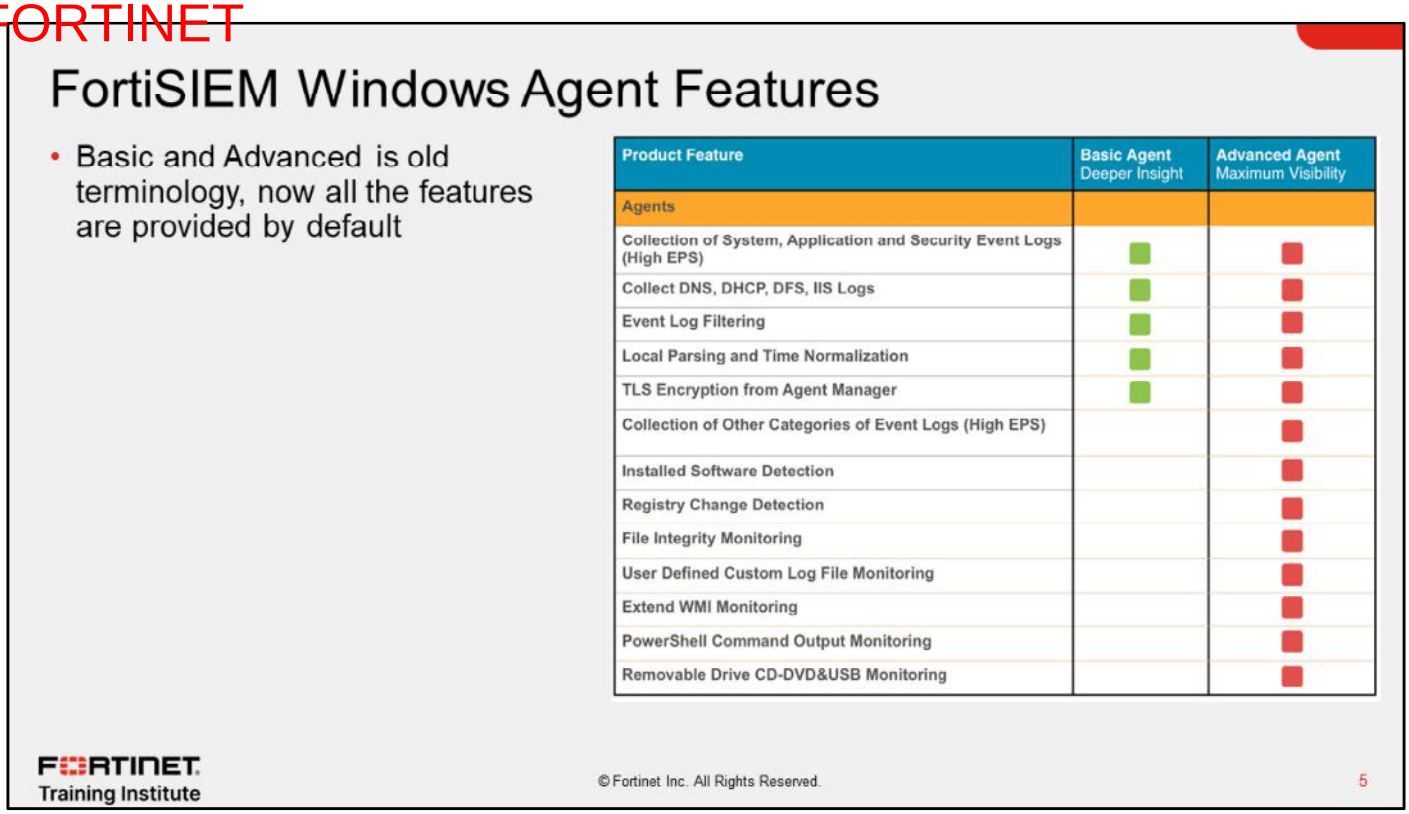
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

4

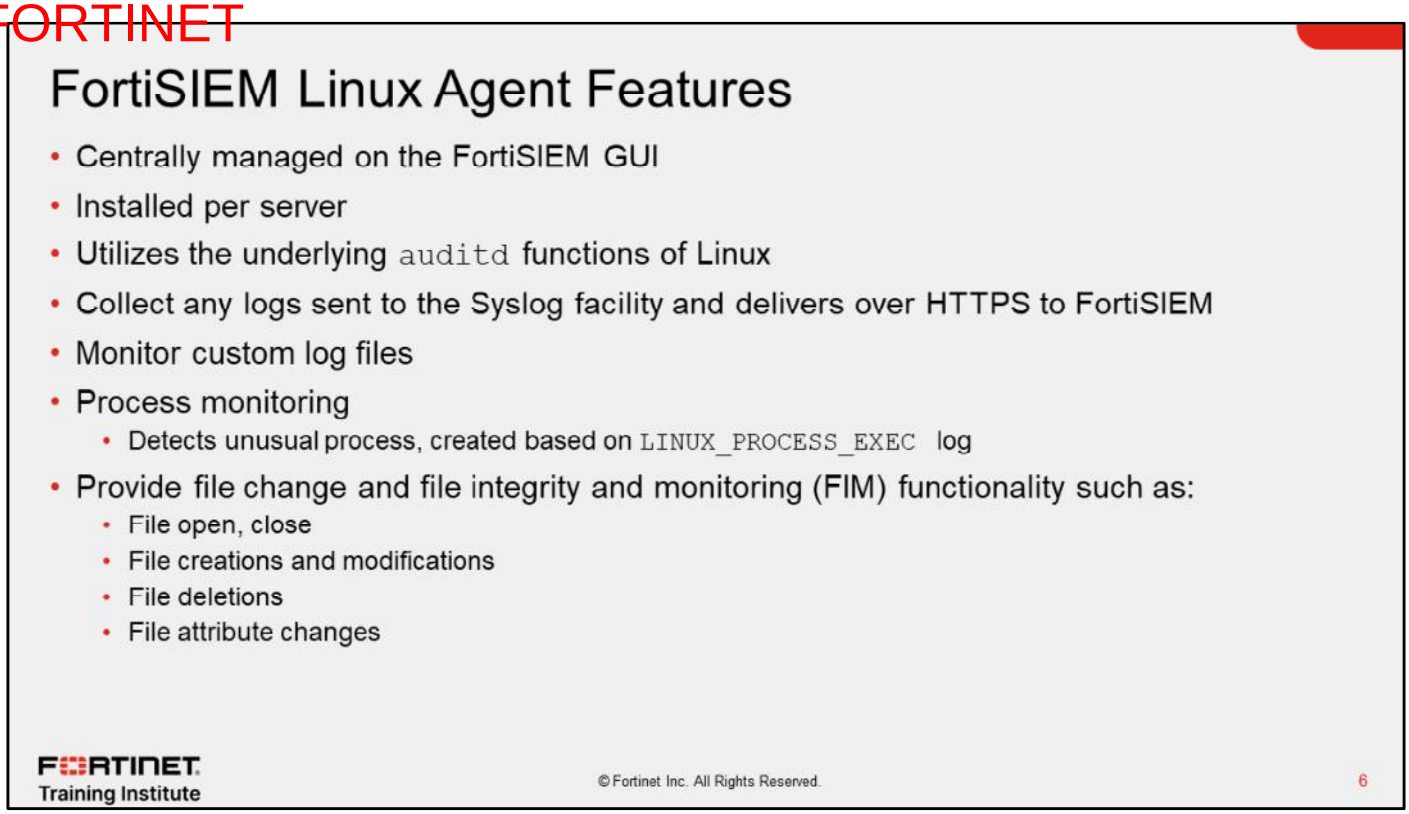
In all types of deployments, there is one supervisor node and optional worker and collector nodes, depending on the type of architecture you've chosen to deploy. Beginning from FortiSIEM 5.2, the supervisor node has an integrated agent manager component that manages FortiSIEM Windows and Linux agents.

DO NOT REPRINT
© FORTINET



Basic and Advanced is old terminology when referring to Windows agents. By default, the new Windows agent supports all features. Features, such as installed software detection, registry change detection, file integrity monitoring, extended WMI monitoring, PowerShell command output monitoring, and removable drive monitoring are supported by the FortiSIEM Windows agent.

DO NOT REPRINT
© FORTINET

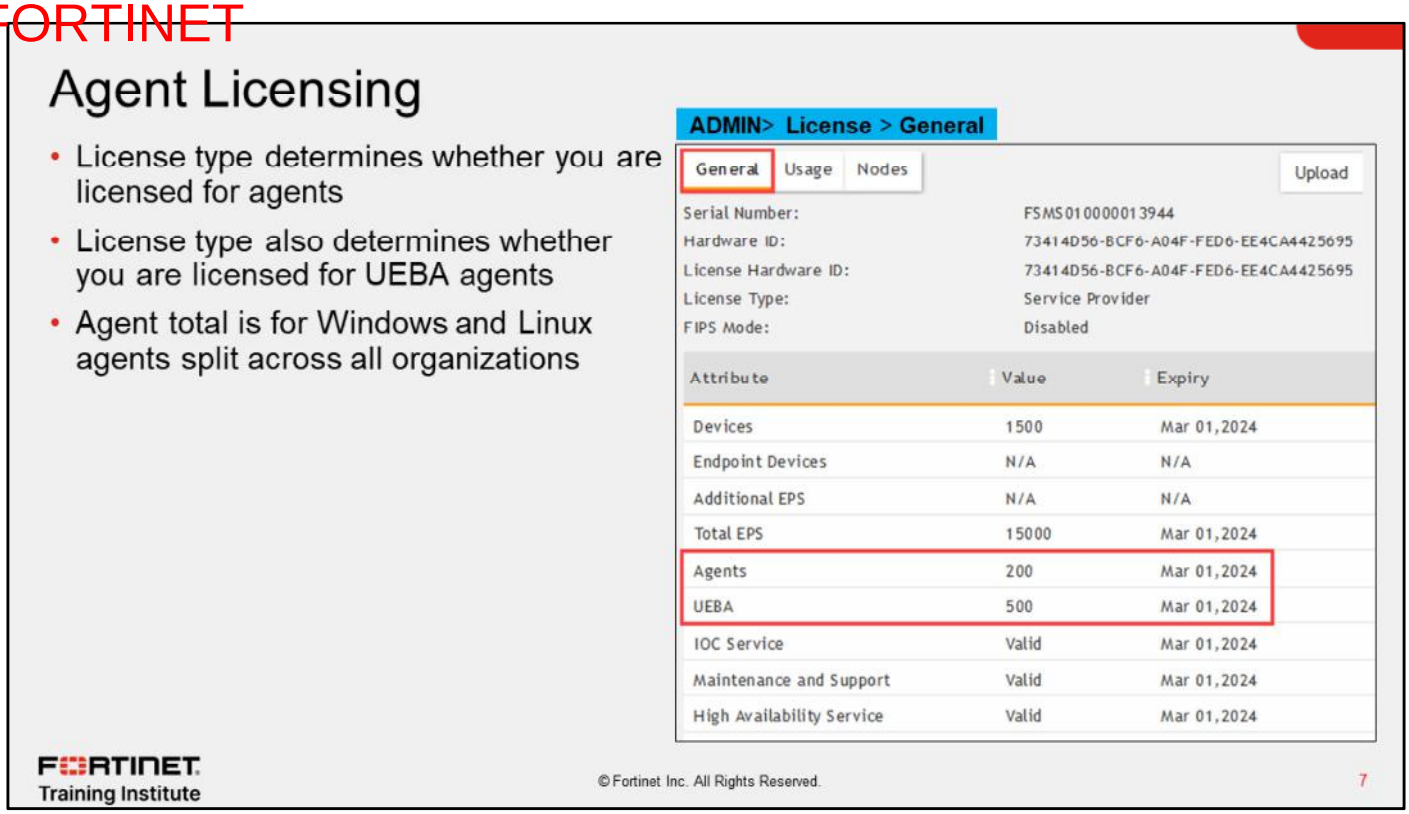


Linux agents can be used for FIM to detect file reads, writes, and edits with added user context. The agents are installed on the servers individually, and centrally managed on the FortiSIEM GUI.

The agents use the `auditd` daemon on Linux. It is a component of the Linux auditing system that is responsible for writing audit records to the disk. The audit daemon itself has some configuration options that you can customize in the `auditd.conf` file.

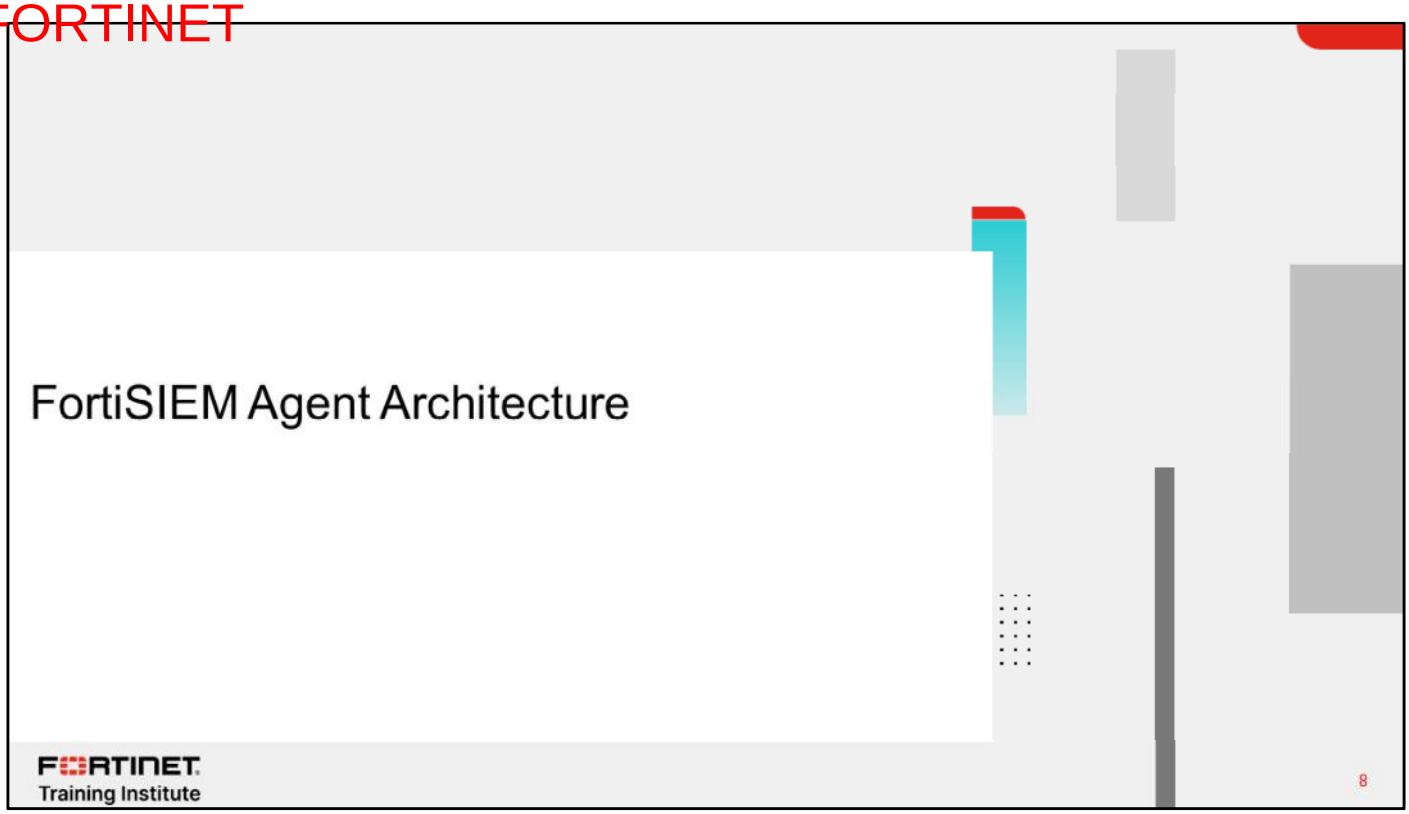
Logs collected by the Linux agent are sent to a Syslog facility, such as a FortiSIEM collector, and the collector delivers the logs over HTTPS to FortiSIEM.

DO NOT REPRINT
© FORTINET



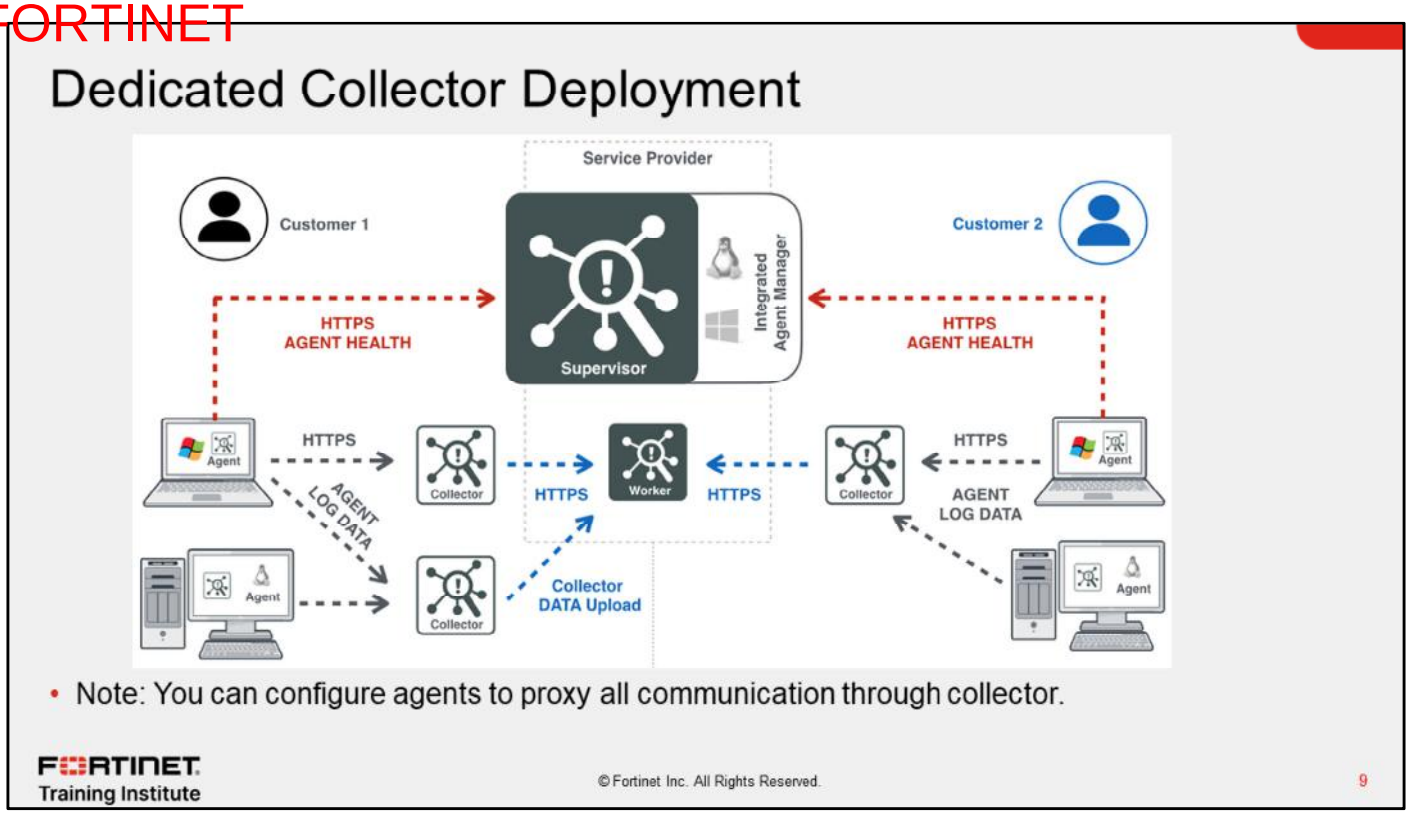
You need to have a license type that supports agents. You can check how many agents are supported in your license on the FortiSIEM GUI. The number of agents supported is a global value and will be shared across organizations. If you have more than 200 agents across all organizations that are managed by the supervisor, then you need to upgrade the license to support additional agents.

DO NOT REPRINT
© FORTINET



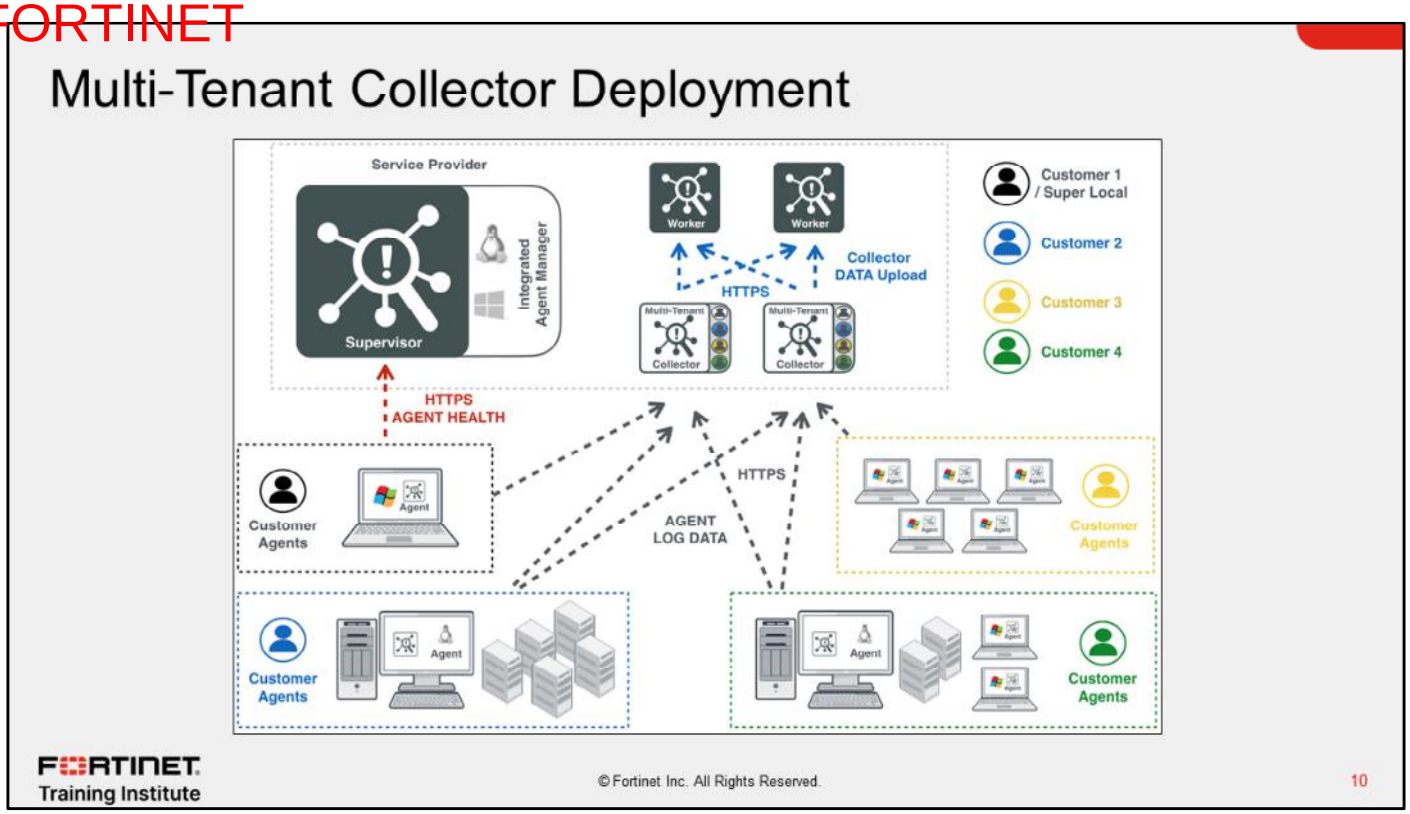
In this section, you will learn about FortiSIEM agent architecture for service providers.

DO NOT REPRINT
© FORTINET



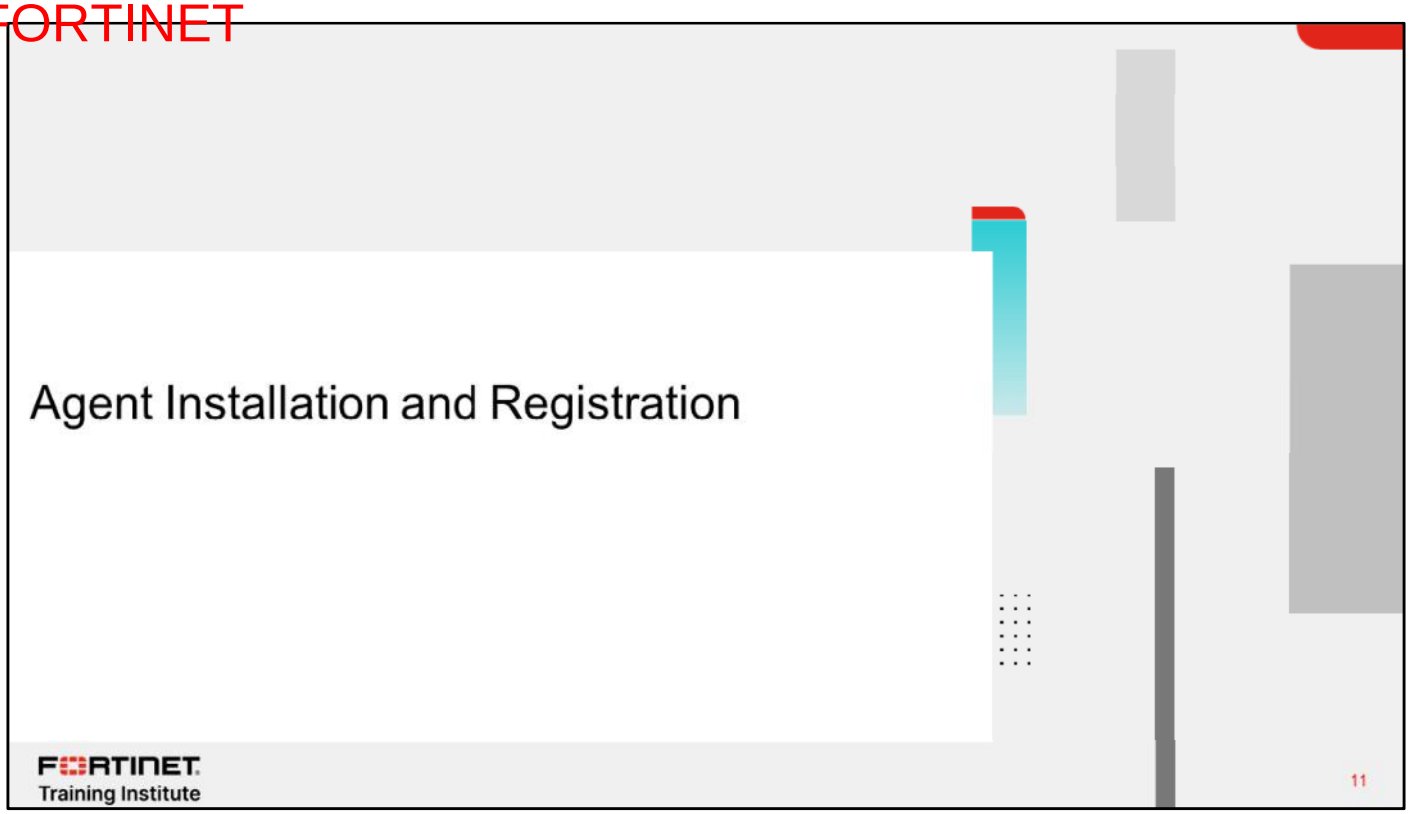
Agents use HTTPS to communicate directly with one or more dedicated customer collectors for log data, while also reporting health information to the supervisor node only. Agents can also proxy all communication through collectors. You need to configure the collector to send traffic for example, all traffic to agent then to collector and then to supervisor or worker.

DO NOT REPRINT
© FORTINET



Agents use HTTPS to communicate directly with one or more multi-tenant collectors to send log data, while also reporting health information to the supervisor node only. Multi-tenant collector is a special use case and only available where the organization does not have its own defined collector, essentially makes the collector almost act as a worker for log collection, still have the same restrictions, no overlapping IP and so on. The agent traffic tags the organization in the log message.

DO NOT REPRINT
© FORTINET



In this section, you will learn about FortiSIEM agent installation on Windows and Linux platform. You will also learn about the agent registration process with the supervisor node.

DO NOT REPRINT
© FORTINET

Supported OSs for Windows and Linux Agents

- CentOS 7.4 and later
- Red Hat Enterprise Linux 7.x
- Ubuntu 14.04, 16.04, 18.04, 20.04 LTS
- Amazon Linux 1 and Amazon Linux 2
- SuSE Enterprise Linux (SLES) 12 and 15
- For CentOS and Red Hat, the version requirements are:
  - curl version later than 7.19.7
  - nss.x86_64 version later than 3.36.0

- Windows 7 Enterprise/Professional
- Windows 8
- Windows 10
- Windows Server 2008 R2
- Windows Server 2012
- Windows Server 2012 R2
- Windows Server 2016
- Windows Server 2019
- Windows Server 2019 Core
- Windows Server 2022





FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

12

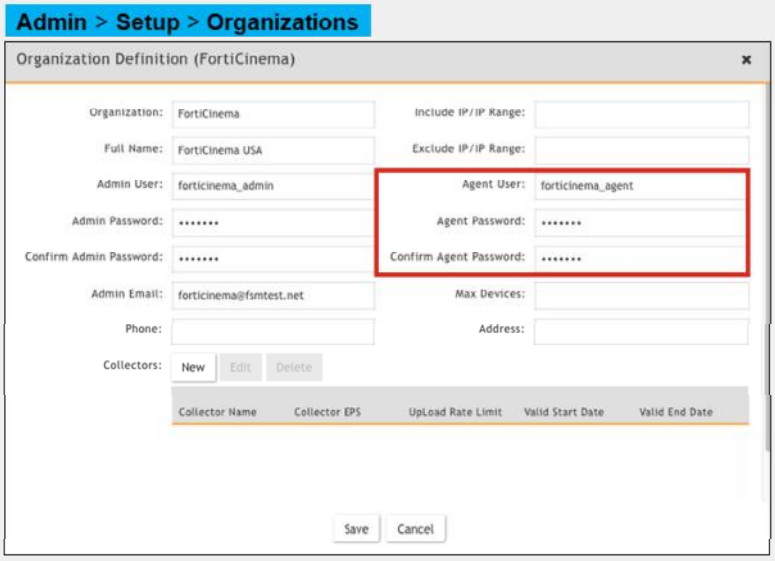
Windows agent is supported on Windows 7 Enterprise and Professional, Windows 8, Windows 10, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, Windows Server 2016, Windows Server 2019, Windows Server 2019 Core, and Windows Server 2022.

Linux agent is supported on CentOS 6.x and 7.x, Redhat Enterprise Linux 6.x and 7.x, Ubuntu 14, 16, 18, 20.04 LTS, Amazon Linux 1, and Amazon Linux 2.

DO NOT REPRINT
© FORTINET

Windows and Linux Agent User

- For service provider mode, on the organization definition page, define a special agent user and agent password
- This is used for agent registration with the supervisor FortiSIEM



- Note: For enterprise mode define agent user in **CMDB** under **Users**

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

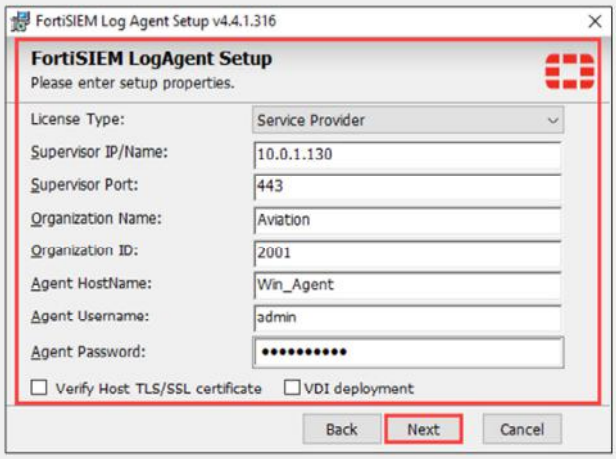
13

Before you deploy an agent for an organization, you must define the agent credentials on the supervisor node. The **Agent User** and **Agent Password** that you define for an organization will be used to validate agents registering from that organization to the supervisor node.

DO NOT REPRINT
© FORTINET

Windows Agent Installation Procedure

- During installation, the Windows agent registers with the FortiSIEM supervisor
- The required parameters are:
  - **License Type**: service provider or enterprise
  - **Supervisor IP/Name**: IP address or host name/FQDN of supervisor node
  - **Supervisor Port**: 443
  - **Organization Name**: FortiSIEM organization name
  - **Organization ID**: FortiSIEM Organization ID to which this agent belongs
  - **Agent Username** : agent username (for registration only)
  - **Agent Password**: agent password (for registration only)
  - **Agent HostName**: This name is displayed in FortiSIEM CMDB. FortiSIEM recommends using an FQDN, especially if SNMP or WMI is also going to be used for this device.
  - Windows agent also supports Virtual Desktop Infrastructure (VDI) as a deployment mechanism



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

14

To install the Windows agent, you must first download the file from the Fortinet Support website—the agent installer. You must edit the parameters for agent registration in the installation agent setup GUI. This includes the organization ID number and name, supervisor IP address and port, and registration username and password. Finally, you can run the agent installer.

Windows agent also supports Virtual Desktop Infrastructure (VDI) as a deployment mechanism. VDI deployment also supports read-only VDI images. In the VDI deployment scenario, device names are added to the CMDB device list as the active session user, separating domain and username with two underscores.

When the user logs in to the VDI environment and downloads a VM from the VDI Server, the VM contains a VDI transient image (containing the Windows Agent). The agent automatically registers on the FortiSIEM supervisor node.

When the user logs out of the VDI environment, the agent automatically unregisters from the FortiSIEM supervisor node. The agent status is set to decommissioned so that it does not consume an agent license.

DO NOT REPRINT
© FORTINET

Linux Agent Installation Procedure

- Download the shell script Linux agent installer from Fortinet Support website
- Install script needs execute permissions and you must install it as a root user
- You must install the following packages on the Linux server before installation:
  - libcap2-bin
  - auditd
  - setools
  - Rsyslog
  - logrotate

```
OPTIONS:
-h  Show this message
-s  Super IP/HostName
-i  Organization Id
-p  Agent Registration Password
-u  Agent Registration User
-o  Organization
-P  Proxy IP <Optional>
-O  Proxy Port <Optional>
```

```
./fortisiem-linux-agent-installer-6.7.4.1733.sh -s <Supervisor-IP> -i
<Organization-Id> -o <Organization-Name> -u <Agent-User> -p <Agent-Password>
```

FORTINET Training Institute

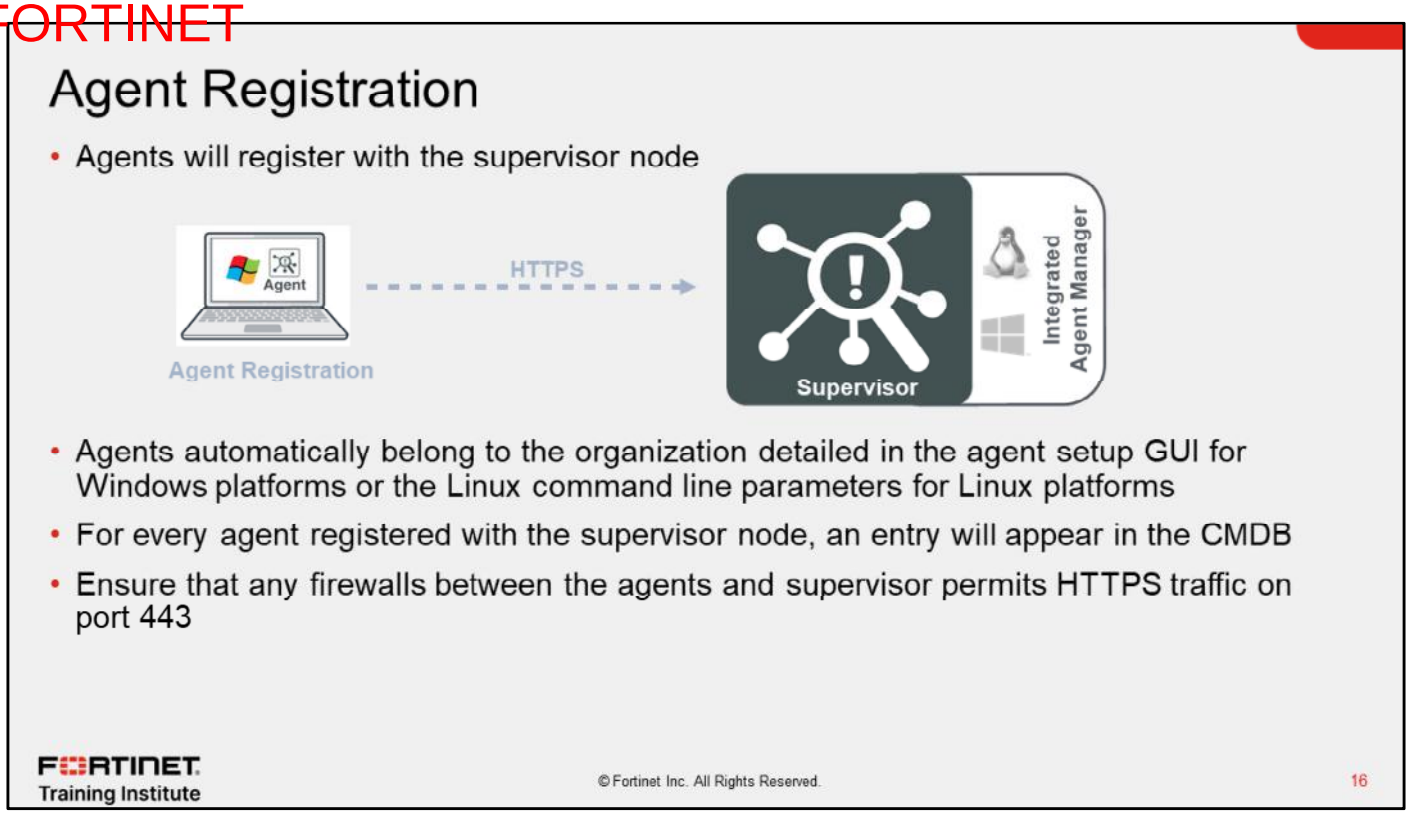
© Fortinet Inc. All Rights Reserved.

15

To install a Linux agent, you have to download the shell script for the Linux agent installer from the Fortinet Support site. There are some dependencies you must be aware of. Make sure you install the packages listed on the slide before running the installer.

Once you have the dependent packages installed, run the shell script using the command shown on this slide. The install script needs execute permissions and you must install it as a root user. Fill in the parameters such as supervisor IP address, organization ID, organization name, agent username, and agent password.

DO NOT REPRINT
© FORTINET



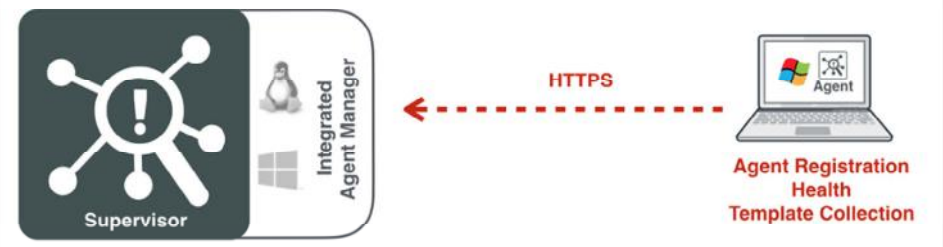
After the installation is complete, the agents will register with the supervisor node. The agent belongs to the organization detailed in the agent setup GUI for Windows platforms, or the command line parameters for Linux platforms. An entry appears in the CMDB for every agent that is installed and registered with the supervisor node.

The registration between the agent and the supervisor node occurs over port 443. Make sure that any firewall between the agent and the supervisor node permits HTTPS traffic on port 443.

DO NOT REPRINT
© FORTINET

Agent Communications After Installation

- Agents will maintain communication with the supervisor after initial installation



- This serves two purposes:
  - Report agent status and health
  - Poll and collect any new agent templates
- By default, each agent will poll the supervisor every minute

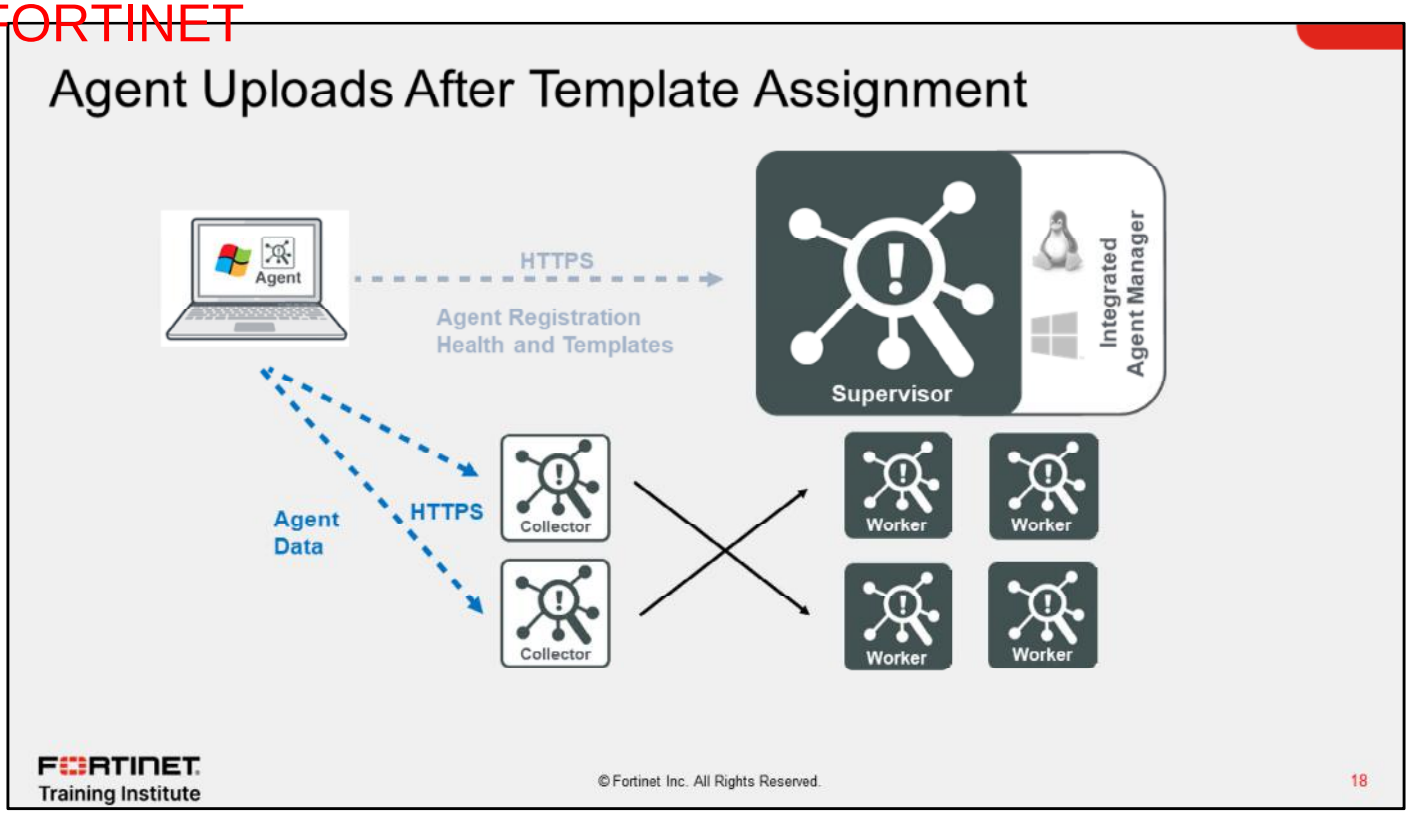
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

17

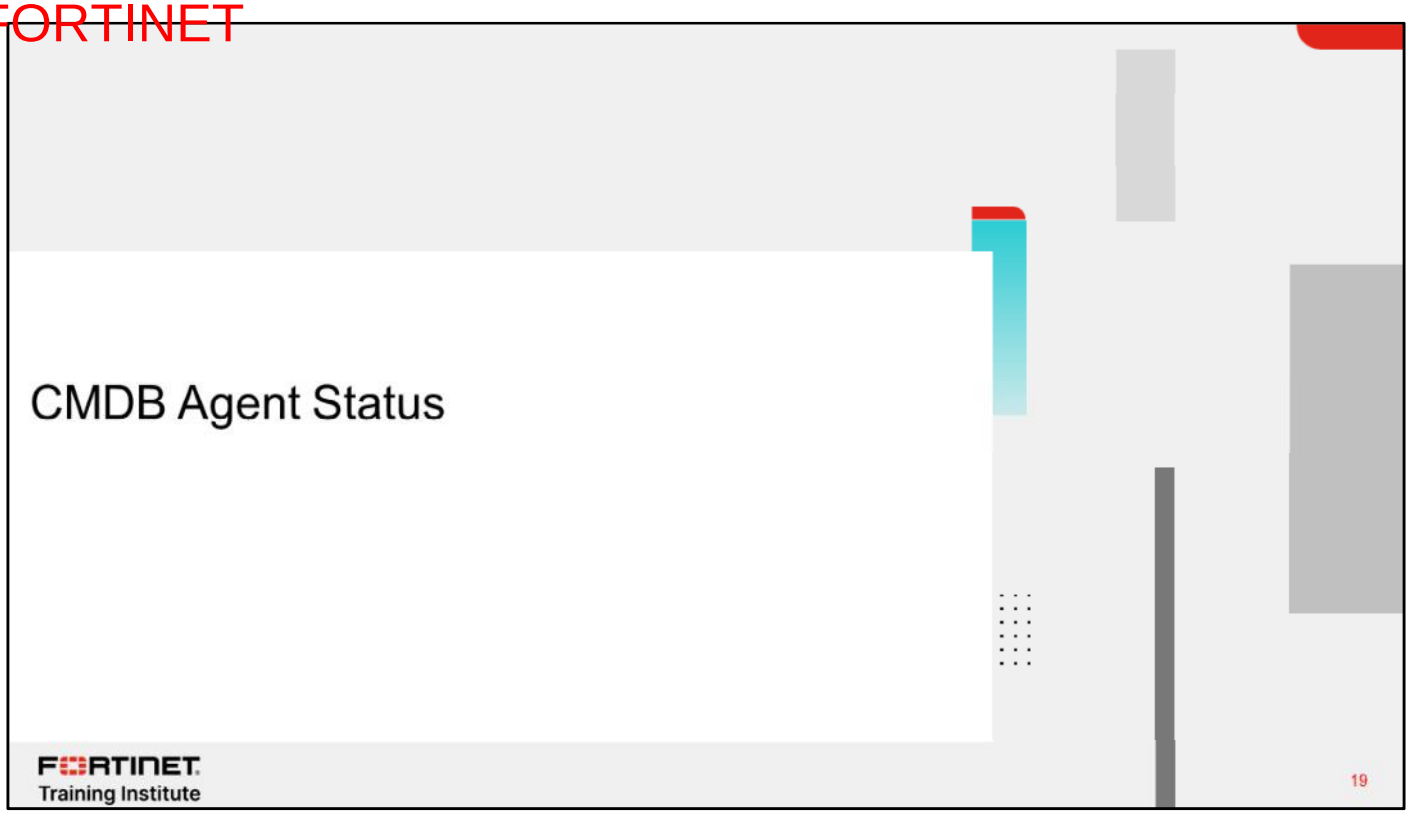
The agents will continue to communicate with the supervisor node after the initial installation. The reason for this communication is to report the agent status and to poll for any new agent templates. By default, every agent will poll the supervisor node every minute.

DO NOT REPRINT
© FORTINET



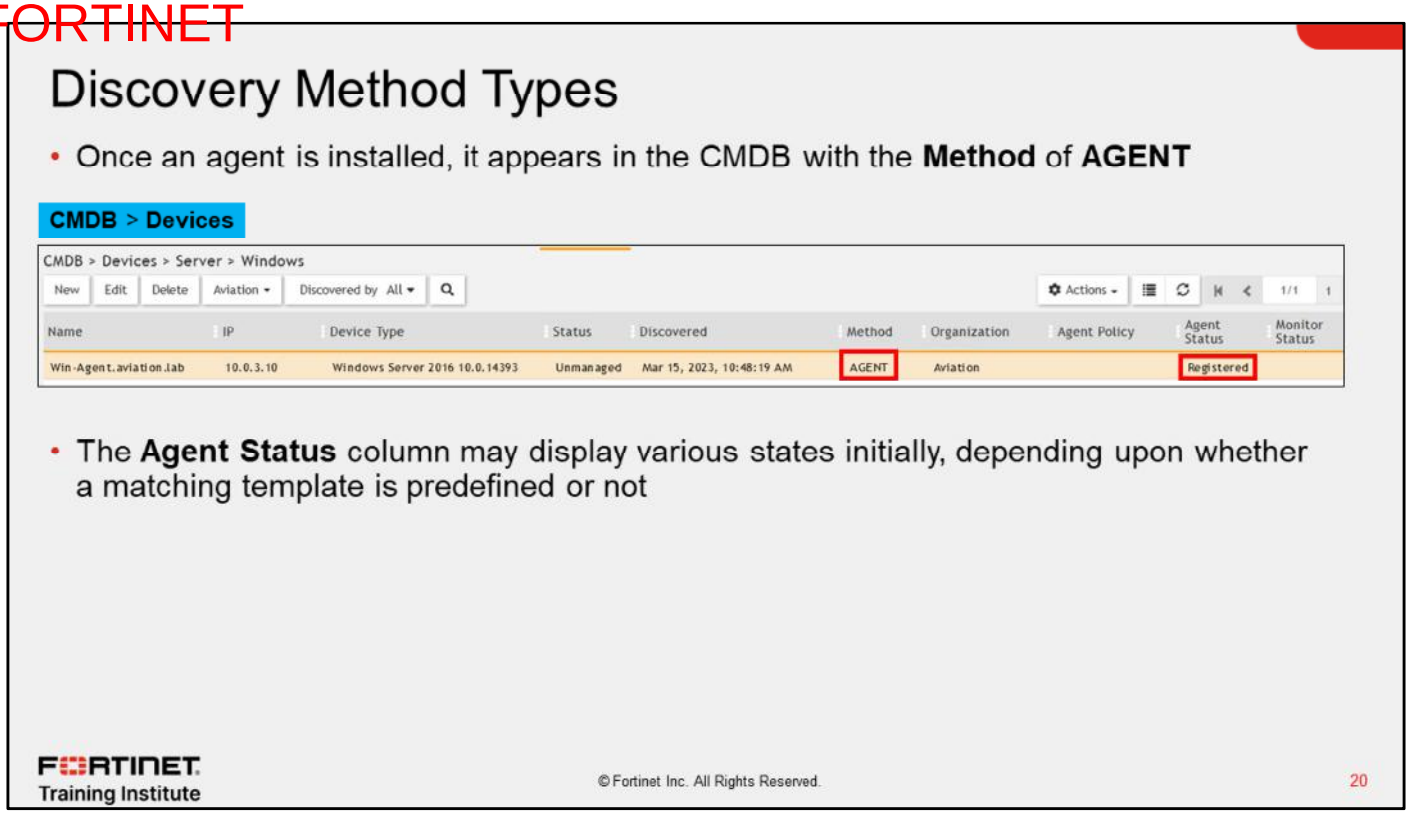
When the agent is registered, there is no template associated with the agent. The agent does not upload any log data until a template is associated with it. When you assign a template to an agent, it will begin uploading log data based on the parameters that are set in the template.

DO NOT REPRINT
© FORTINET



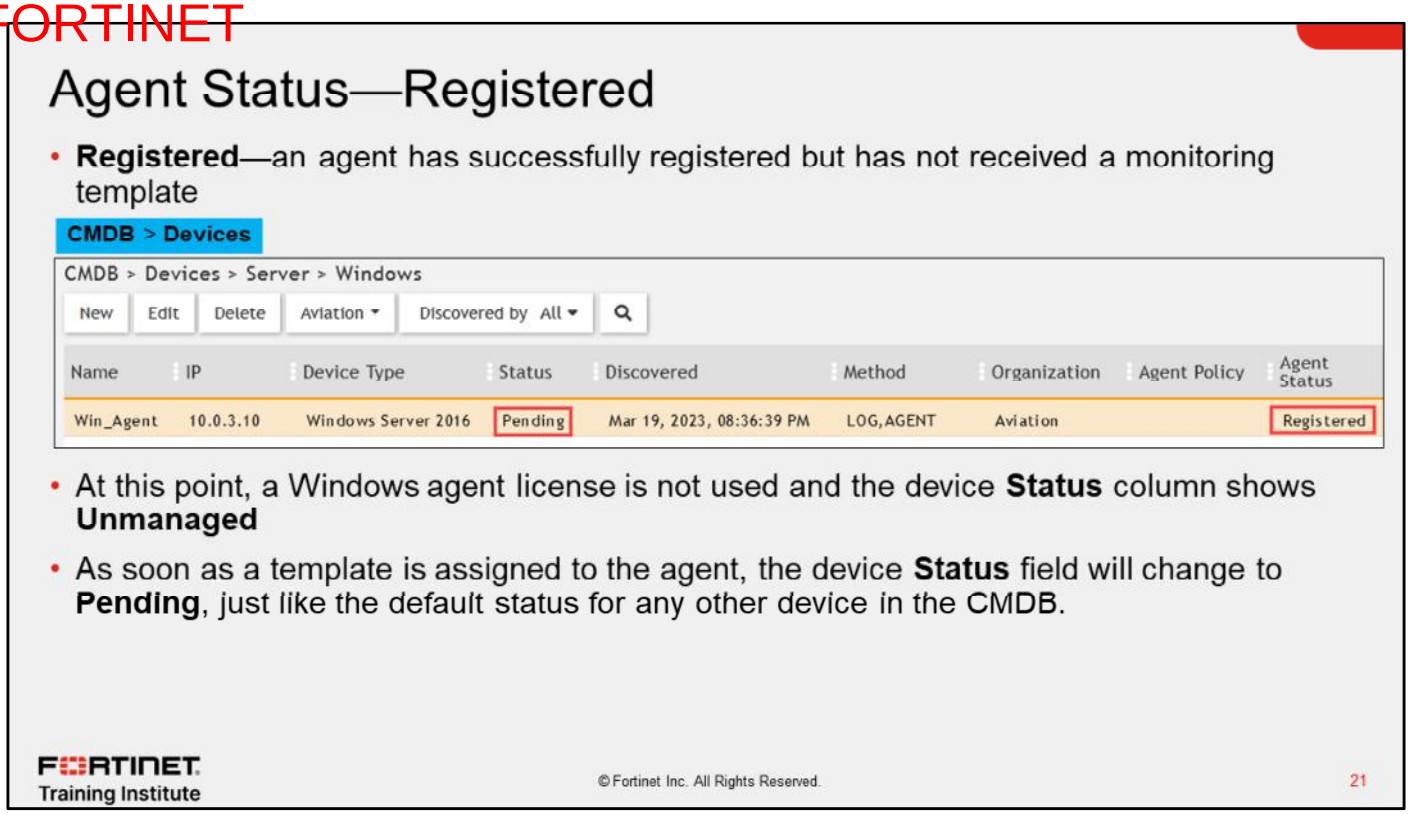
In this section, you will explore various CMDB agent statuses from the time the agent is registered until a template is associated with the agent.

DO NOT REPRINT
© FORTINET



Once an agent is installed, it appears in the CMDB with the **Method** of **AGENT**. The **Agent Status** column may display various states initially, depending upon whether a matching template is predefined or not.

DO NOT REPRINT
© FORTINET

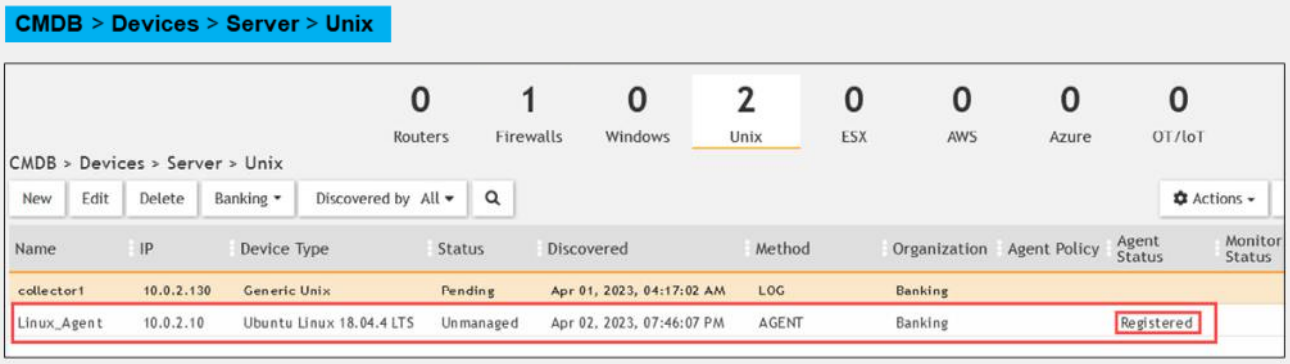


When an agent is registered successfully but has not received a monitoring template, the **Agent Status** column will show **Registered**. At this point, a Windows agent license is not used and the device **Status** column shows **Unmanaged**. As soon as a template is assigned to the agent, the device **Status** column will change to **Pending**, just like the default status for any other device in the CMDB.

DO NOT REPRINT
© FORTINET

Linux Agents Status

- The same agent status messages apply to both Windows and Linux agents



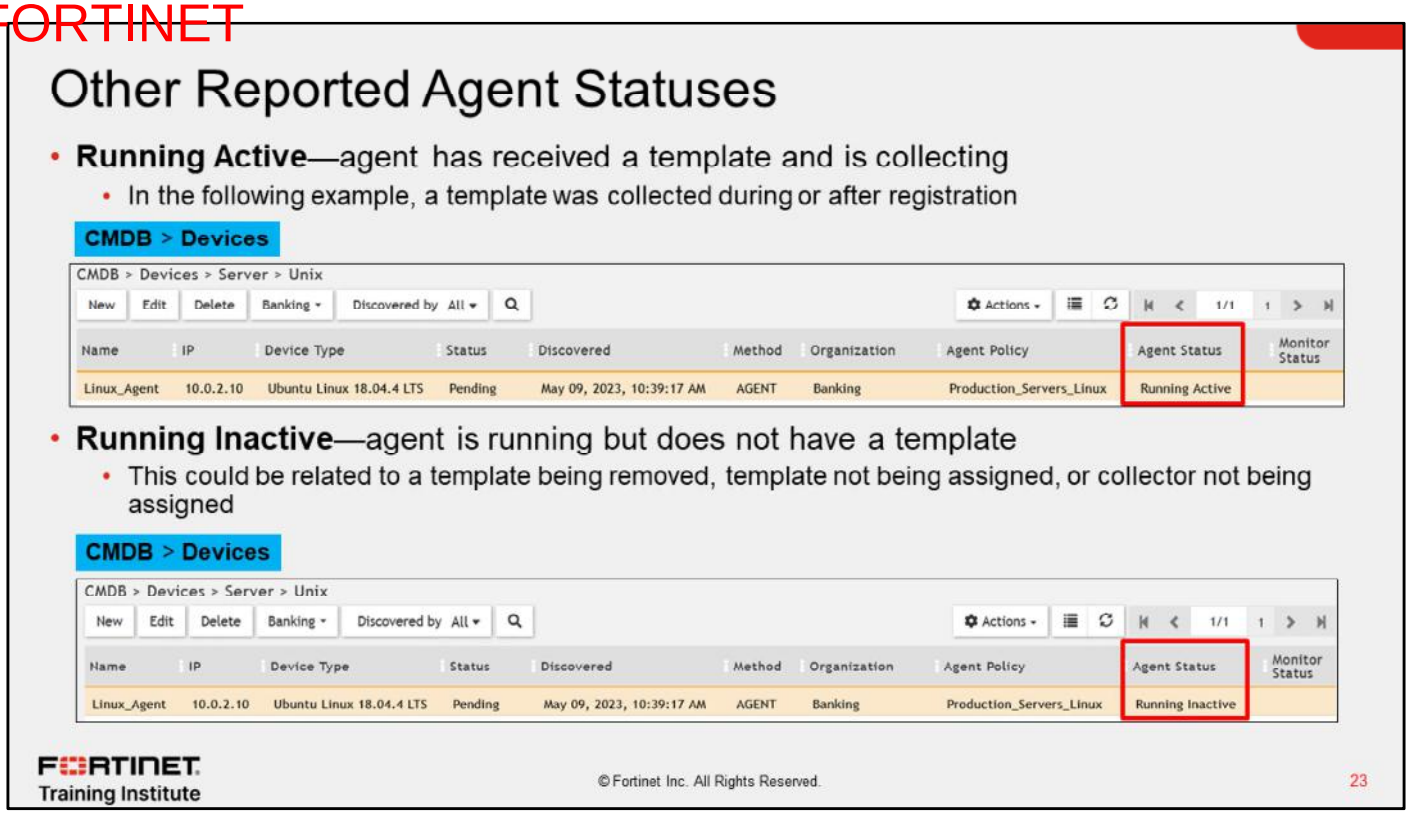
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

22

The same agent status messages apply to both Windows and Linux agents. This slide shows an example of a Linux agent that has just been installed.

DO NOT REPRINT
© FORTINET



When a template is associated with the agent, the **Agent Status** column will change to **Running Active**. At this point, the agent is actively collecting logs from the device it is installed on and sending them to the collector or supervisor node.

When the agent is running but does not have a template associated with it, the status will change to **Running Inactive**. This could be related to a template being removed, template not being assigned, or collector not being assigned.

DO NOT REPRINT
© FORTINET

Other Reported Agent Statuses (Contd)

- **Stopped**—agent service has been stopped on the device

CMDB > Devices

CMDB > Devices > Server > Unix

New | Edit | Delete | Banking | Discovered by All | Actions | 1/1

Name	IP	Device Type	Status	Discovered	Method	Organization	Agent Policy	Agent Status
Linux_Agent	10.0.2.10	Ubuntu Linux 18.04.4 LTS	Pending	May 09, 2023, 10:39:17 AM	AGENT	Banking	Production_Servers_Linux	Stopped

- **Disconnected**—the agent has not been communicating with the supervisor for more than 10 minutes

CMDB > Devices

Name	IP	Type	Status	Discovered	Method	Agent Policy	Agent Status	Monitor Status	Event Status
FortiBank_Collector	10.10.2.64	Generic Unix	Pending	Apr 28, 2023, 10:00:01 PM	LOG				Warning
fortibank_dc.fortibank.net	10.10.2.63	Windows	Pending	May 09, 2023, 10:39:17 AM	AGENT		Disconnected		

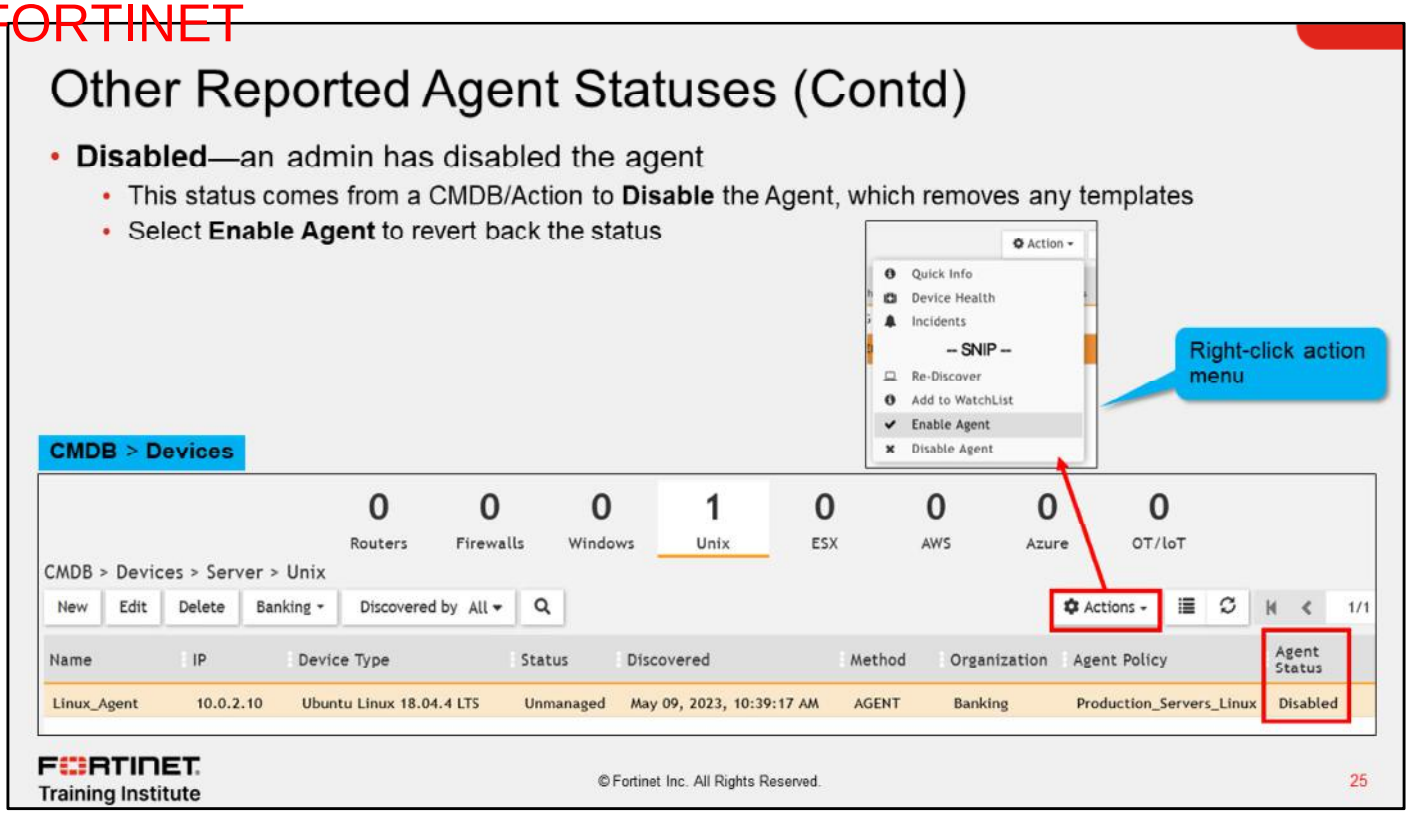
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

24

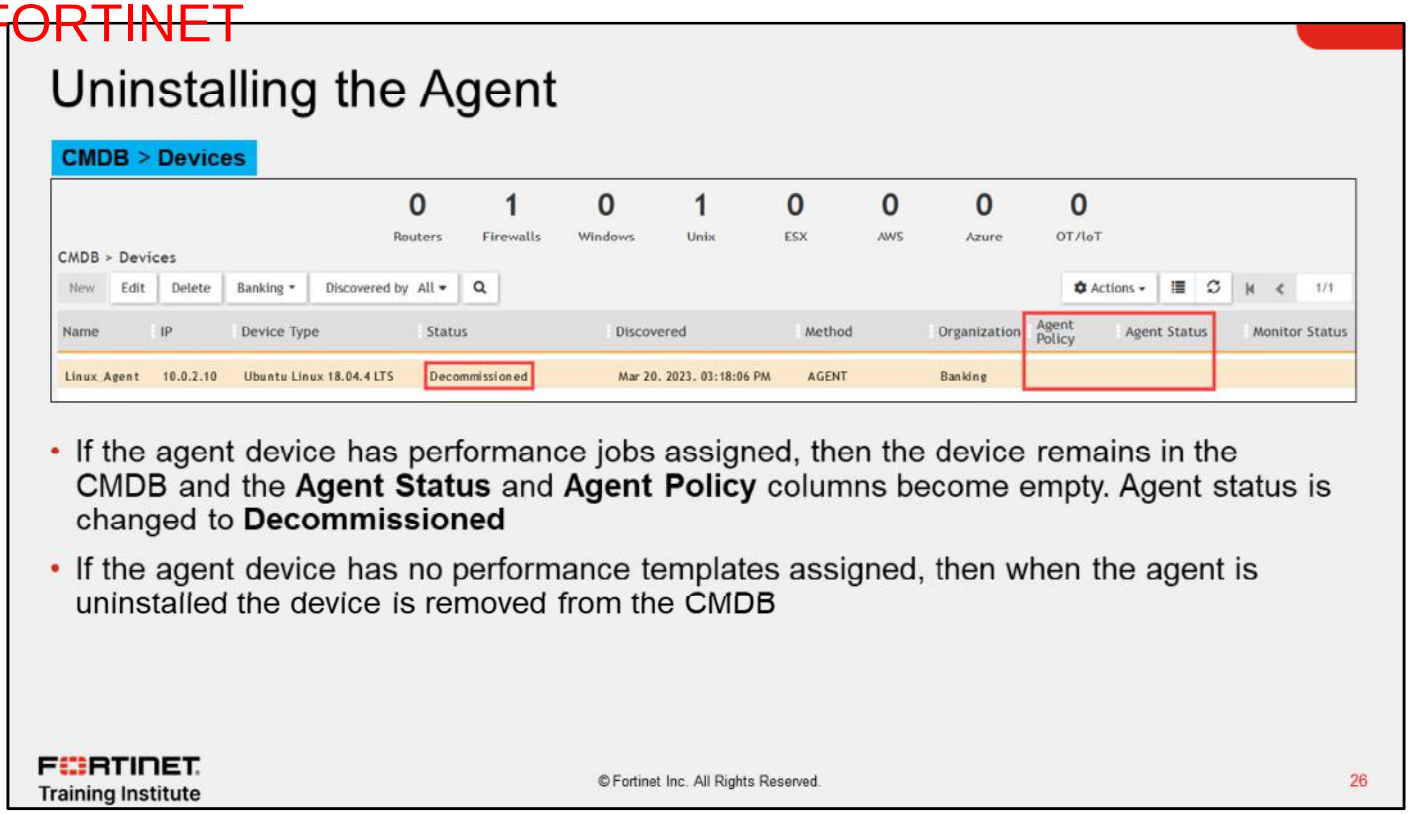
If the agent service is stopped on the device, then the **Agent Status** column will show **Stopped**. If, for some reason, the agent is not communicating with the supervisor node because of network connectivity issues, then the **Agent Status** column will show **Disconnected**. Any change in status may take up to one minute because of the polling interval.

DO NOT REPRINT
© FORTINET



The **Agent Status** column will show **Disabled** if an administrator has disabled the agent. This status comes from a CMDB action to **Disable Agent**, which removes any templates. Select **Enable Agent** to revert the status.

DO NOT REPRINT
© FORTINET



If the agent device has no performance templates assigned, then when the agent is uninstalled the device is removed from the CMDB. If the agent device has performance jobs assigned, then the device remains in the CMDB and the **Agent Status** and **Agent Policy** columns become empty, and agent status is changed to **Decommissioned**.

DO NOT REPRINT
© FORTINET

Summary of Agent Statuses

- Agent status reported in CMDB

Agent Status	Description
Registered	Agent has completed registration but has not received the monitoring template
Running Active	Agent has received a monitoring template and it is performing correctly
Running Inactive	Agent is running but does not have a monitoring template. This can be due to: • no license • incomplete definition—no collector or template is defined for that device
Stopped	Agent is stopped on the server
Disconnected	Supervisor has not received any status updates from the agent for the last 10 minutes

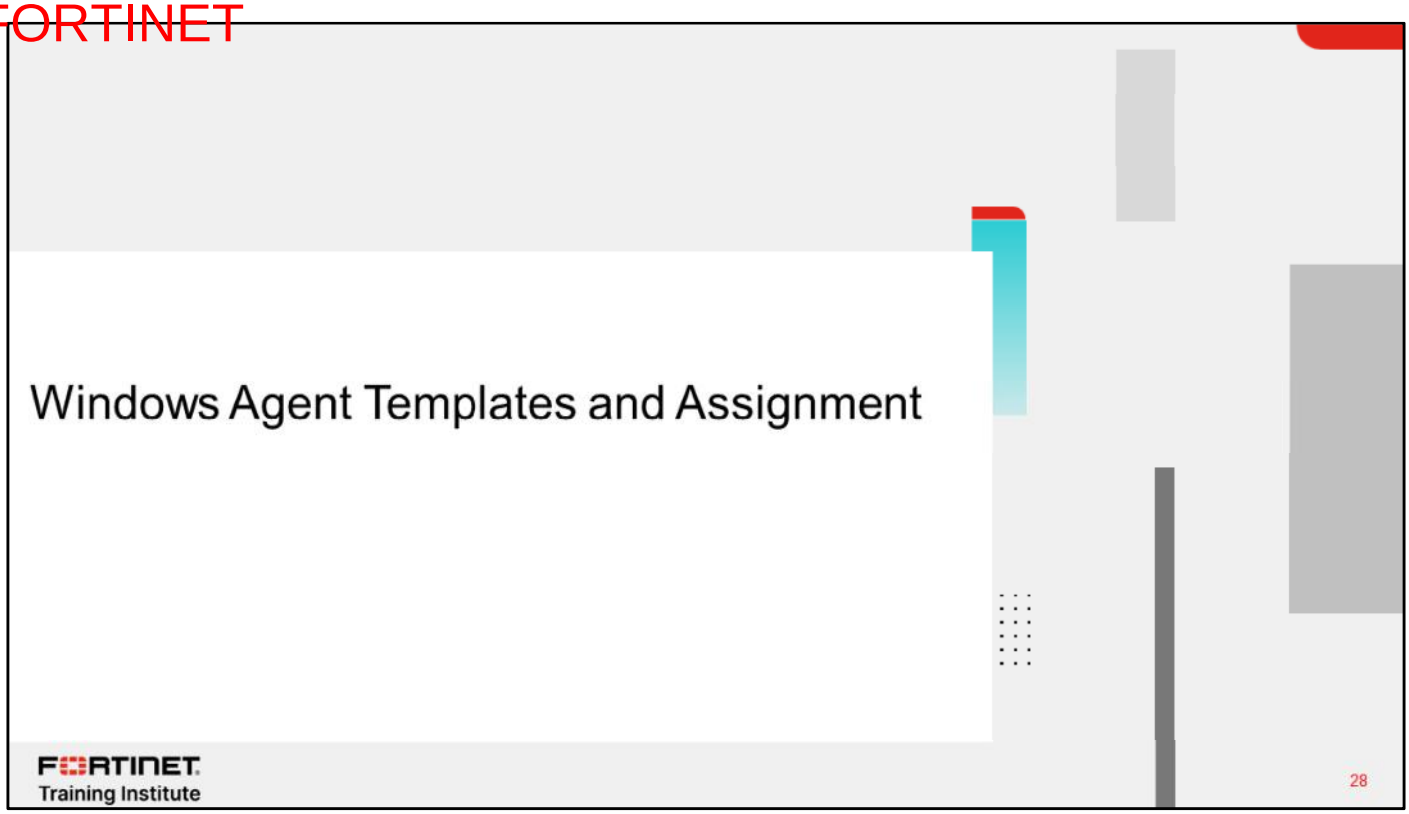
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

27

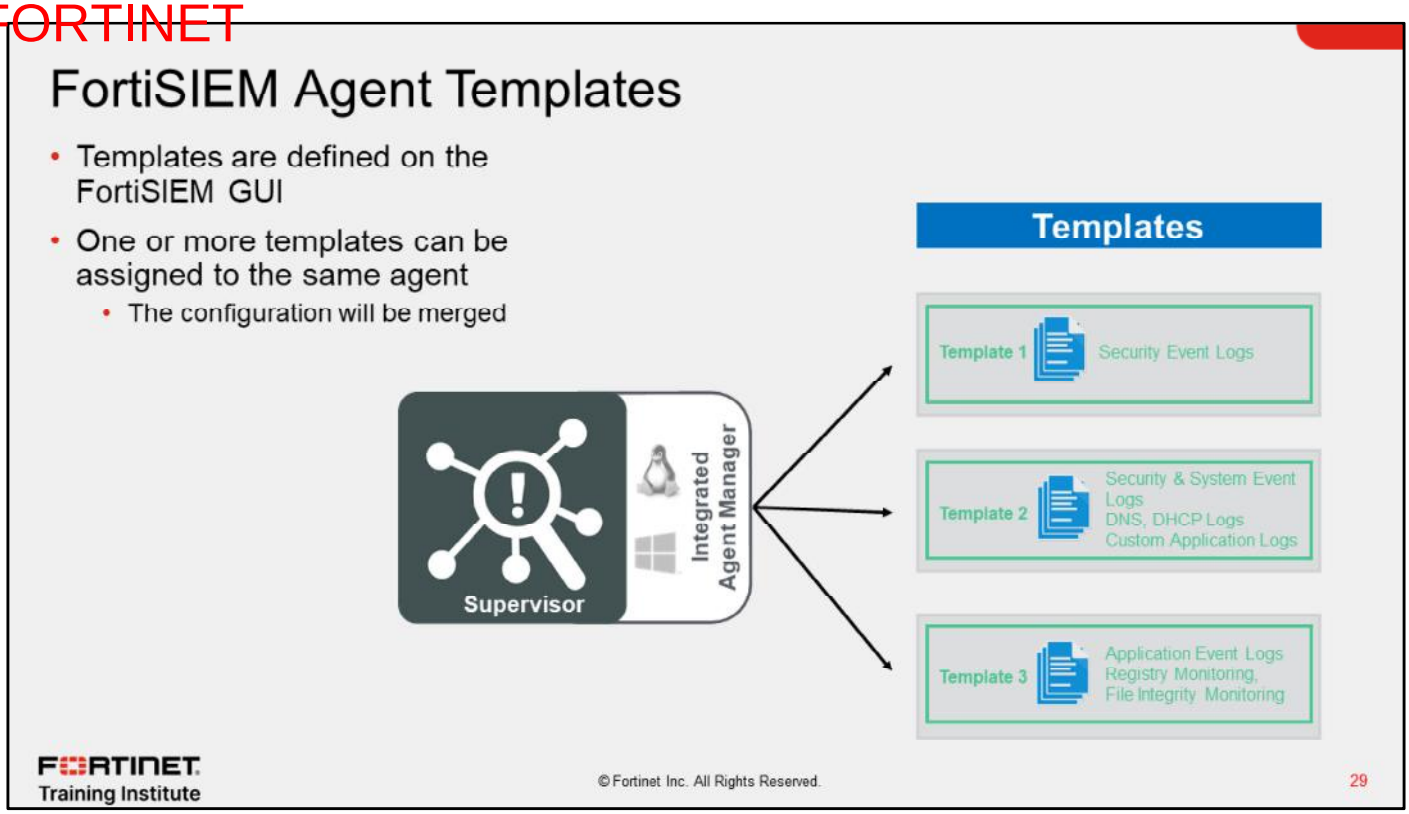
This slide shows a summary of the various agent statuses. These statuses apply to both Windows and Linux agents.

DO NOT REPRINT
© FORTINET



In this section, you will learn about Windows agent templates and the types of templates that can be assigned to a Windows agent.

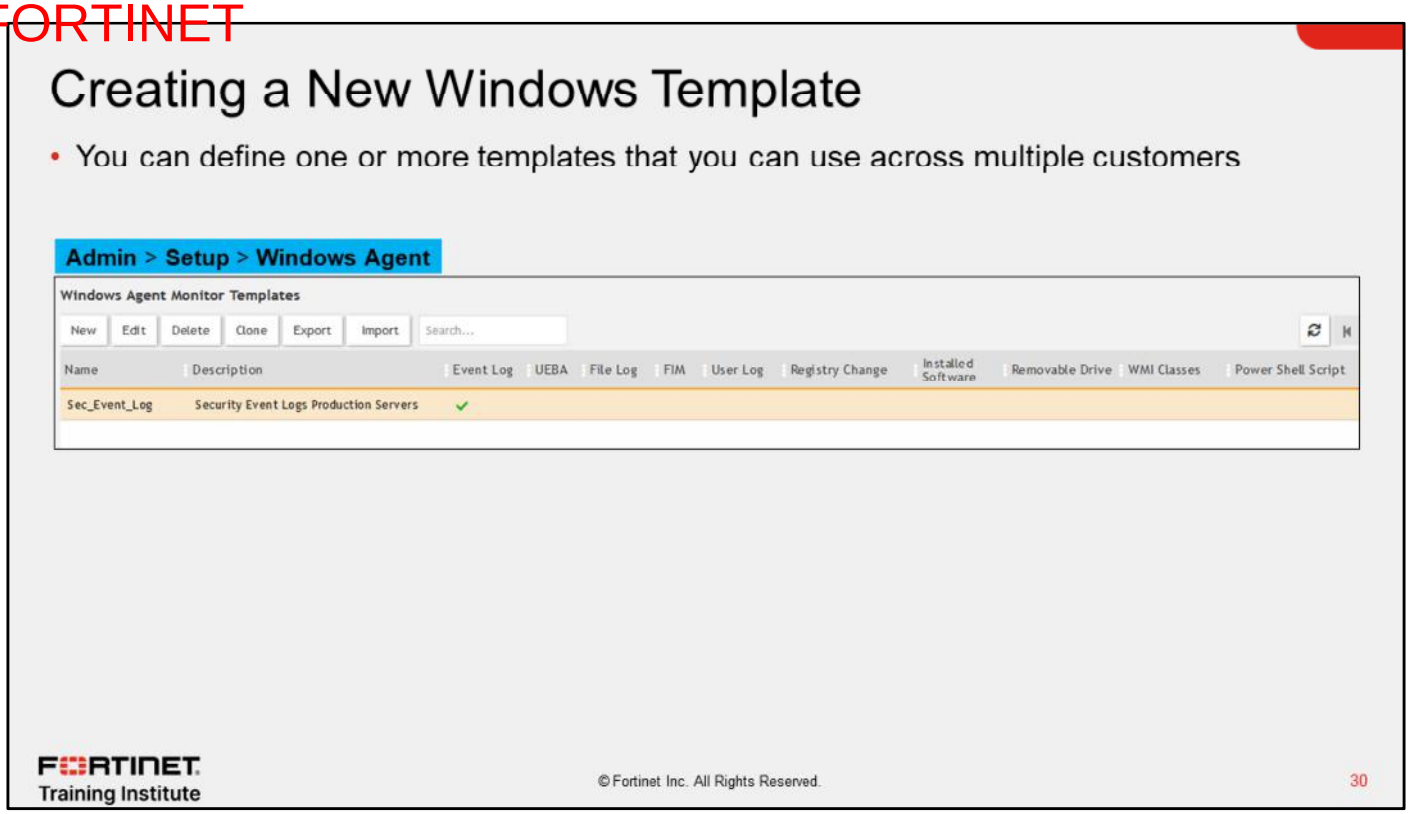
DO NOT REPRINT
© FORTINET



Templates are defined on the FortiSIEM GUI. Templates define what type of logs the agent will monitor and upload, such as security event logs, system event logs, DNS logs, DHCP logs, custom application logs, file integrity monitoring logs, and so on.

You can assign one or more templates to the same agent and the configuration will be merged. For example, you might use template 1 to collect Windows event logs, and use template 2 to collect Windows event logs and application logs from DNS and DHCP, and so on.

DO NOT REPRINT
© FORTINET

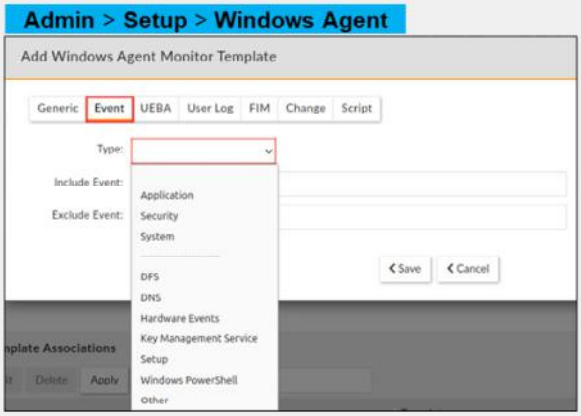


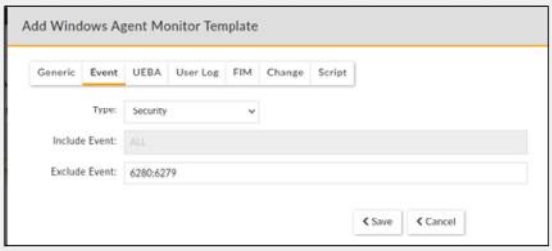
You can create one or more templates and use it across multiple customers. You must be logged in as FortiSIEM super admin to create templates, and the template name must not contain spaces.

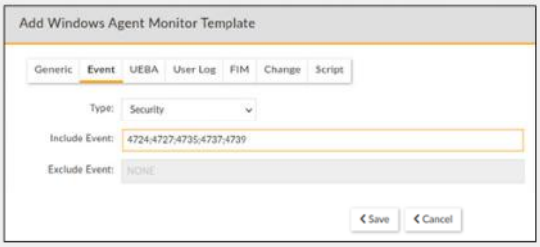
DO NOT REPRINT
© FORTINET

Defining Basic Logging for the Template

- The **Event** tab allows the definition of which event logs to collect such as security, system, or application







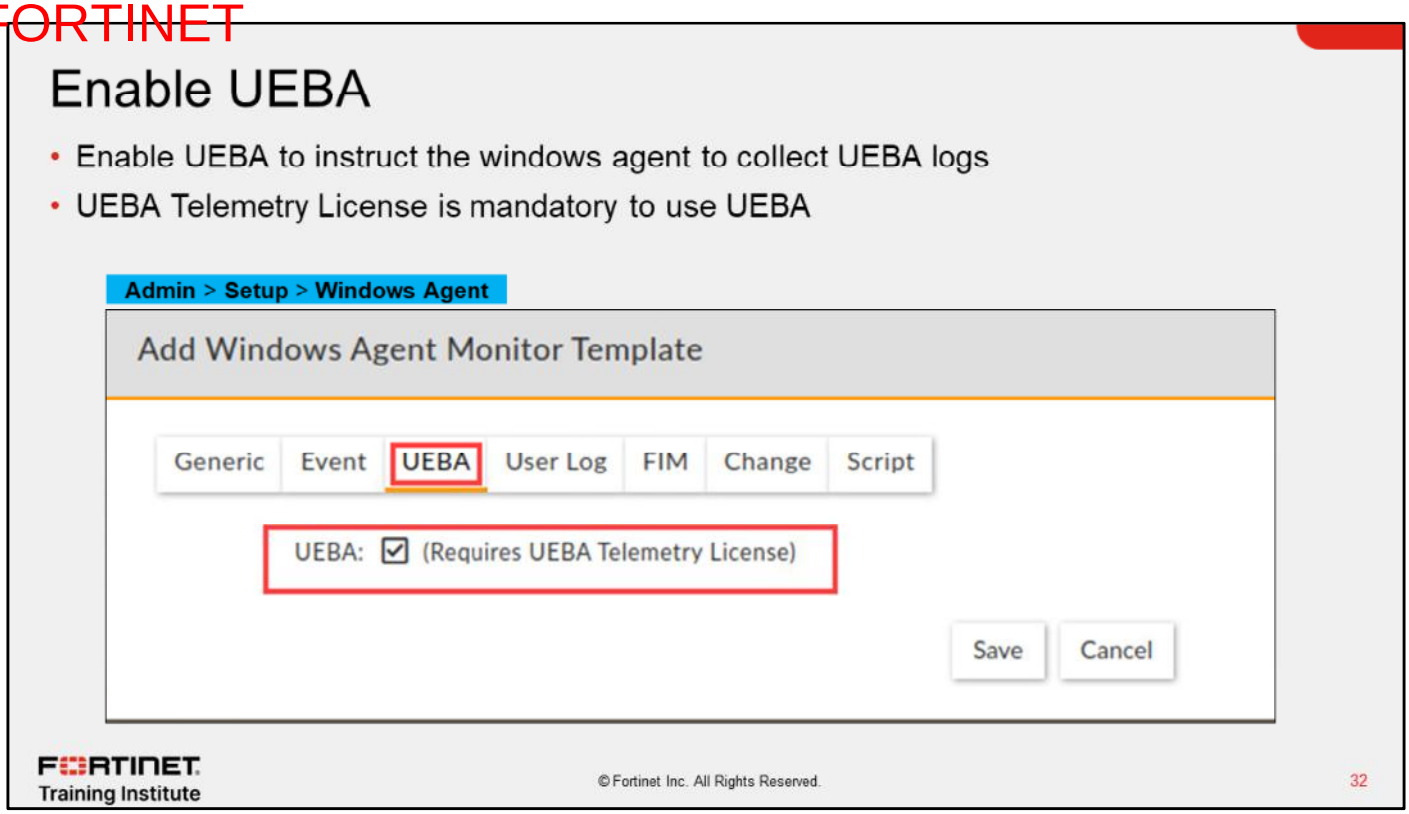
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

31

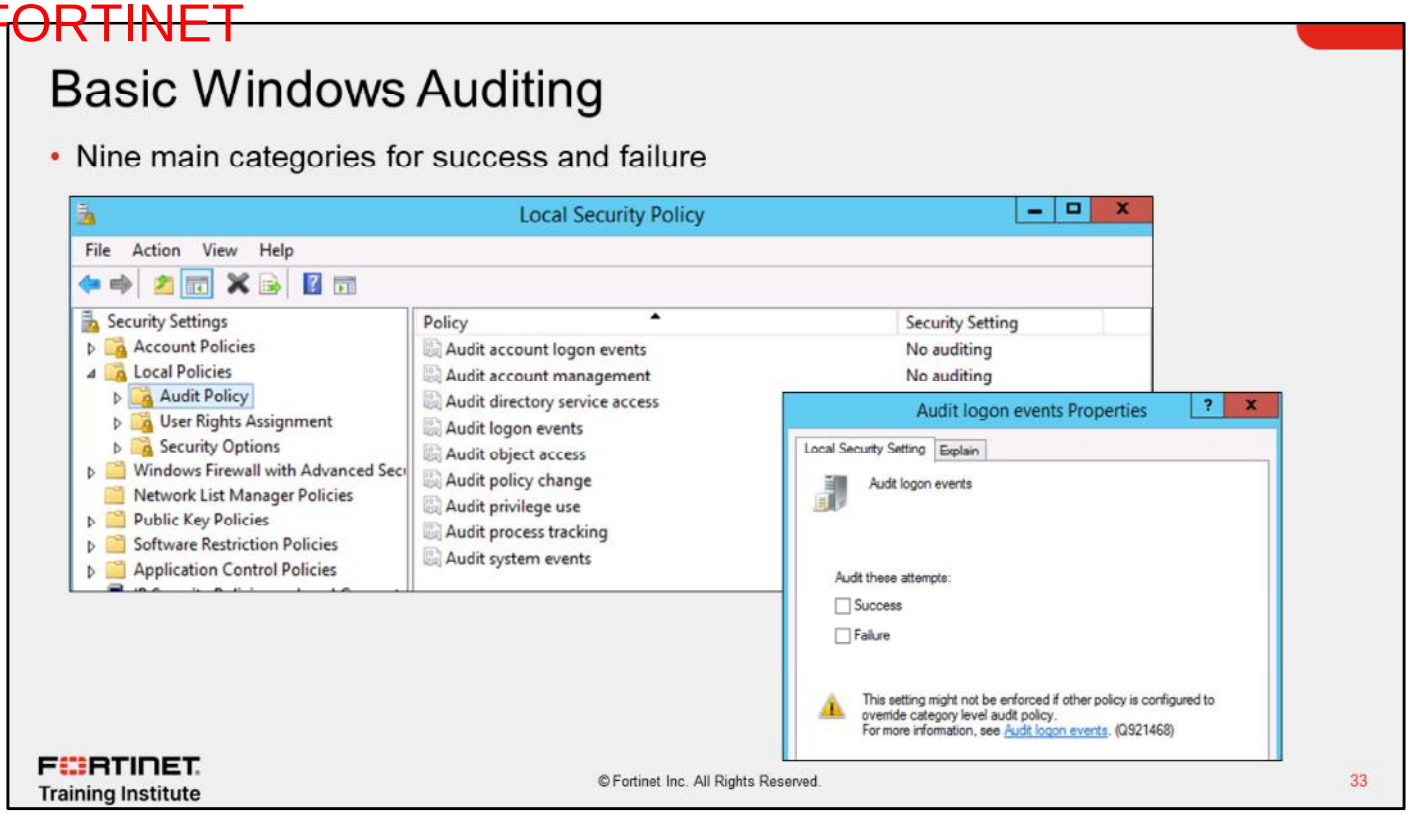
On the **Event** tab, you can define which event logs to collect, such as security, system, or application. You can filter all event logs by specifying a semi-colon separated list of event IDs, as either an include or exclude list.

DO NOT REPRINT
© FORTINET



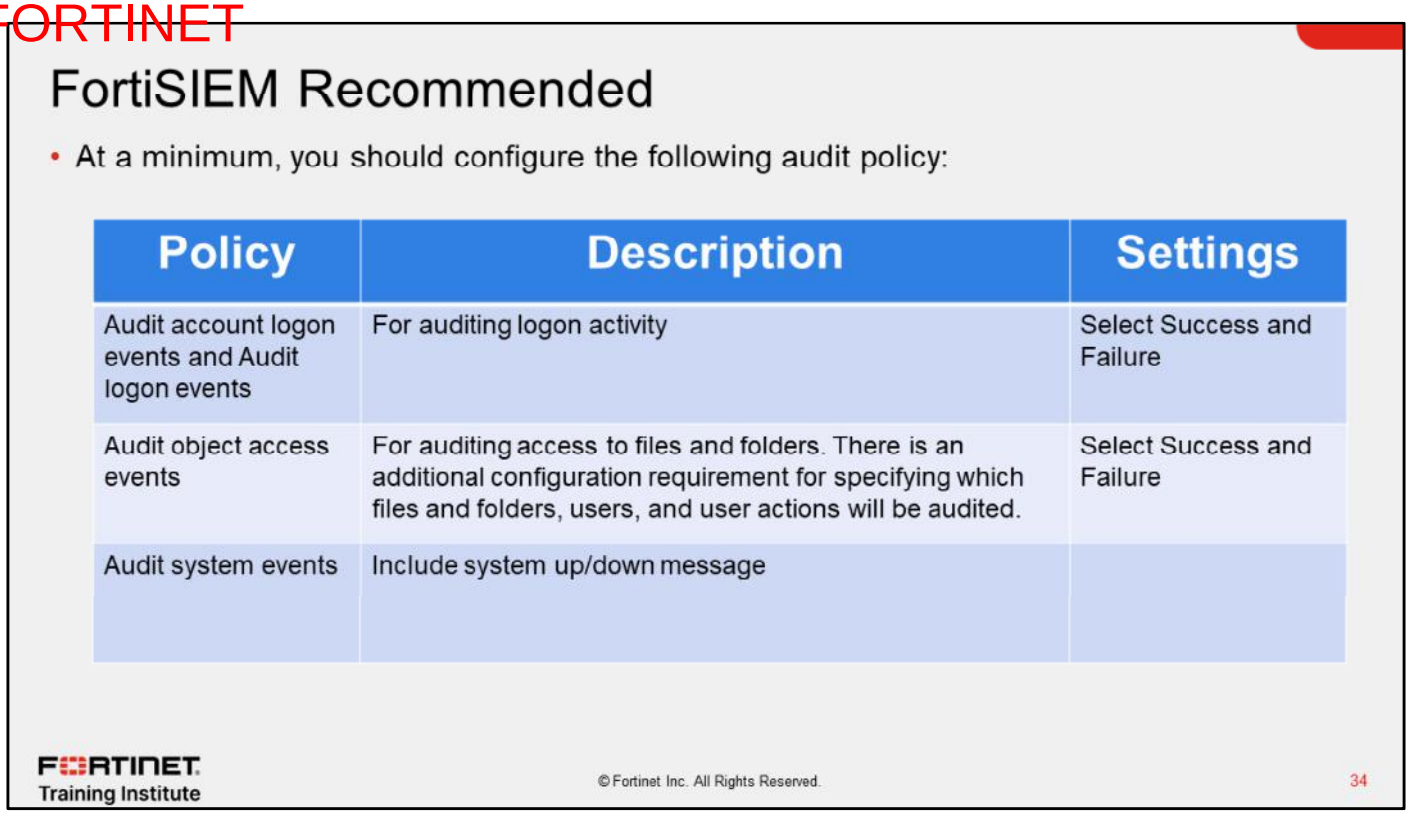
On the windows agent template enable UEBA to instruct the windows agent to collect UEBA logs. You must have UEBA Telemetry License to allow the windows agent to collect UEBA logs.

DO NOT REPRINT
© FORTINET



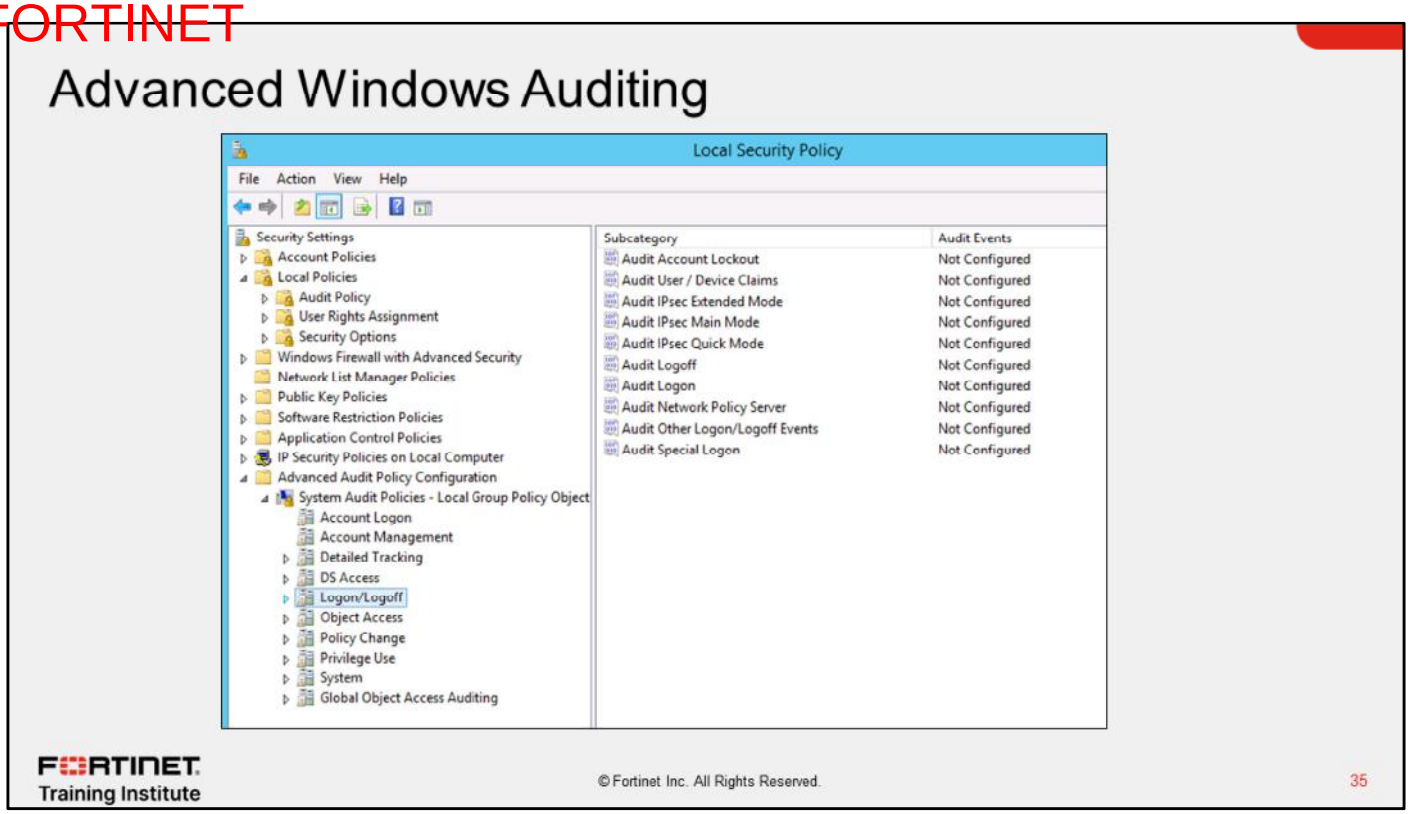
A Windows system audit policy determines which type of information about the system you'll find in the security logs. Windows uses nine audit policy categories and 50 audit policy subcategories to give you more granular control over which information is logged.

DO NOT REPRINT
© FORTINET



At a minimum, Fortinet recommends that you enable the **Audit account logon events** and **Audit logon events** policies for auditing logon activity. Enable **Audit object access** events for auditing access to files and folders. Enable **Audit system events** for system up and down status messages. Most forensic data requires you to enable additional policies.

DO NOT REPRINT
© FORTINET



The purpose of security auditing is to ensure that events are logged whenever an activity occurs. However, when every activity is audited, event logs become flooded with irrelevant information that makes it difficult for network administrators to separate critical events from insignificant ones. Advanced audit policy settings help administrators exercise granular control over which activities get recorded in the logs, helping cut down on event noise.

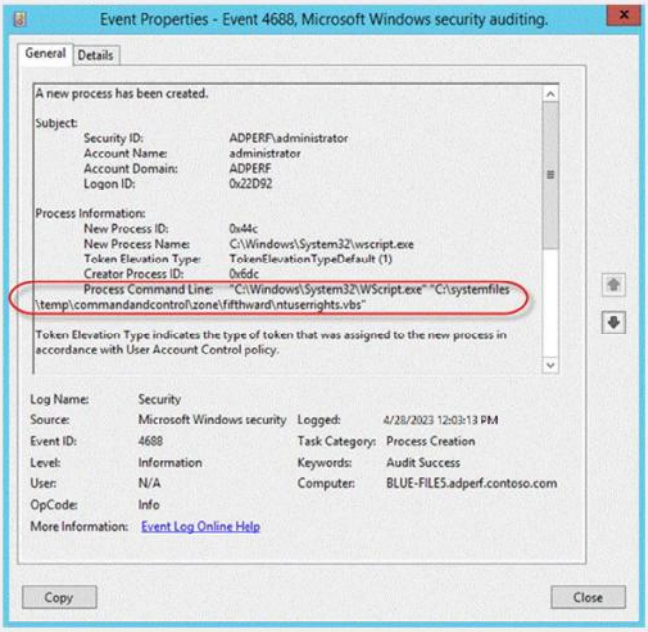
For example, instead of turning on the **DS Access** audit policy category to troubleshoot a replication problem—which would generate around eight events every time this activity occurs—an administrator could turn on the advanced audit policy subcategory for directory service replication, which would only generate one event instead of eight.

Do not fall for the recommendation to enable only failure events for each category. A common misconception is that a failure-only audit policy will alert you to all suspicious events. In reality, many of the most important events in the security log—events such as changes to critical user accounts and groups, account lockouts, and changes to security settings—are success events.

DO NOT REPRINT
© FORTINET

Process Command Line Auditing

- Log the process command line when a new process is created
- Crucial for determining malicious use of standard operating system commands
- Needs to be enabled in auditing
  - Include command line in process creation events



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

36

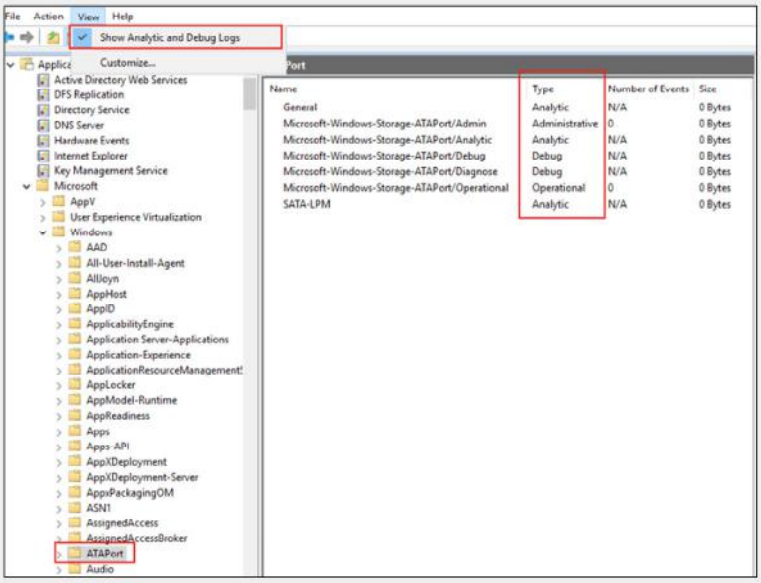
Event 4688 documents each program that is executed, who the program ran as, and the process that started this process.

When you start a program, you are creating a process that stays open until the program exits. This process is identified by the process ID. You can correlate this event to other events by the process ID to identify what the program did while it ran and when it exited. This is crucial for identifying malicious use of standard operating system commands.

DO NOT REPRINT
© FORTINET

Windows Advanced Event Logging

- From Windows Server 2008 R2 and later, Windows provides additional event logging capabilities
- **Applications and Services Logs** are new categories of event logs that store events from a single application or component
- These categories of logs include four main subtypes—admin, operational, analytic, and debug logs—but there are others
- Each event log may have one or more of the above subtypes, and some event logs appear only after installing particular applications



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

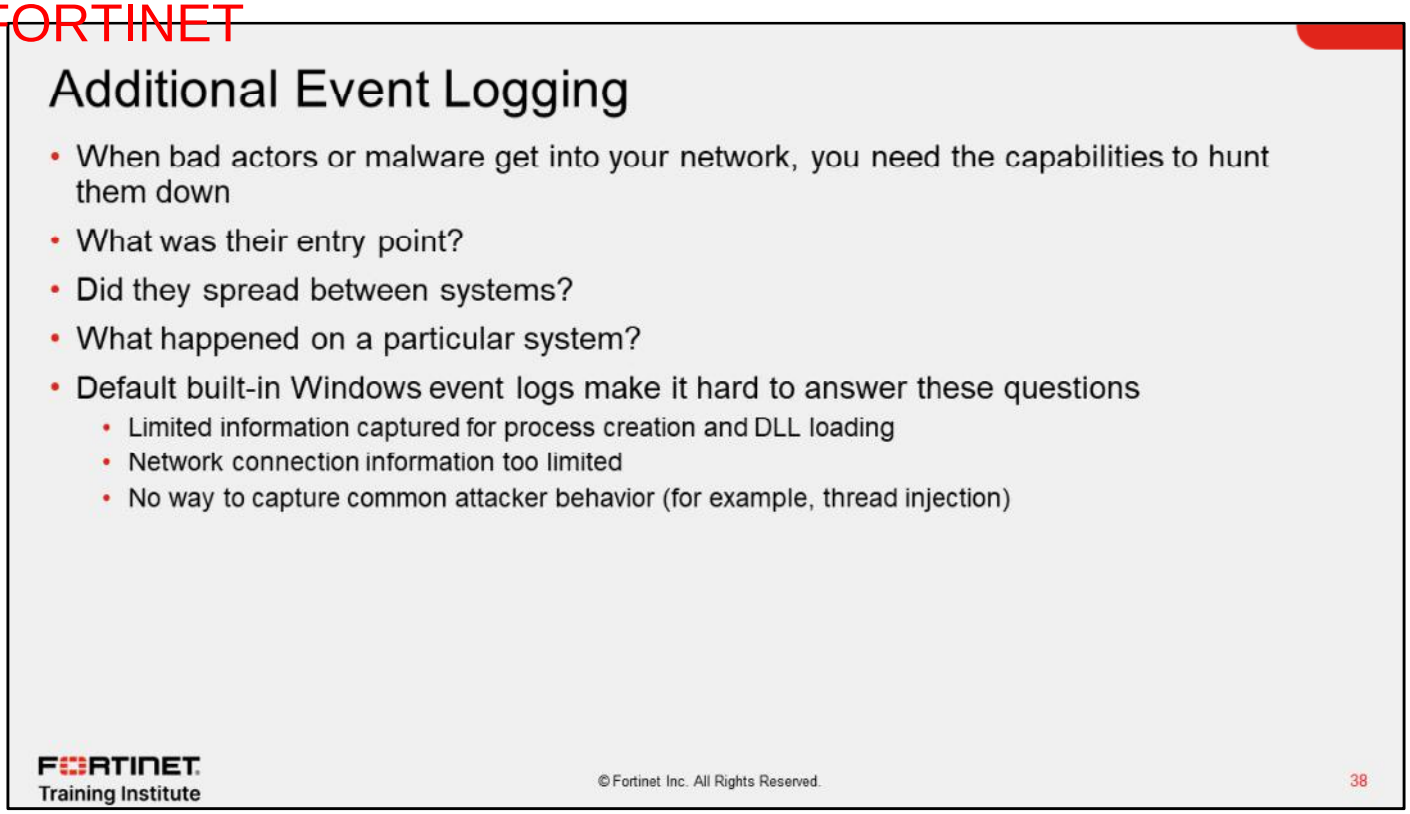
37

Not all these categories are suitable for collection by FortiSIEM because some are capable of creating huge volumes of events just like a debug command on a busy firewall.

Admin log types are useful for troubleshooting, targeted at end users, administrators, and support personnel. **Operational** log types are useful for analyzing and diagnosing a particular problem or occurrence. **Analytic** log types are events published in a high volume that trace an issue. **Debug** log types are used by developers to troubleshoot issues with applications.

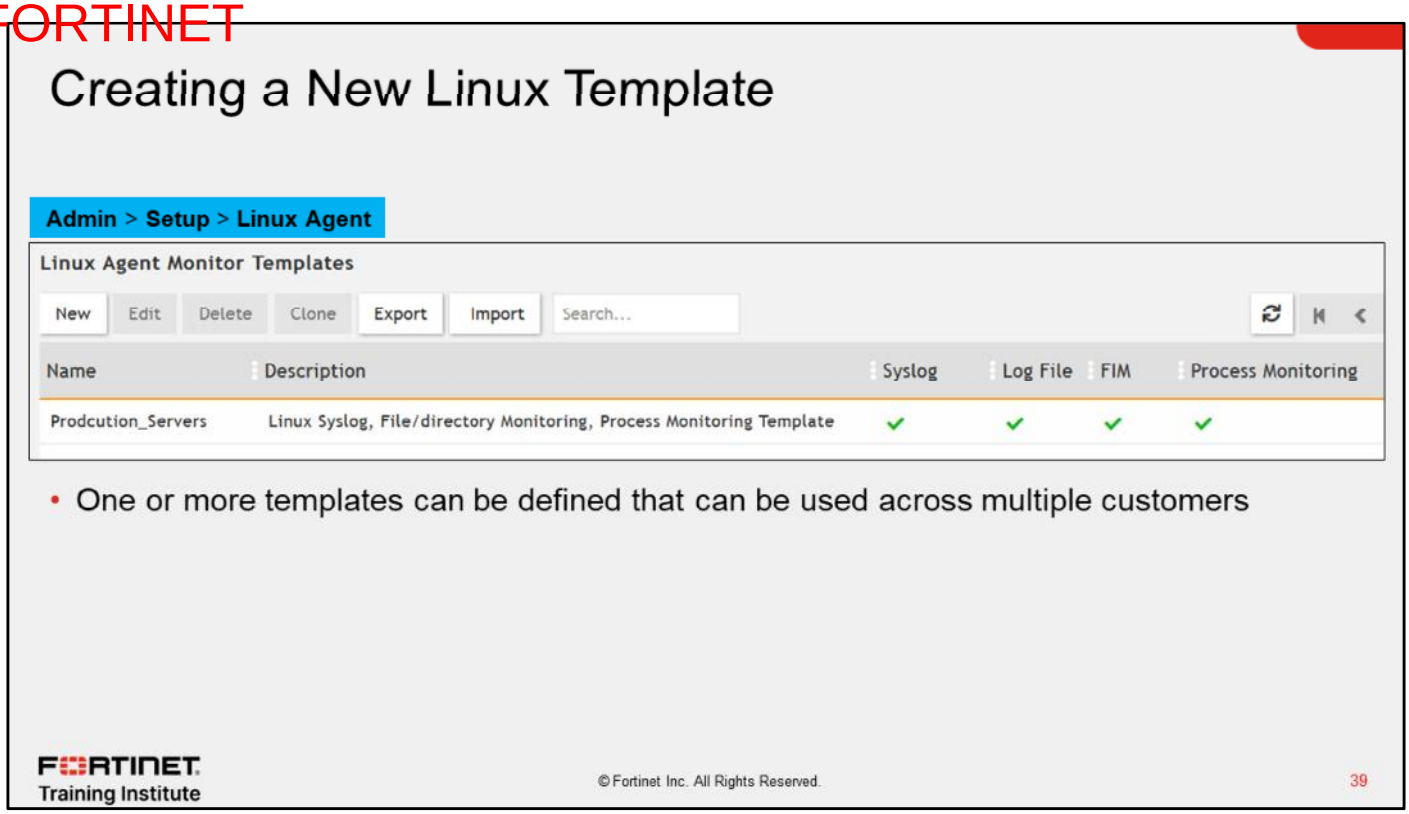
By default, both analytic and debug logs are hidden and disabled.

DO NOT REPRINT
© FORTINET



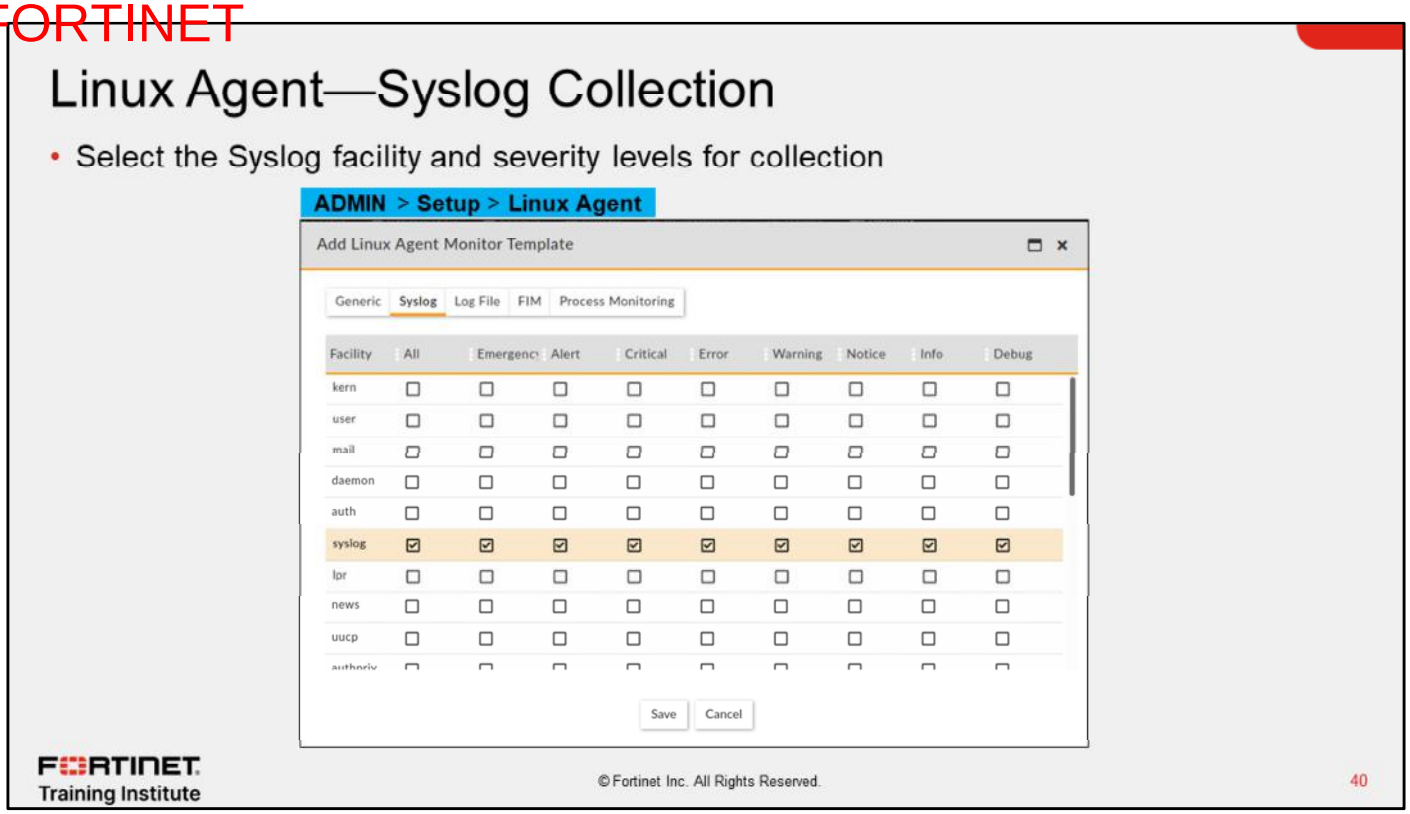
You need additional event logging to identify bad actors or malware that might get into your network. By analyzing the logs, FortiSIEM should be able to track their entry point, identify whether they spread between systems, and identify what happened on a particular system. The default built-in Windows event logs make it hard to answer these questions. There is limited information captured for process creation and DLL loading. Network connection information is also limited. There is also no way to capture common attacker behavior, such as thread injection.

DO NOT REPRINT
© FORTINET



Similar to Windows agents, you can also create one or more templates for Linux agents that can be used across multiple customers.

DO NOT REPRINT
© FORTINET



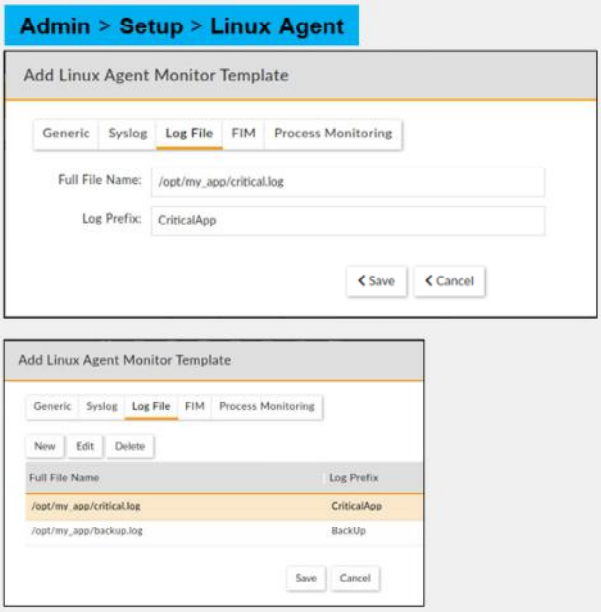
On the Linux template, you can select a Syslog facility and the severity levels for collection.

The facility represents the machine process that created the Syslog event. For example, is the event created by the kernel, mail system, security, or authorization processes? The facility represents a filter instructing the agent to forward only those events whose facility matches the one defined in this field. So, by changing the facility and the severity level, you change the number of messages that are sent by the agent.

DO NOT REPRINT
© FORTINET

Linux Agent—Custom Log Files

- The **Log File** specifies the path to the files or directories to be monitored
- The **Log Prefix** will be inserted into the Syslog header when delivered to FortiSIEM



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

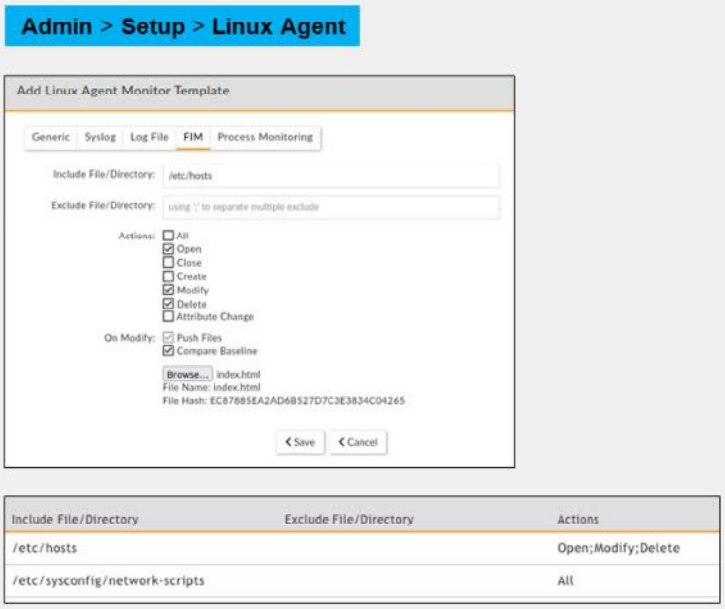
41

The **Log File** specifies the path to the files or directories that will be monitored. The **Log Prefix** will be inserted into the Syslog header when delivered to FortiSIEM.

DO NOT REPRINT
© FORTINET

Linux Agent—File Integrity Monitoring

- Available monitoring actions are:
  - **Open**, **Close**
  - **Create**, **Modify**, **Delete**
  - **Attribute Change**
    - Changing the directory or file permissions (that is, chmod)
  - **All**
- The **Modify** action will contain an MD5 hash code for the new file
- Compare modified file with baseline file
- Agent will not do file integrity monitoring on `root` directory.



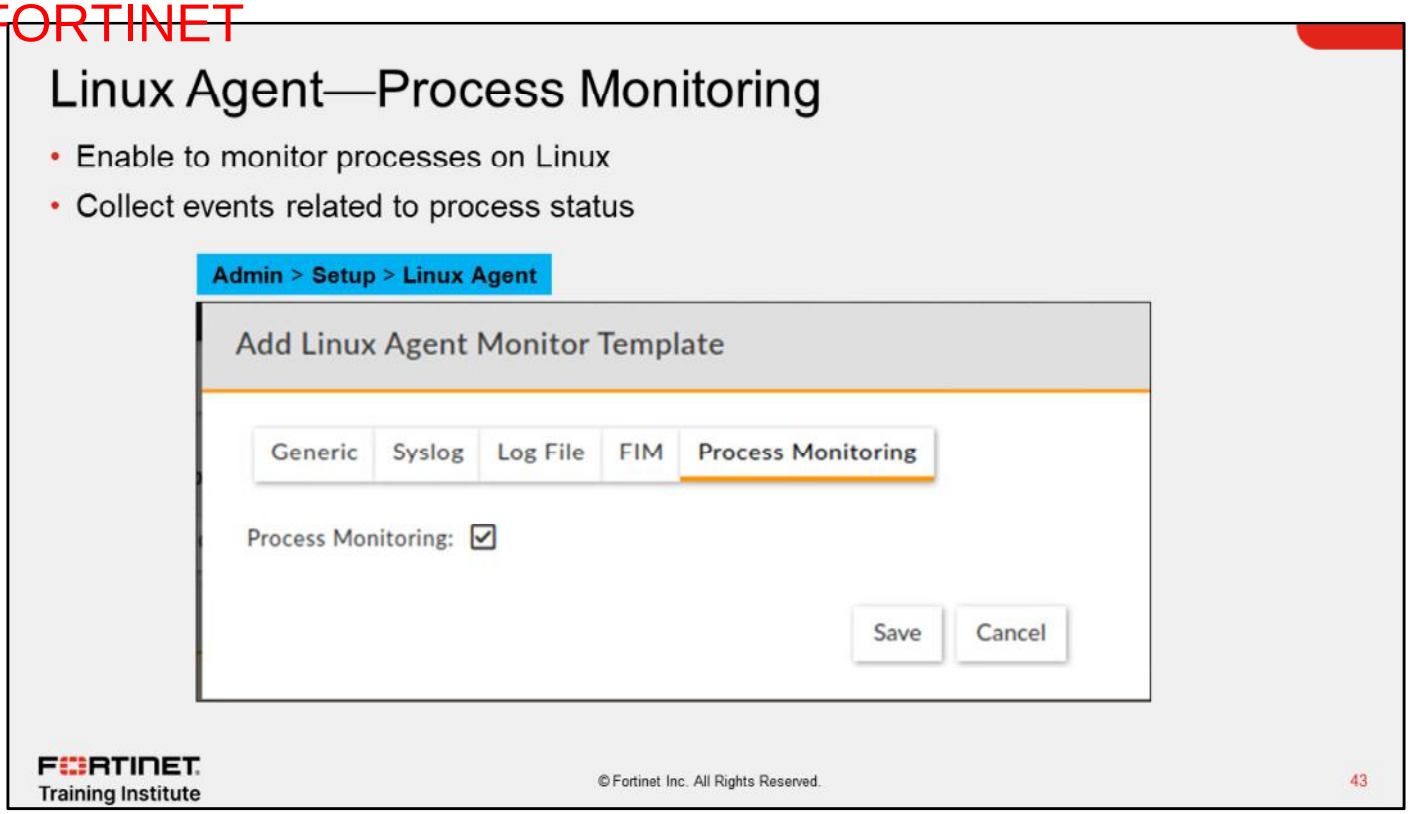
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

42

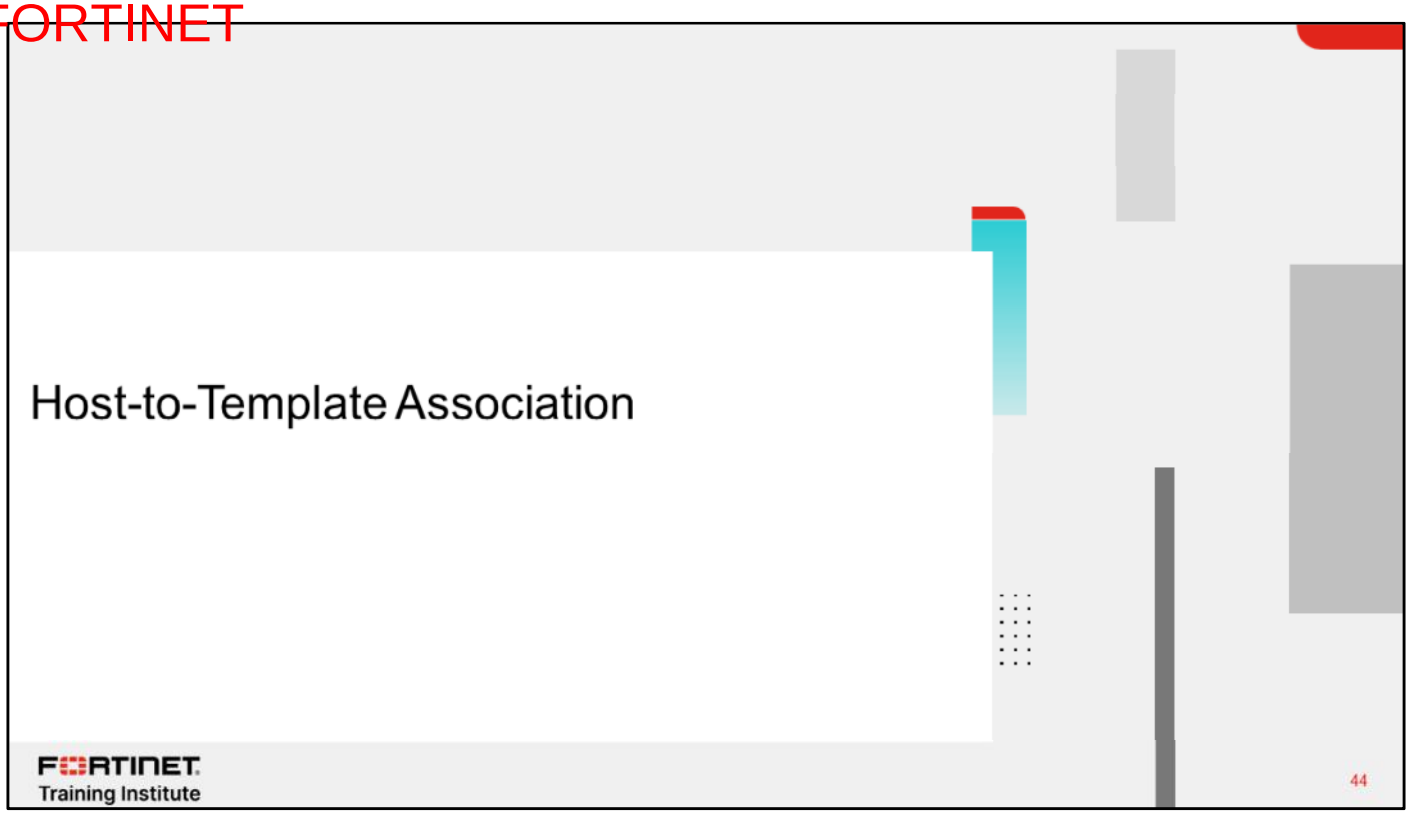
You can use file integrity monitoring to make sure that critical files and directories on servers are not modified. The **Modify** action will contain an MD5 hash code for the new file. By default, the agent will not perform any file integrity monitoring on the `root` directory. You can also push the modified file to FortiSIEM and you could also compare the hash of the modified file with a baseline file.

DO NOT REPRINT
© FORTINET



The Process Monitoring on Linux template is used collect events and logs related to process status on Linux devices.

DO NOT REPRINT
© FORTINET

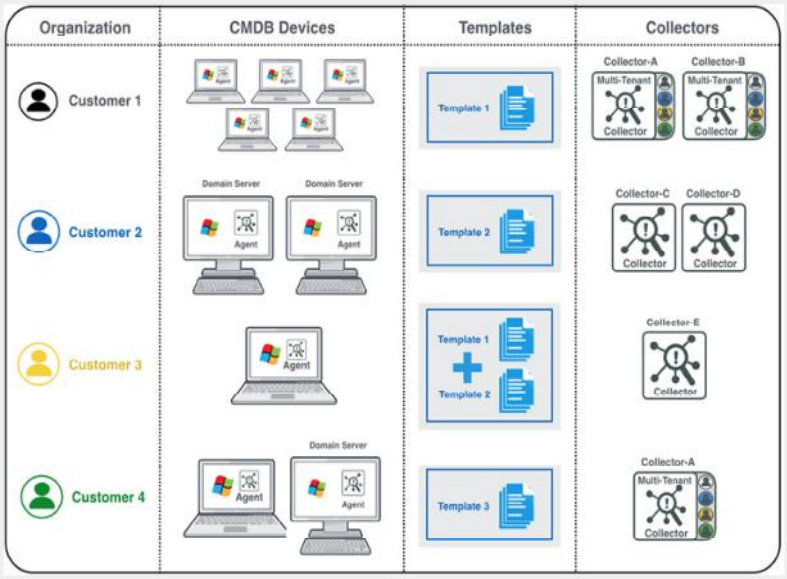


In this section, you will learn about associating a template with a host.

DO NOT REPRINT
© FORTINET

Associating Templates, Devices, and Collectors

- Devices from each organization are assigned one or more templates and one or more collectors
- Agents will load balance across multiple collectors



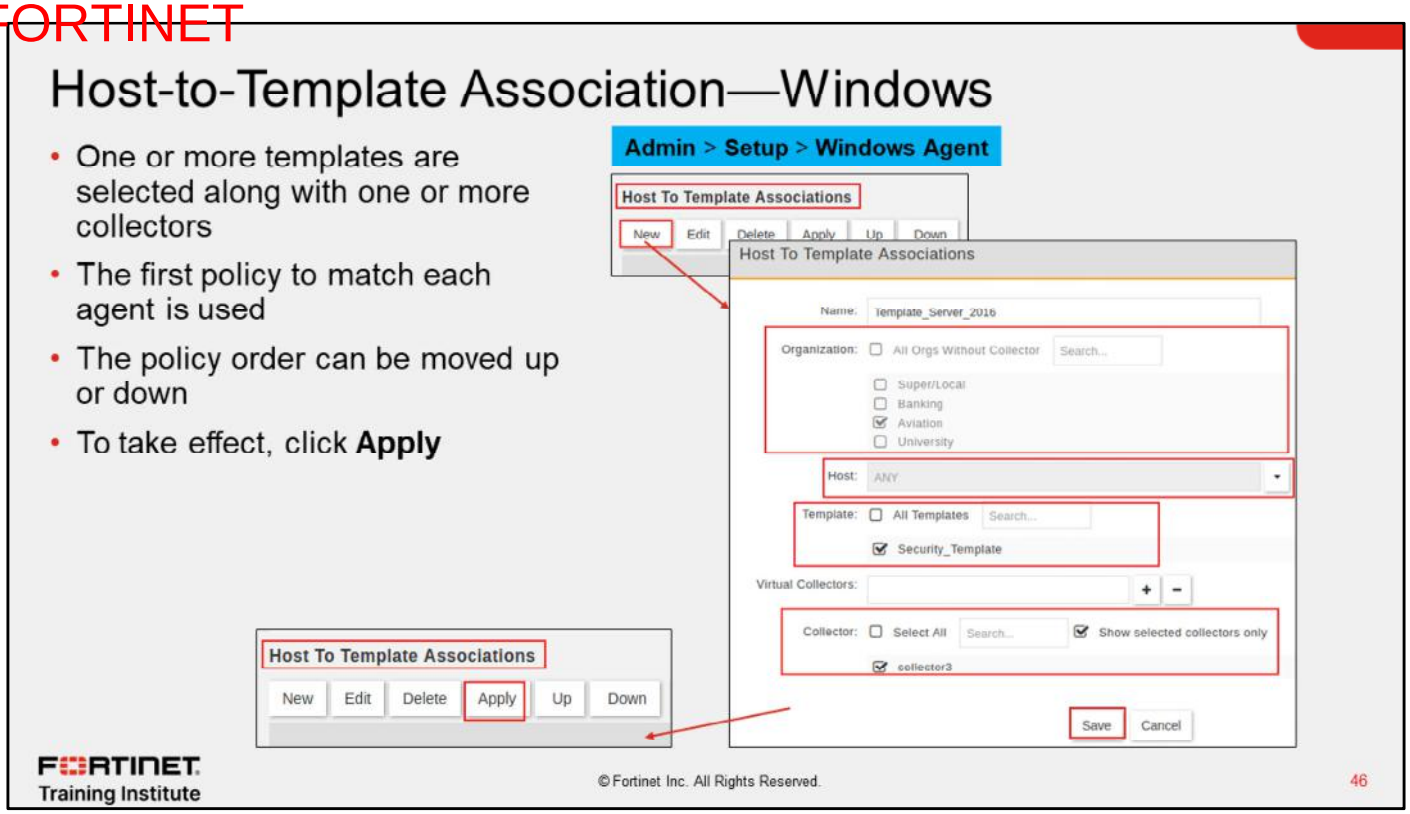
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

45

Devices from each organization are assigned one or more templates and one or more collectors. Agents will load balance logs across multiple collectors. For example, if a template is assigned to multiple devices for a customer with multi-tenant collectors, then the agents will load balance logs across all the collectors defined for that customer.

DO NOT REPRINT
© FORTINET

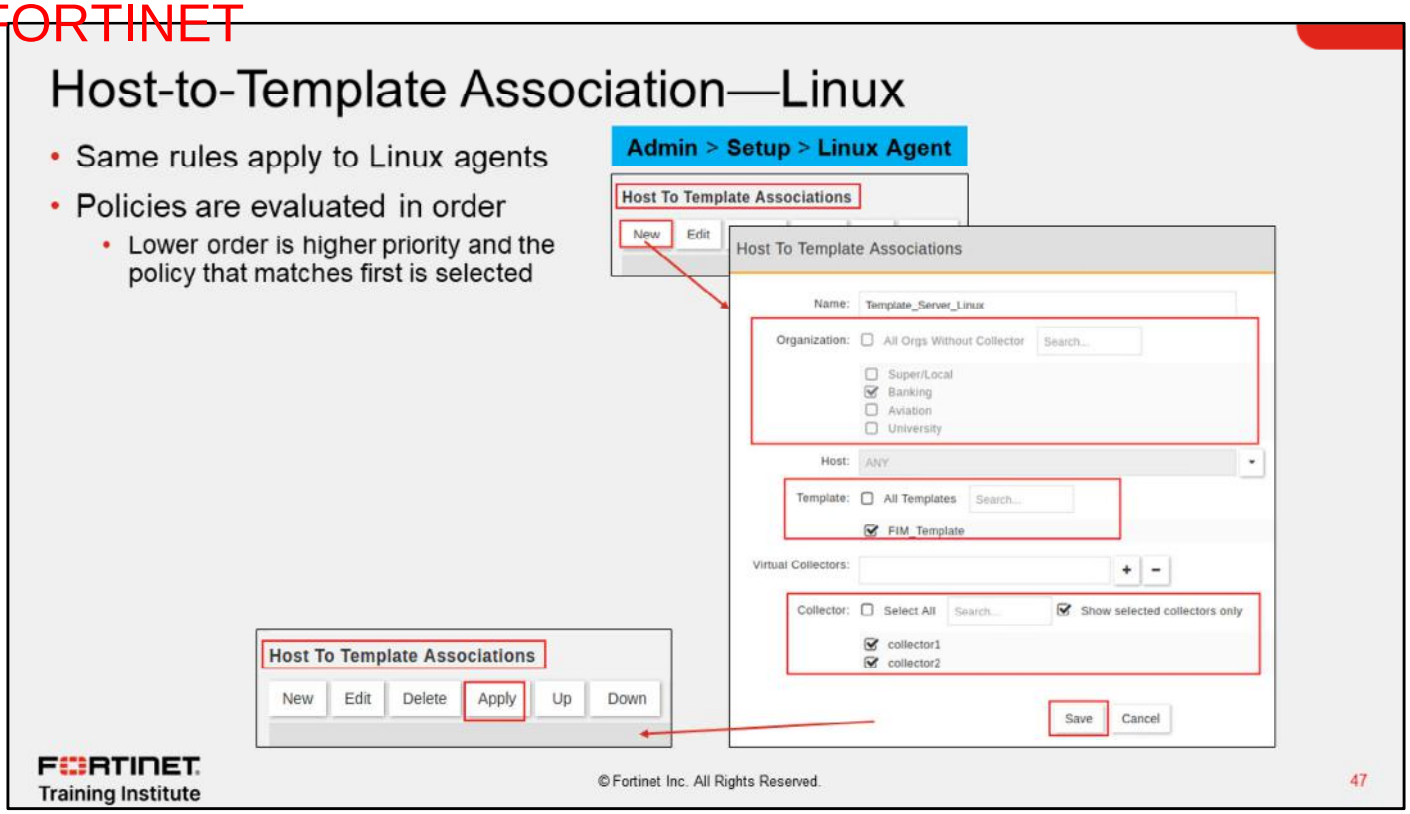


When you associate a host with a template, you have the option to select an organization with or without a collector. If an organization is selected with a dedicated collector, other organizations will be grayed out. If you select organizations without a collector, then FortiSIEM will allow any multi-tenant collector to be assigned. You can assign the template to *any* device within the organization, or you could select specific devices from the CMDB.

You can select a single template or multiple templates. If you select multiple templates, then the configuration will be merged. You can select a specific collector or multiple collectors for that organization. If you select multiple collectors, then logs sent by the agents will be load balanced across the collectors.

The first policy to match each agent will be used. You can move the policy order up or down. For the policy to take affect you must click **Apply**.

DO NOT REPRINT
© FORTINET

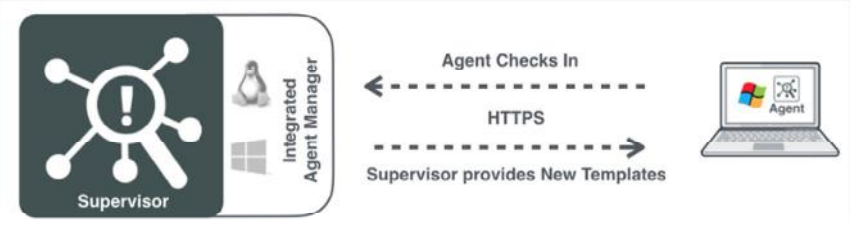


The same rules apply to Linux agents. Policies are evaluated in order, where lower order is higher priority. The policy that matches first is selected.

DO NOT REPRINT
© FORTINET

Agent Check In

- Agents will periodically check in with the supervisor every minute, and collect any new template associations
- It may take more than a minute for a new template to take effect



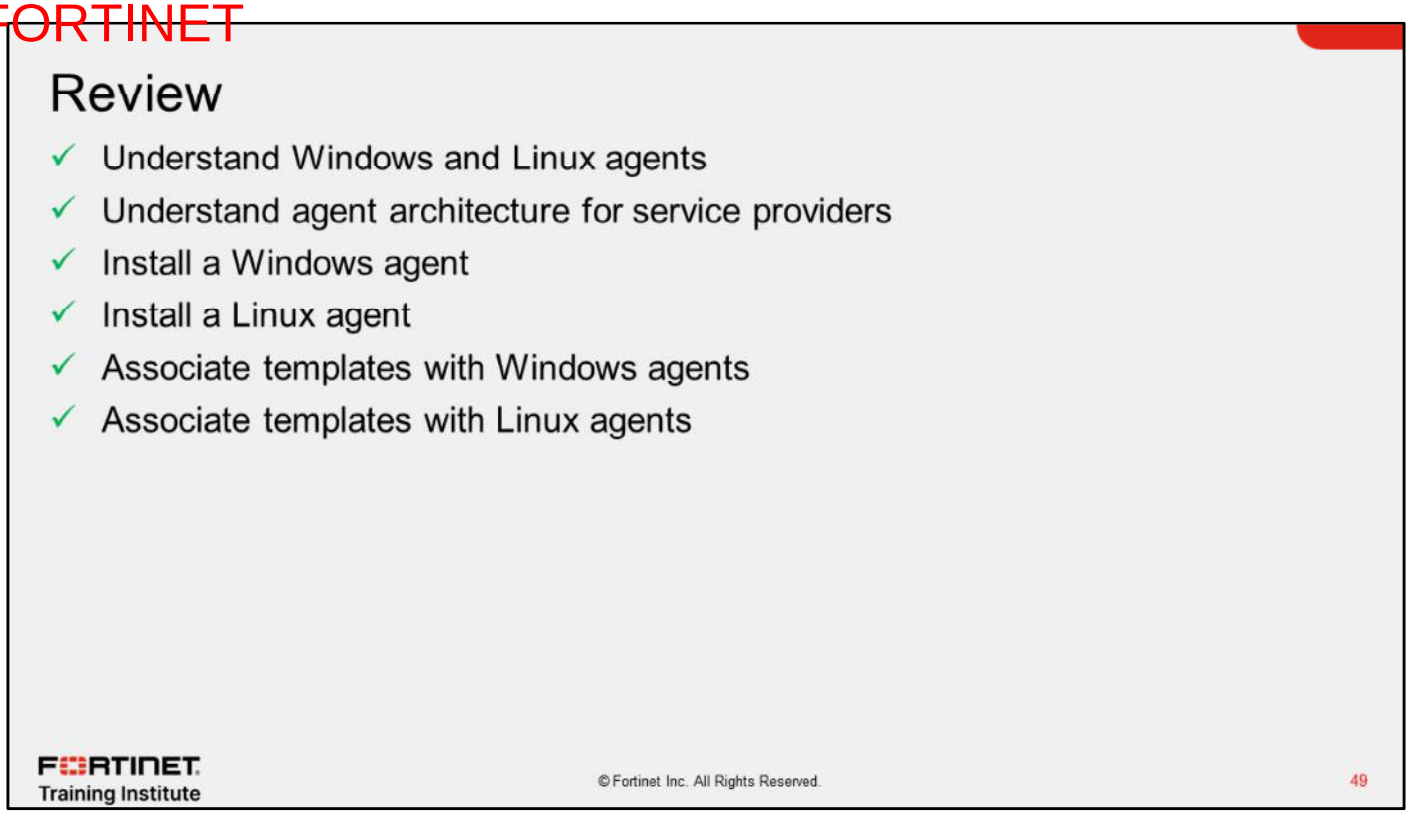
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

48

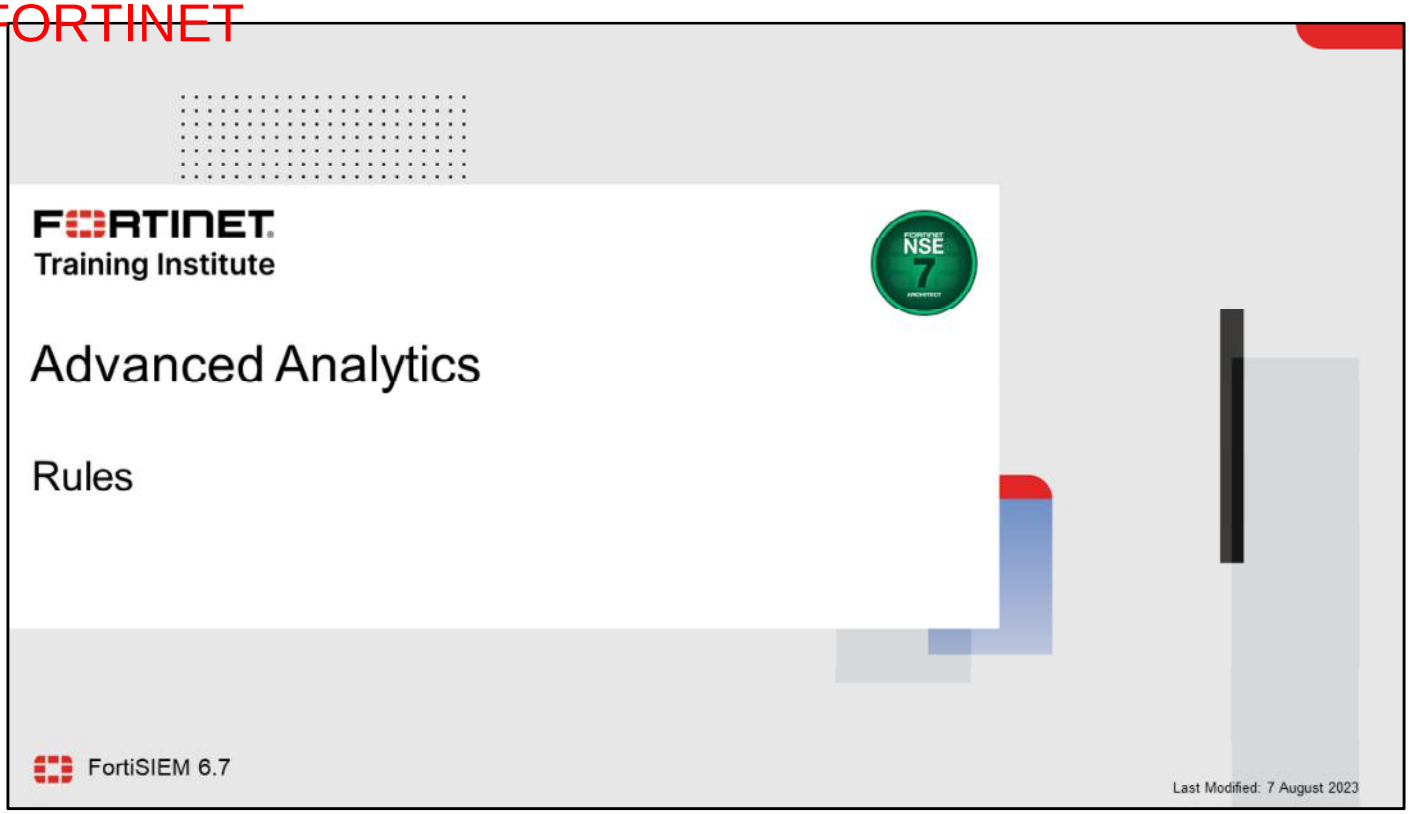
Agents will periodically check in with the supervisor every minute and collect any new template associations. It may take more than a minute for a new template to take effect.

DO NOT REPRINT
© FORTINET



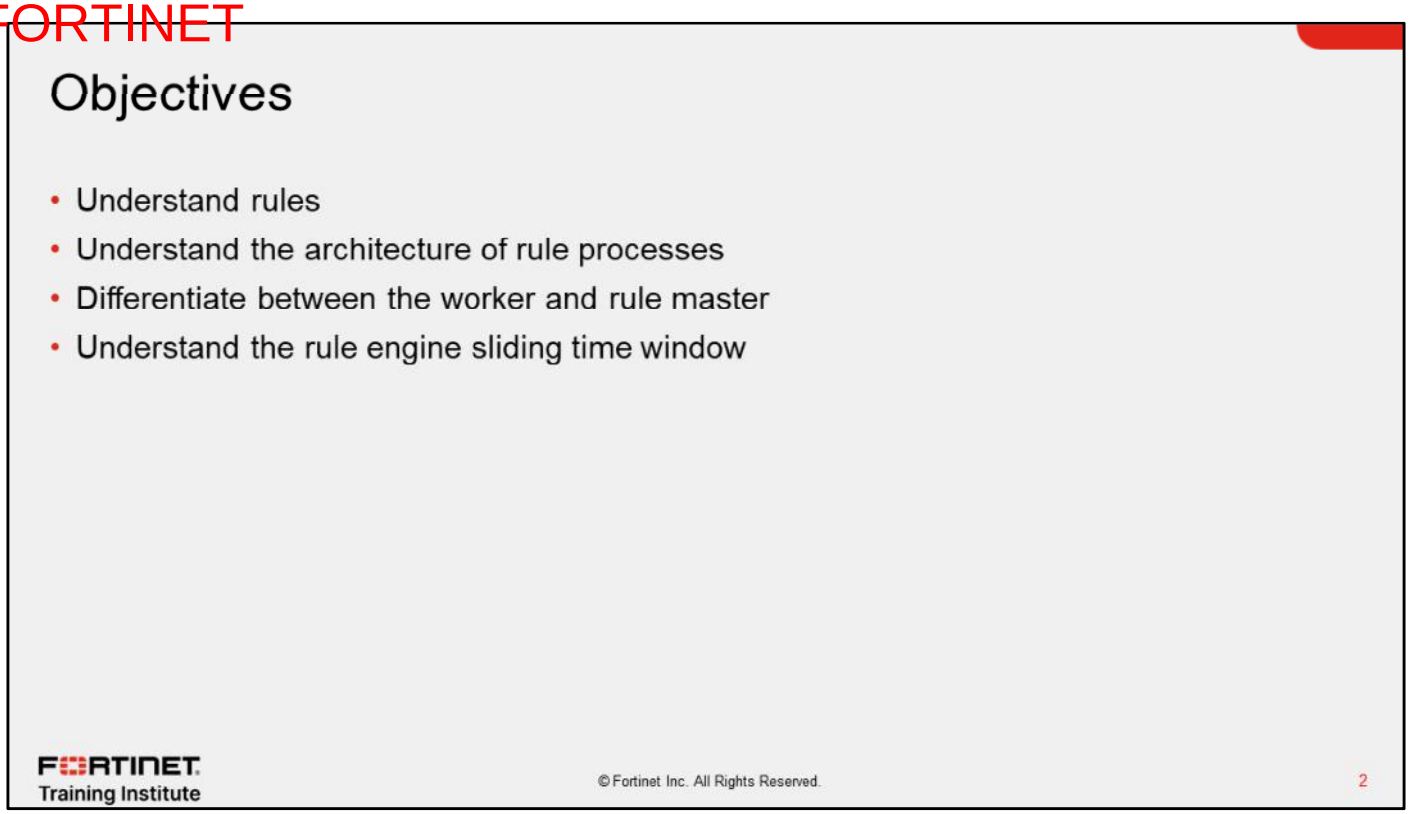
By mastering the objectives covered in this lesson, you learned about installing and managing FortiSIEM Windows and Linux agents.

DO NOT REPRINT
© FORTINET



In this lesson, you will learn about rules, including subpatterns in a rule and the back-end process that is involved in evaluating rules and triggering incidents. You will also learn about the sliding time window in a rule, out-of-box rules, and how each rule time window may be different.

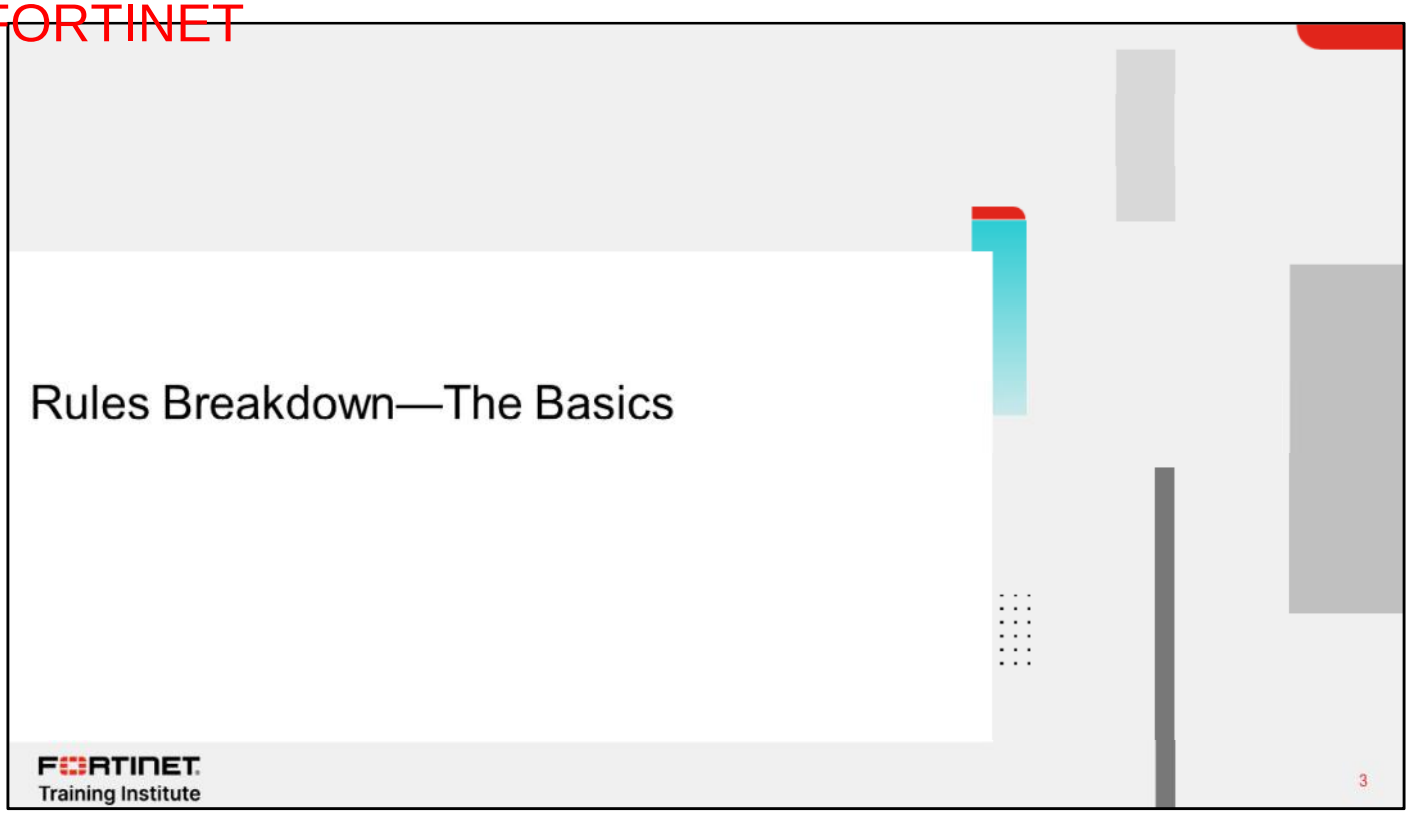
DO NOT REPRINT
© FORTINET



After completing this lesson, you should be able to achieve the objectives shown on this slide.

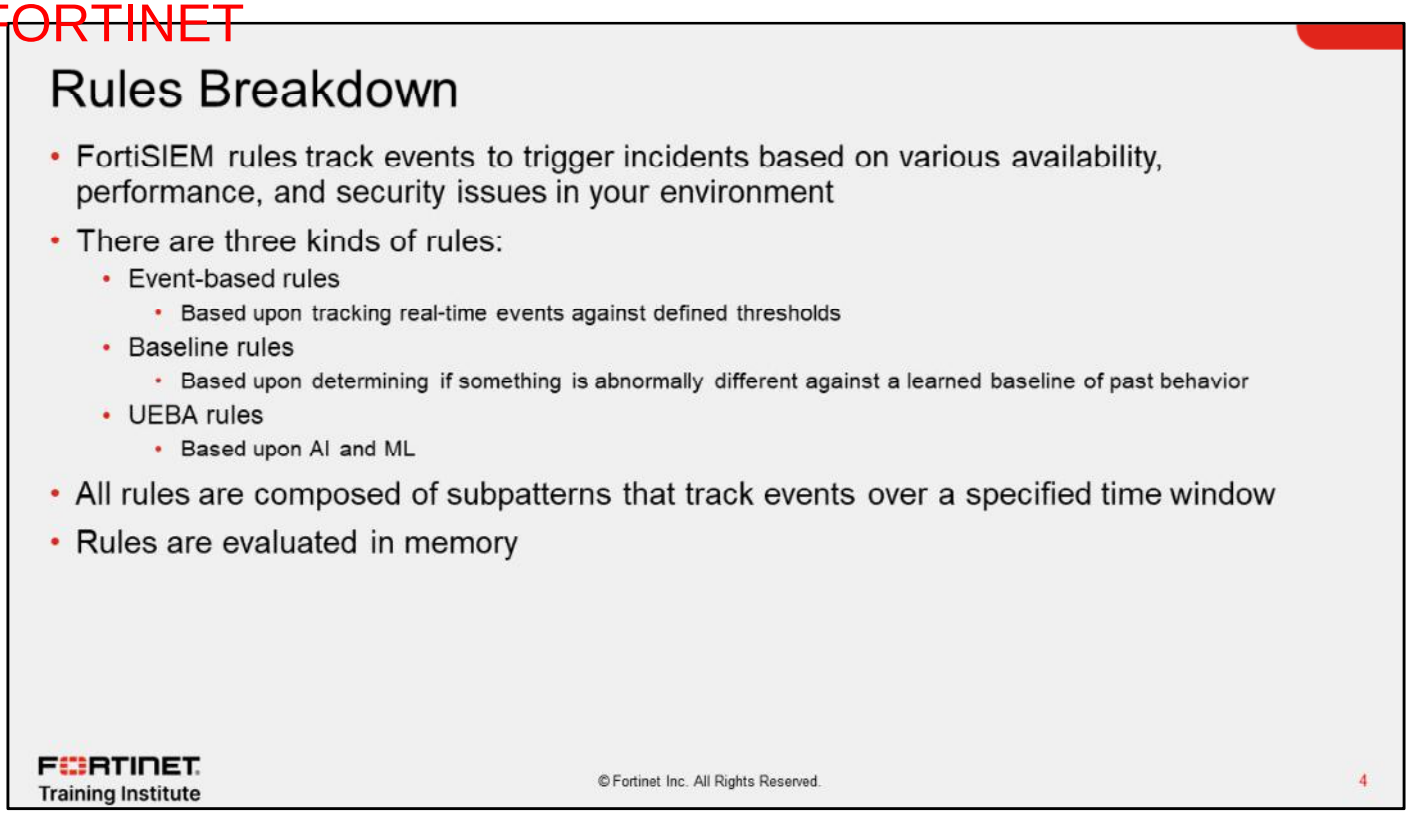
By demonstrating competence in understanding rules, you will be able to identify the purpose of each rule that is built into FortiSIEM. Also, by understanding the structure of rules you will be able to create your own custom rules.

DO NOT REPRINT
© FORTINET



In this section, you will learn the basics of rule breakdown.

DO NOT REPRINT
© FORTINET

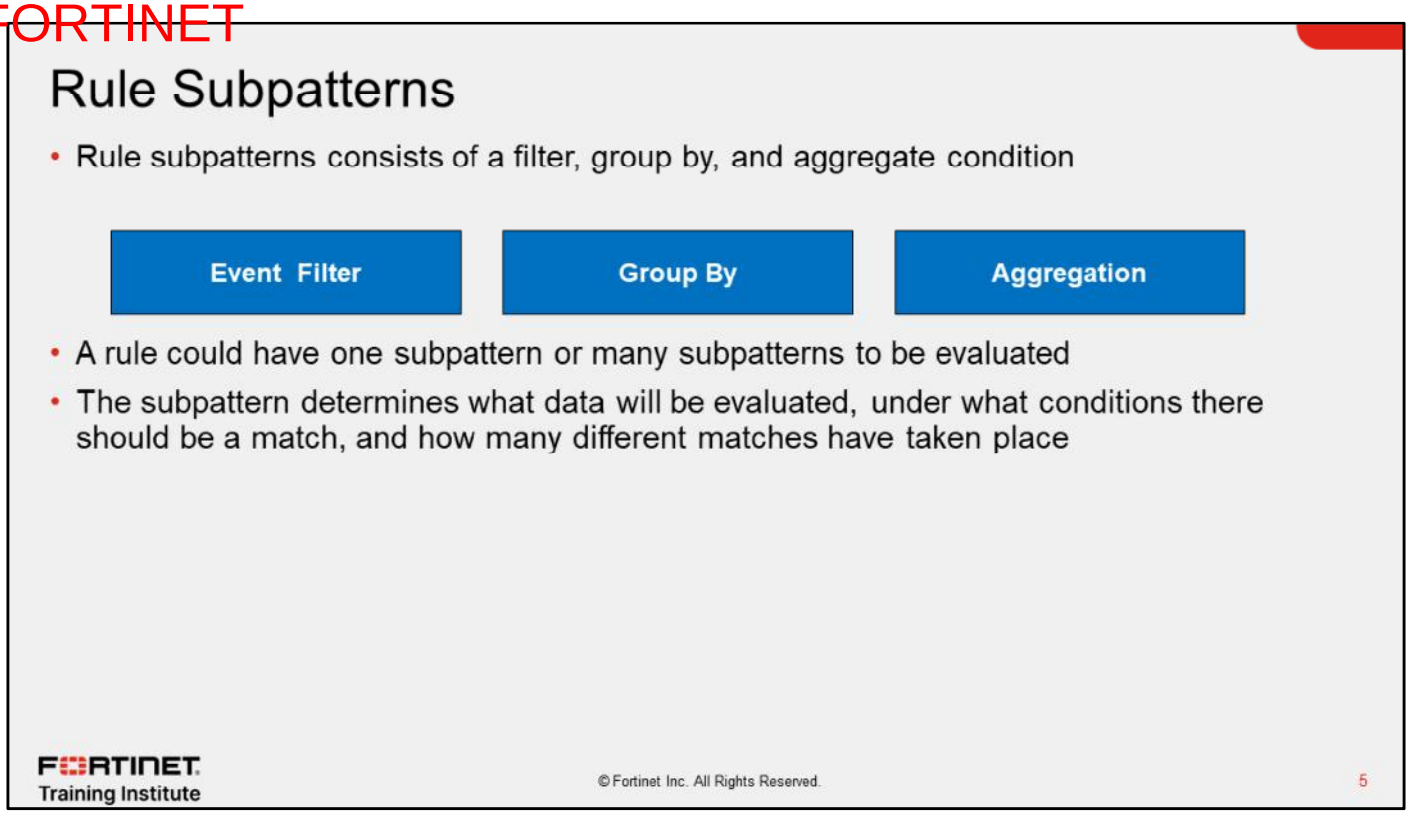


FortiSIEM rules track events to trigger incidents based on various availability, performance, and security issues in your environment.

There are three kinds of rules—event-based, baseline, and UEBA. Event-based rules are based on tracking real-time events against defined thresholds. Baseline rules are based upon determining if something is now abnormally different against a learned baseline of past behavior. UEBA rules are based upon AI and ML algorithm.

All rules are composed of subpatterns that track events over a specified time window and are evaluated in memory.

DO NOT REPRINT
© FORTINET



Rule subpatterns consists of a filter, group by, and aggregate condition. A rule could have one or many subpatterns to be evaluated. The subpattern determines what data will be evaluated, under what conditions there should be a match, and how many different matches have taken place.

DO NOT REPRINT
© FORTINET

Subpattern Filter

- You can think of the subpattern filter just like a SQL query, where the whole FortiSIEM event database is a huge table with many columns

Event Receive Time	Event Type	Source IP	Host IP	Destination TCP/UDP Port	Customer ID	Win Logon Id	User	Reporting Vendor	CPU Util	Raw Event Log
08:48:57 02/02/2022	FortiGate-Traffic-Allowed	10.0.0.1		80	1			Fortinet		<189>date=2022-02-02 time=08:47:00 dev
08:49:01 02/02/2022	FortiGate-event-login-success	10.0.2.15			1		Admin	Fortinet		<189>date=2022-02-02 time=08:49:00 dev
08:49:24 02/02/2022	Win-Securiity-4624	10.0.5.17			1	0x3B48215C	Bryan	Microsoft		Fri Feb 2 AccelOps-WinLog-WMI

- There may be many attributes that are not all populated with data (NULL), because different events contain different attributes
- The subpattern filter determines which events the rule should be evaluating

FORTINET Training Institute

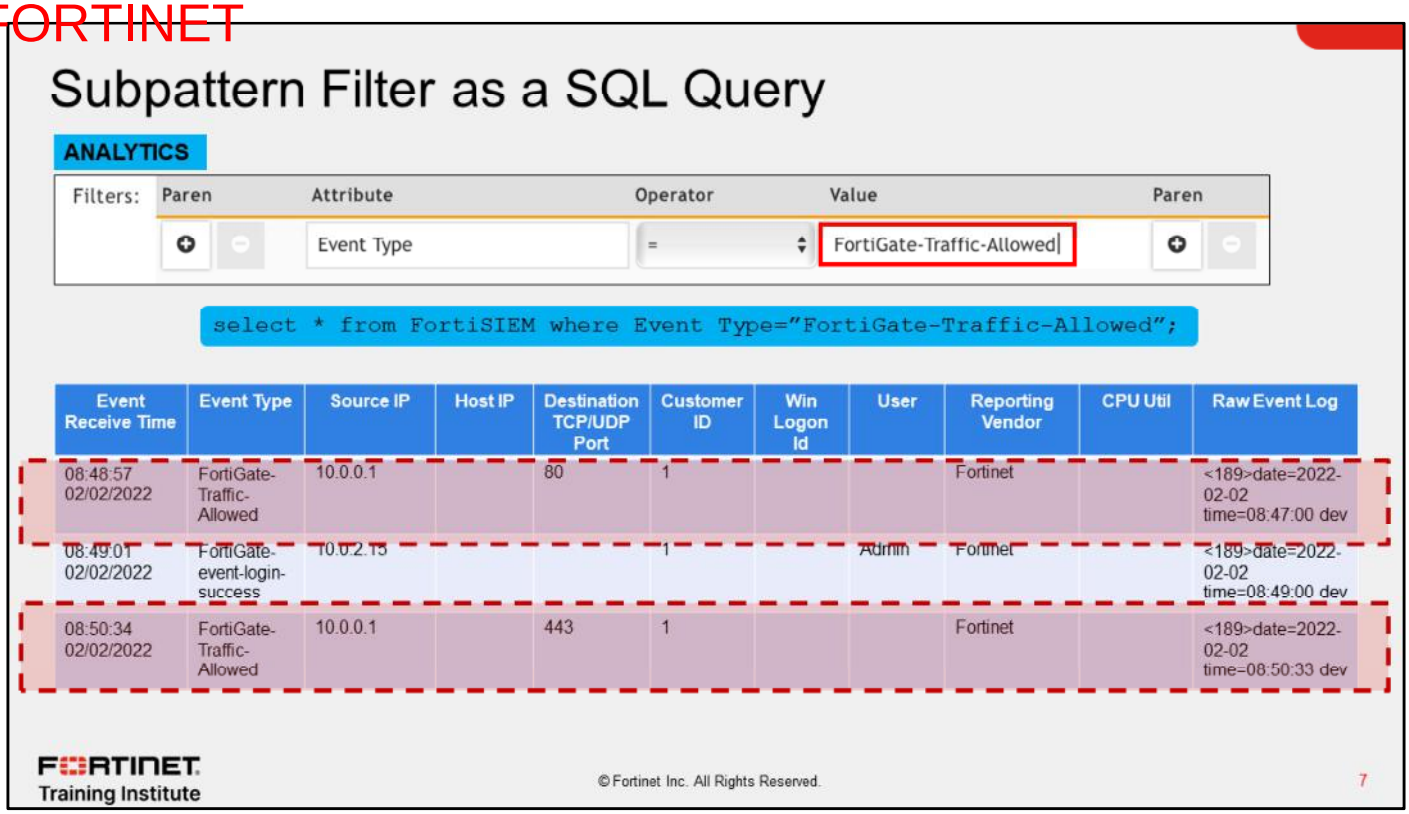
© Fortinet Inc. All Rights Reserved.

6

A subpattern defines the characteristics of events that will cause a rule to trigger an incident. A subpattern involves defining event attributes that will be monitored, and then defining the threshold values for aggregations of event attributes that will trigger an incident.

You can think of the subpattern filter just like a SQL query, where the whole FortiSIEM event database is a huge table with many columns that are not all populated with data because every event contains different values.

DO NOT REPRINT
© FORTINET



The parameters that you specify in the event filter eventually translate to a SQL query. This slide shows multiple columns with different event attributes and values. If a certain attribute does not apply to an event and there is no value for that attribute, then it will be populated with a NULL value.

In the **Event Filter**, if you select **Event Type** as the attribute and enter the value shown on this slide, FortiSIEM translates that into a SQL statement to query the table. It will filter out other events and evaluate those events where traffic was allowed by FortiGate.

DO NOT REPRINT
© FORTINET

Subpattern Filter—SQL Example Concept

ANALYTICS

Filters: | Paren | Attribute | Operator | Value | Paren |
| | Event Type | = | PH_DEV_MON_SYS_CPU_UTIL | |

```
select * from FortiSIEM where Event Type="PH_DEV_MON_SYS_CPU_UTIL";
```

Event Receive Time	Event Type	Source IP	Host IP	Destination TCP/UDP Port	Customer ID	Win Logon Id	User	Reporting Vendor	CPU Util	Raw Event Log
08:48:57 02/02/2022	FortiGate-Traffic-Allowed	10.0.0.1		80	1			Fortinet		<189>date=2022-02-02 time=08:47:00 dev
08:50:31 02/02/2022	PH_DEV_MON_SYS_CPU_UTIL		10.0.5.15		1			Microsoft	8.70	[PH_DEV_MON_SYS_CPU_UTIL]:[eventSeverity]=P
08:50:34 02/02/2022	FortiGate-Traffic-Allowed	10.0.0.1		443	1			Fortinet		<189>date=2022-02-02 time=08:50:33 dev

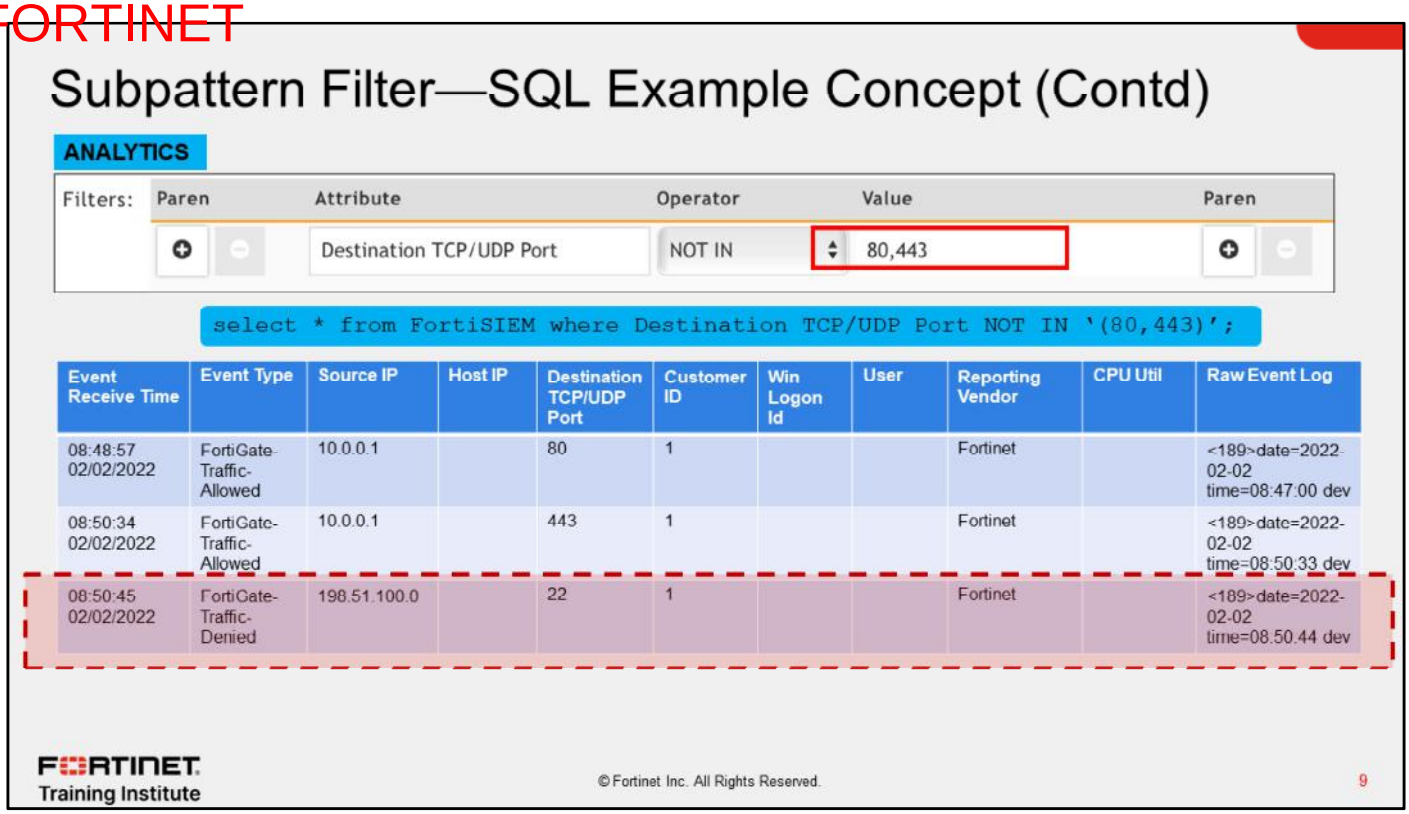
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

8

This slide shows another performance monitoring event type related to CPU utilization. FortiSIEM will translate the filter into a SQL query, which is shown on this slide. It will filter out other events and evaluate only those events that are related to CPU utilization.

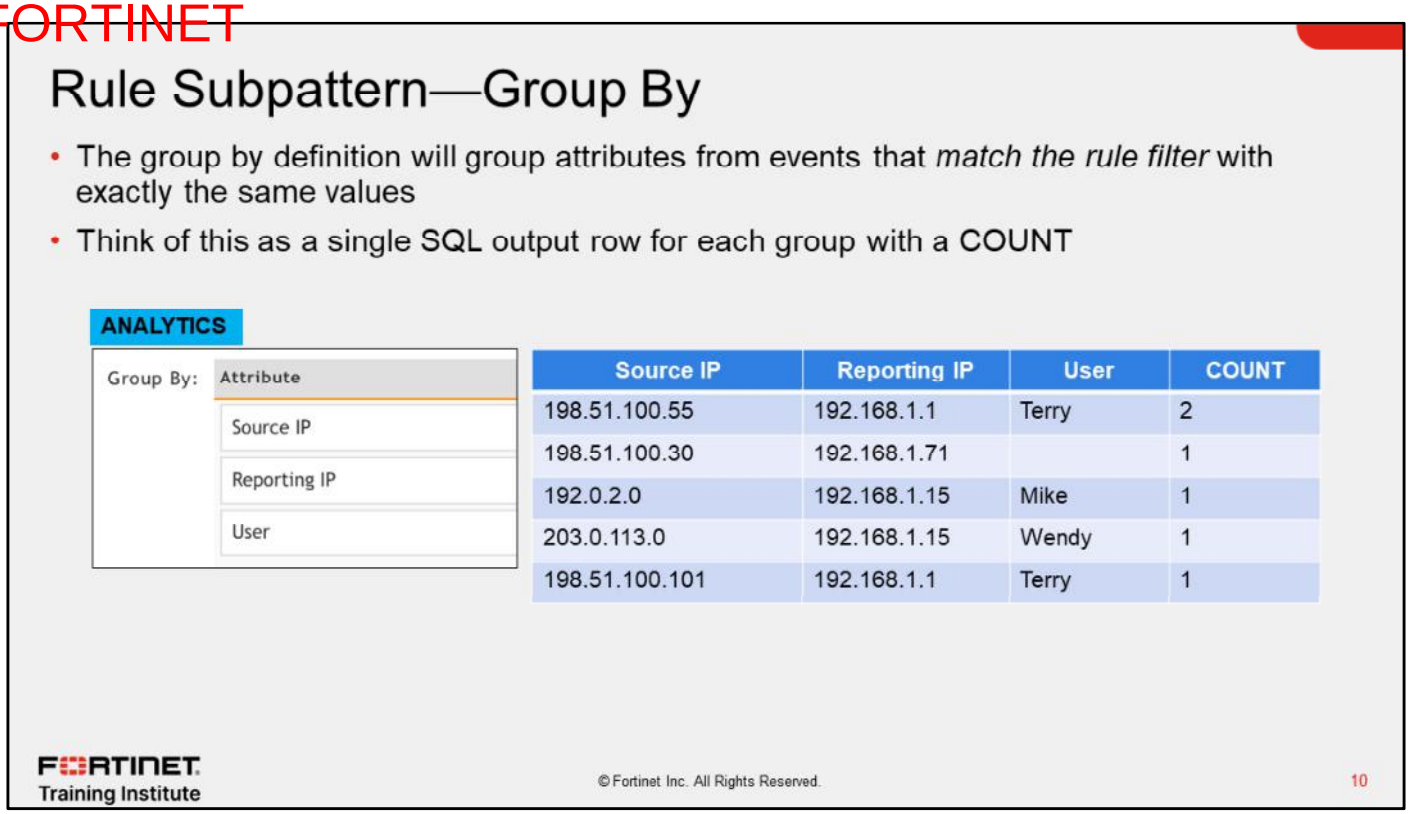
DO NOT REPRINT
© FORTINET



This slide shows another example using the **NOT IN** operator with the **Attribute** set as **Destination TCP/UDP Port** with a **Value** of `80,443`. The resulting SQL statement is shown on this slide.

This query will return all events whose destination TCP and UDP port is *not* 80 or 443, which means all events whose destination port is not related to HTTP or HTTPS traffic.

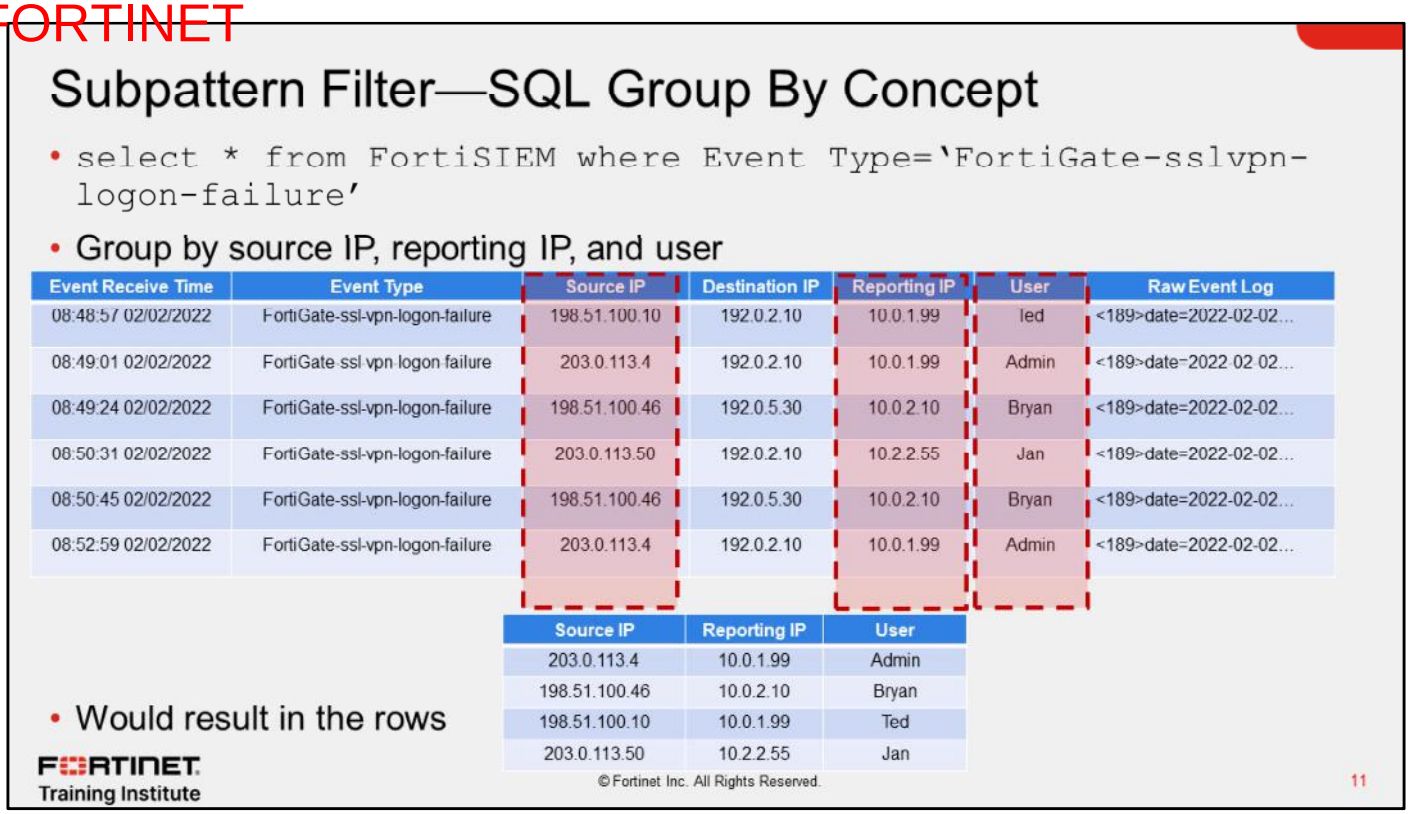
DO NOT REPRINT
© FORTINET



The rule filter will filter out events that are of interest for evaluating a rule, and group by will group all those events based on the attributes that you define in the **Group By** field.

In this example, user Terry was found in three events but, from those three events, two of the events match the group by attribute with a common source IP, reporting IP, and user. Think of this as a single SQL output row for each group with a COUNT.

DO NOT REPRINT
© FORTINET

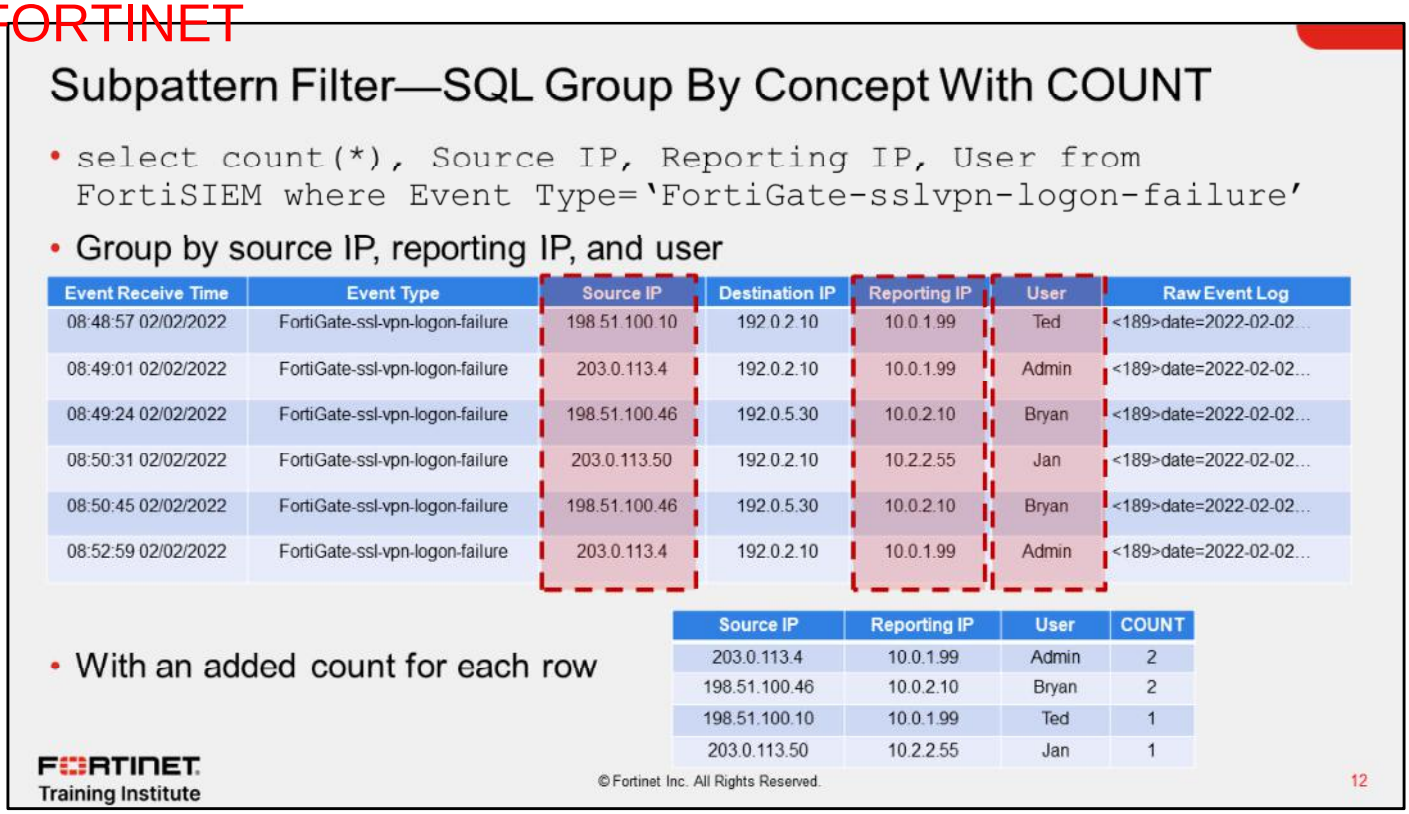


This slide shows an example event type filter for FortiGate SSL VPN logon failures. The group by attributes are set as source IP, reporting IP and user.

In the example shown on this slide, the user Admin has two entries with the same source IP and reporting IP. So, based on the group by definition, those two entries will be grouped as one.

Similarly, user Bryan has two entries with the same source IP and reporting IP. So, those two entries will also be grouped as one.

DO NOT REPRINT
© FORTINET



This slide shows the same event filter example, but this time there is an additional count column. The filter is still searching for FortiGate SSL VPN logon failures, and the group by attributes are set as source IP, reporting IP, and user.

The count column will display the number of times the event was repeated for the same source IP, reporting IP, and user.

In the results, the user Bryan had SSL VPN logon failures from the same source IP and it was reported twice, which you can see in the corresponding COUNT column.

DO NOT REPRINT
© FORTINET

Rule Subpattern—Aggregate Condition

- The aggregate condition is used to perform calculations on the matching data and return a single result
- Typical SQL aggregate functions are supported such as:

Function	Description
COUNT(Matched Events)	The row count of the matching events in the group by condition
COUNT DISTINCT(attribute)	Counts the number of distinct attribute values
SUM(attribute)	Sum of the numerical values of the specified attribute
AVG(attribute)	Returns the average value of the specified attribute
MIN(attribute)	Returns the smallest value
MAX(attribute)	Returns the largest value
FIRST(attribute)	Returns the first value
LAST(attribute)	Returns the last value

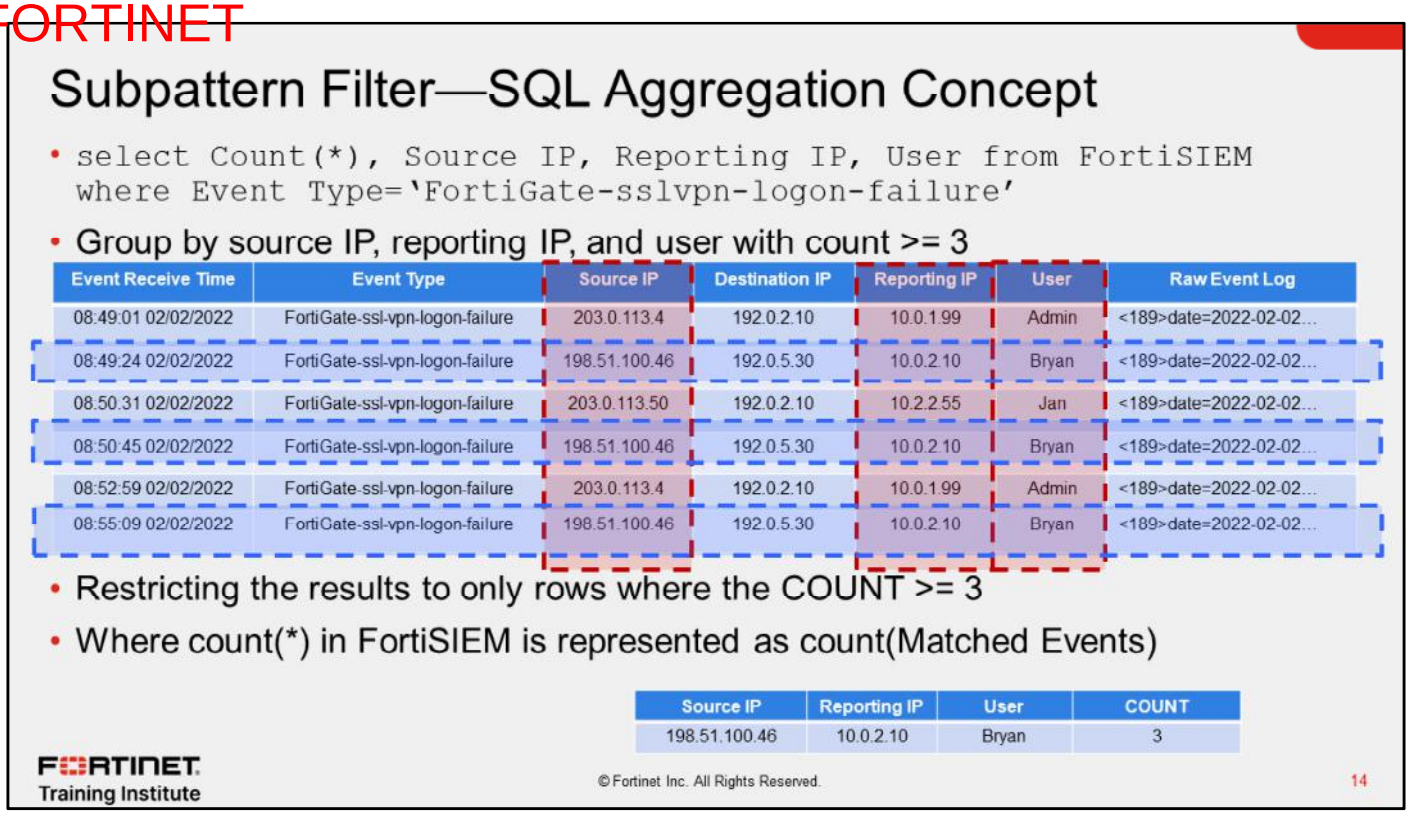
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

13

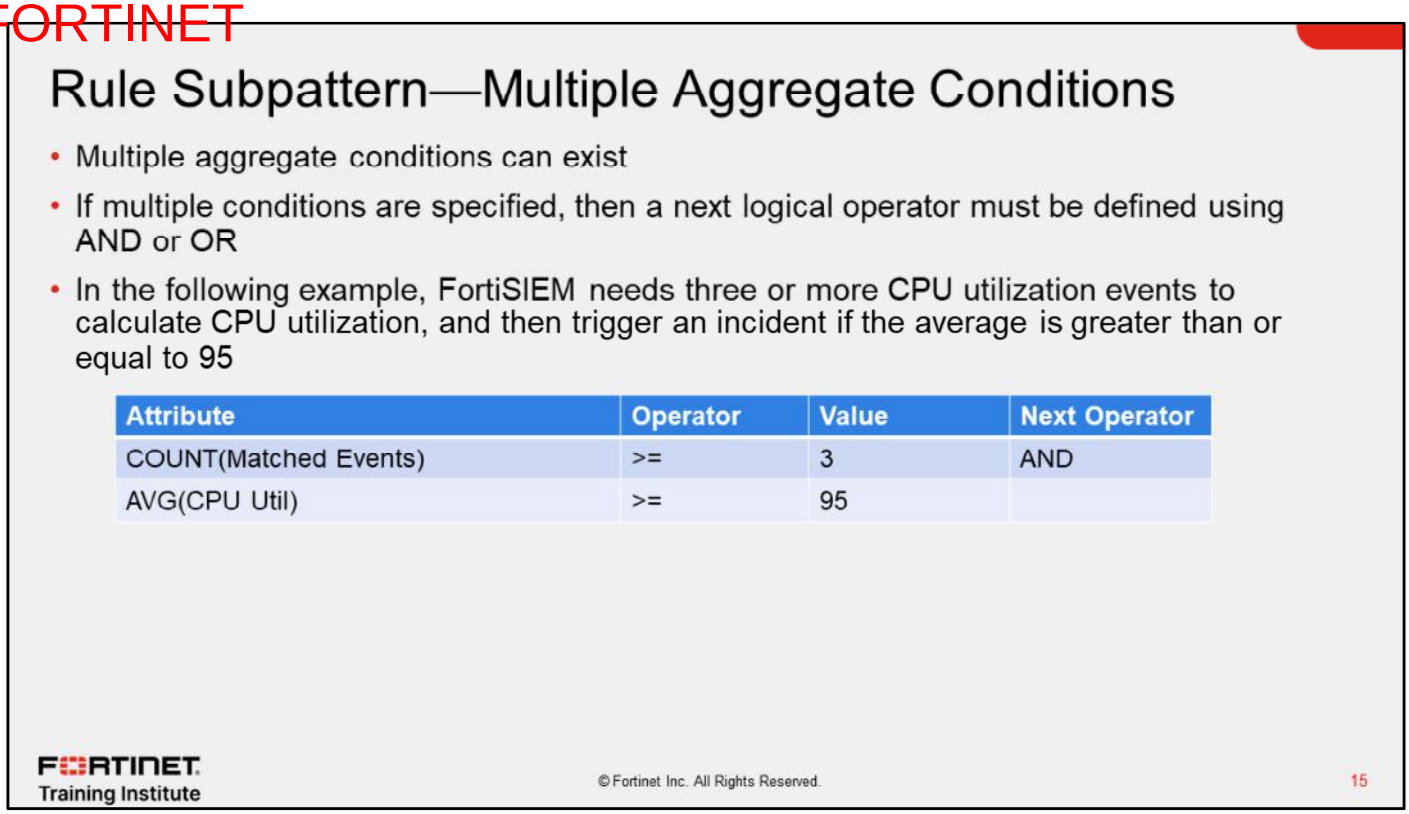
The aggregate condition is used to perform calculations on the matching data and return a single result. The typical SQL aggregate functions are supported. Some of these supported functions are shown on this slide.

DO NOT REPRINT
© FORTINET



Take a look at the same event filter example shown on this slide. This time, the group by condition is configured to restrict the number of results, which means that only events with same source IP, reporting IP, and user, as long as the count is equal to or greater than 3, will be evaluated by the rule.

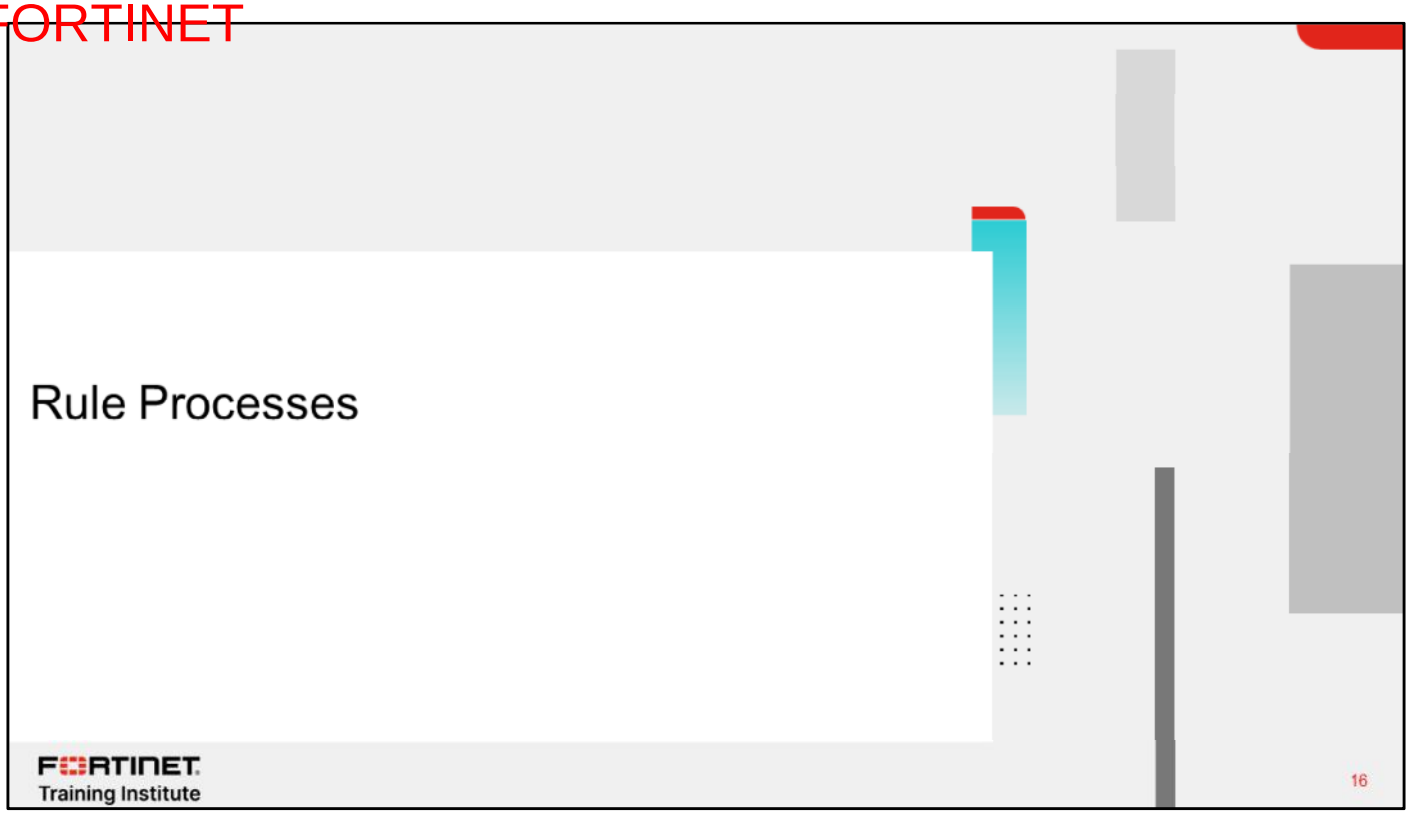
DO NOT REPRINT
© FORTINET



You can also have multiple aggregate conditions using the next operator. If multiple conditions are specified, then a next logical operator must be defined using AND or OR.

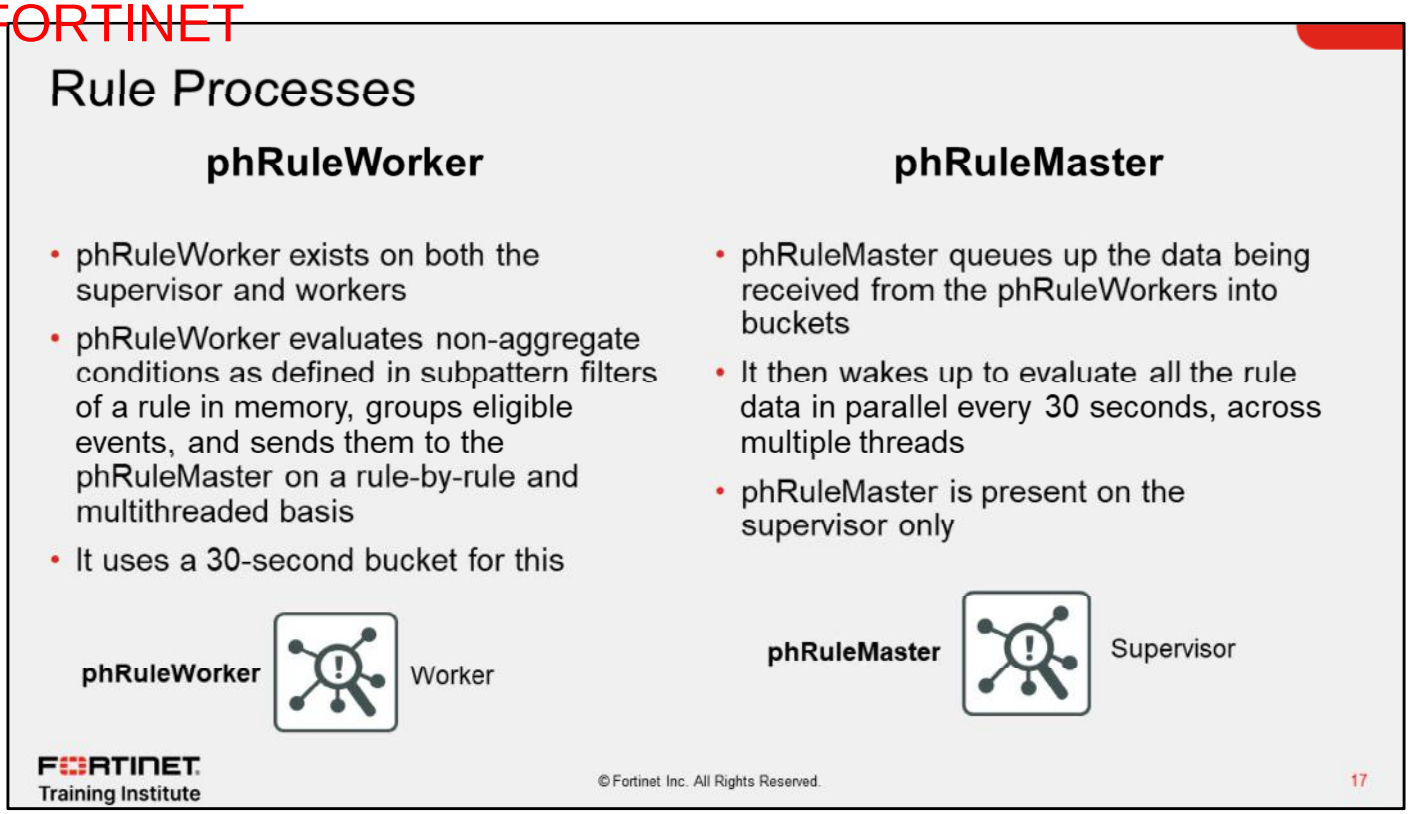
In the example table shown on this slide, FortiSIEM needs three or more CPU utilization events to calculate average CPU utilization, and then trigger an incident if the average is greater than or equal to 95. This compound rule is built using the AND operator.

DO NOT REPRINT
© FORTINET



In this section, you will learn about rule processes and rule process architecture.

DO NOT REPRINT
© FORTINET

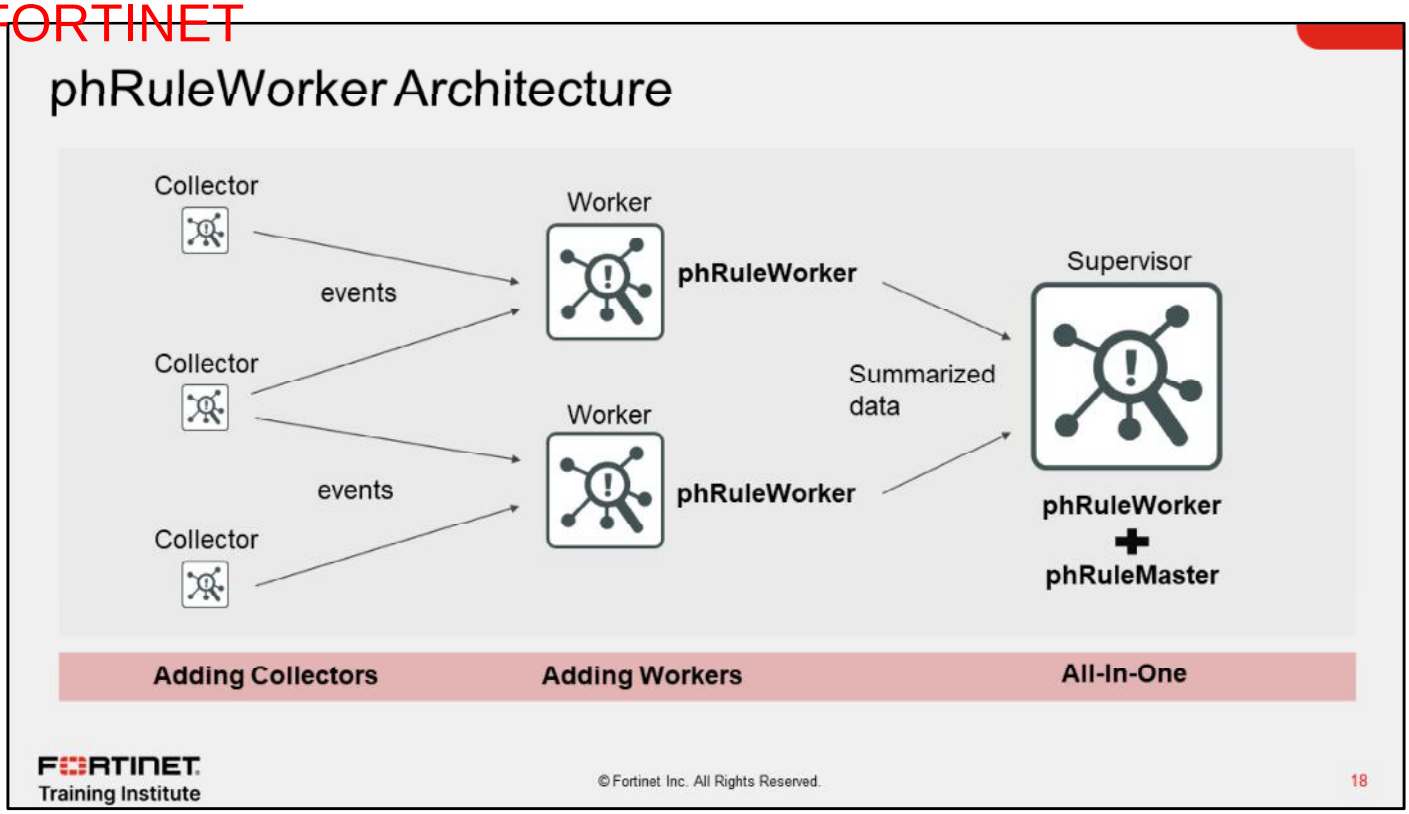


All FortiSIEM rules processing is performed in memory. There are two backend processes involved—phRuleWorker and phRuleMaster. There can be many phRuleWorker processes but only one phRuleMaster in a FortiSIEM deployment. This is how FortiSIEM is able to correlate events from multiple devices and locations no matter where the data was actually received.

The phRuleWorker exists on both the supervisor and workers. The phRuleWorker process evaluates non-aggregate conditions as defined in the subpattern filters of a rule in memory, groups eligible events, and sends them to the phRuleMaster on a rule-by-rule and multithreaded basis. It uses a 30-second bucket for this.

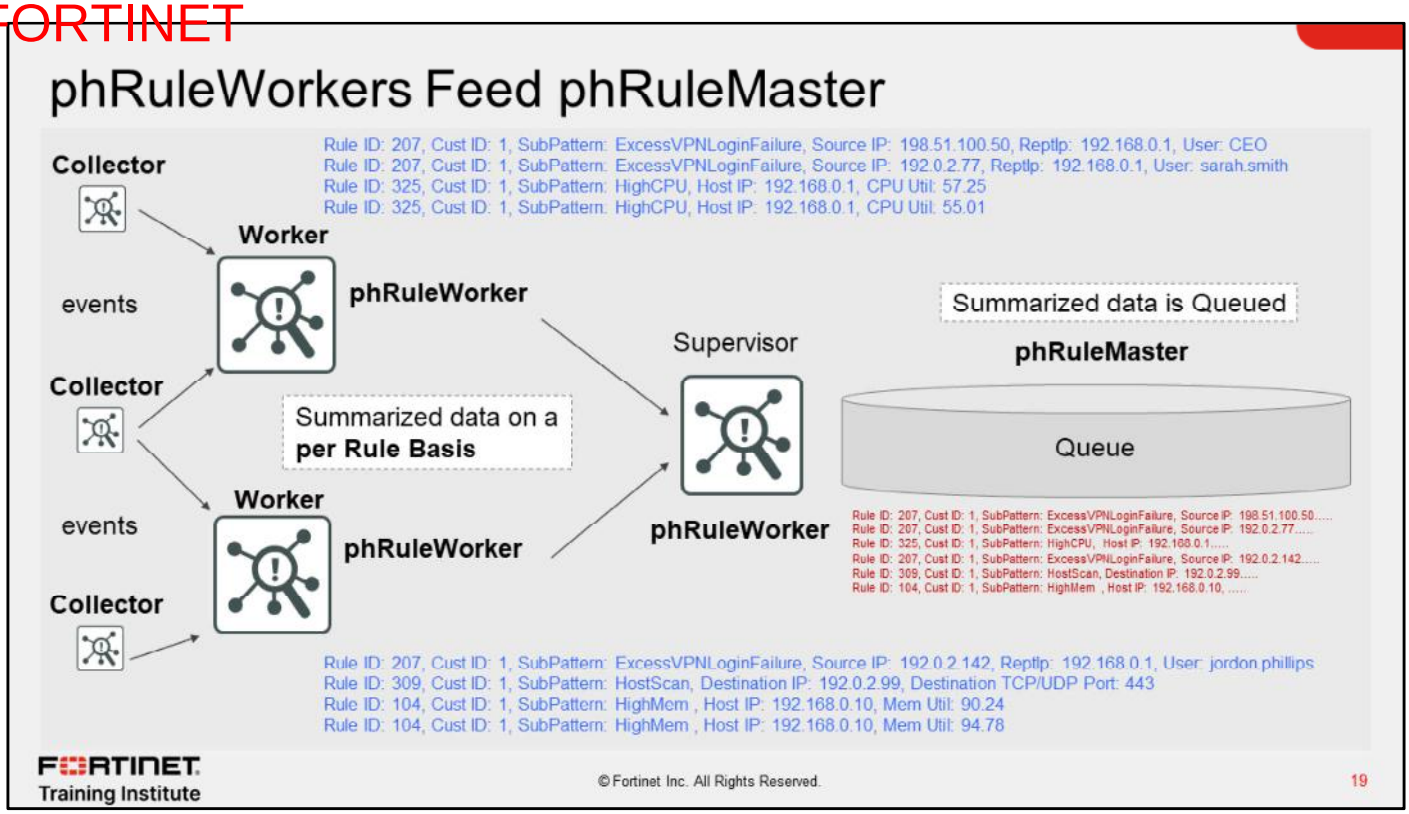
The phRuleMaster process queues up the data being received from the phRuleWorkers into buckets. It then wakes up to evaluate all the rule data in parallel every 30 seconds, across multiple threads. The phRuleMaster process is present on the supervisor only.

DO NOT REPRINT
© FORTINET



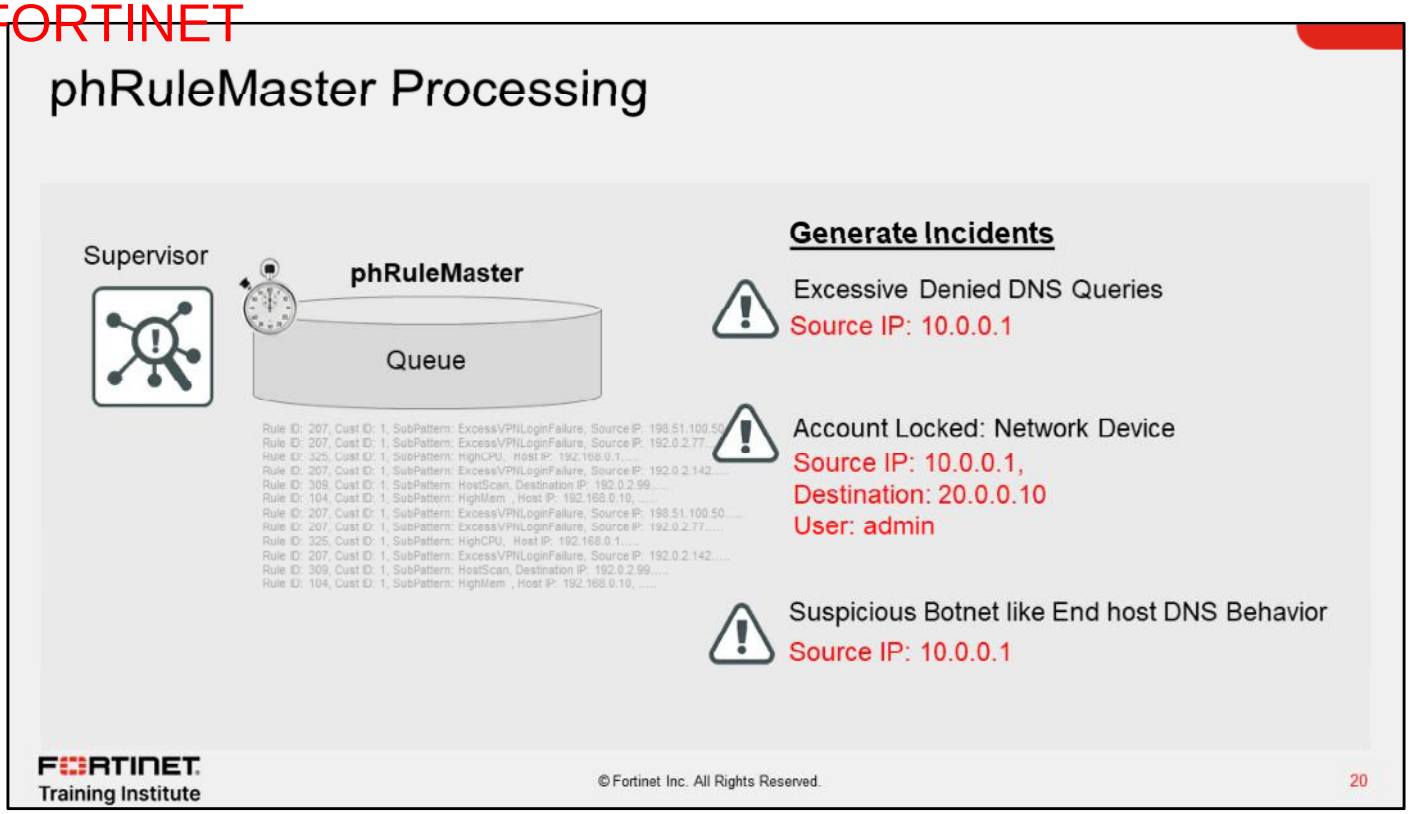
Scaling out the rule worker would require adding additional workers. The collectors will send events to the workers, and the phRuleWorker process on an individual worker will forward those events in a summarized form to the supervisor node, which also has a phRuleWorker. The task of phRuleWorker on the supervisor is to summarize events that are sent from organizations directly to the supervisor node and assets that belong to the super organization.

DO NOT REPRINT
© FORTINET



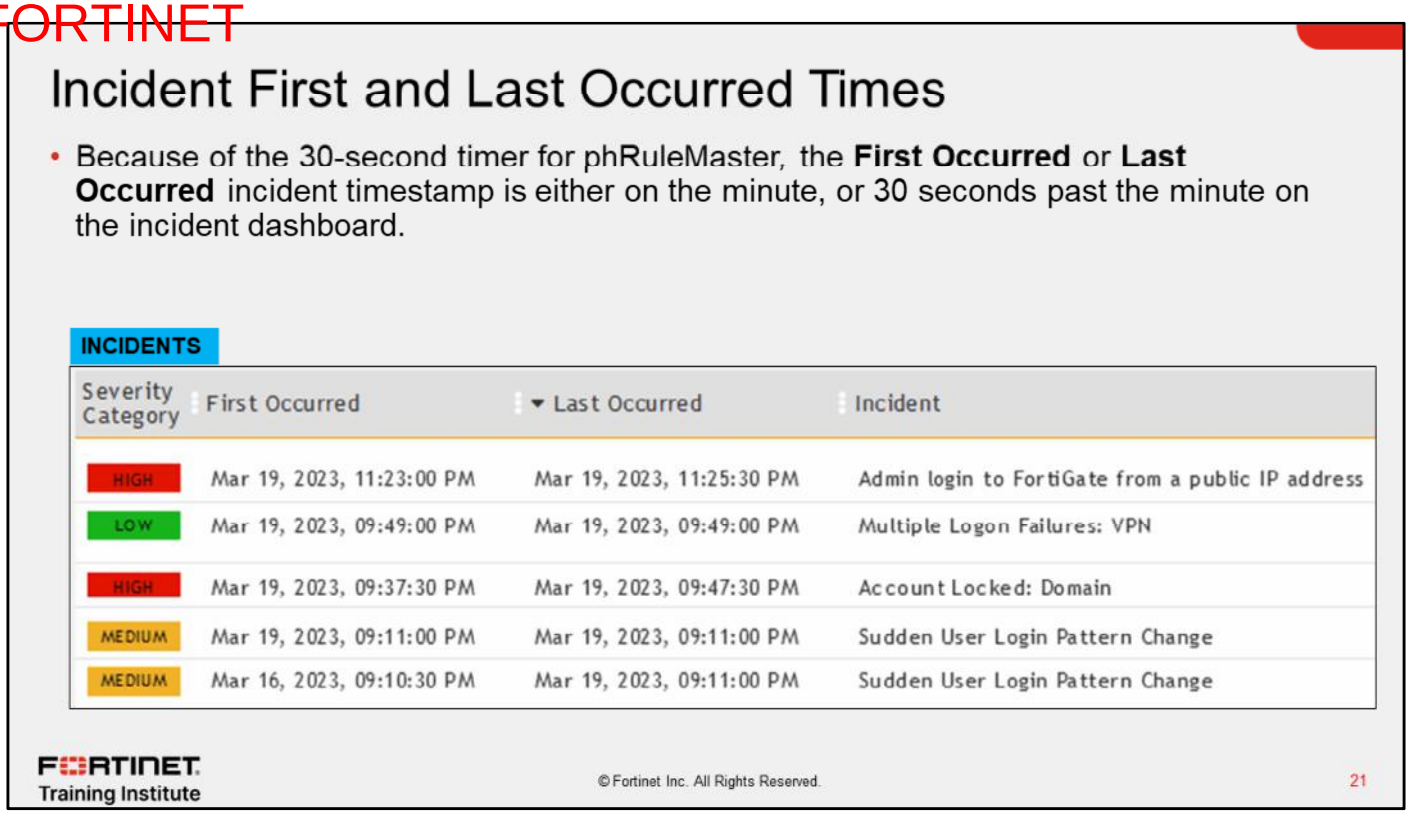
The phRuleWorkers feed summarizes data based on the event filter and group by conditions of the rule subpattern conditions on a per-rule basis to the supervisor node. The summarized data is queued by the supervisor node, waiting to be evaluated by the phRuleMaster.

DO NOT REPRINT
© FORTINET



The phRuleMaster wakes every 30 seconds to analyze the queue. It is the phRuleMaster that evaluates the aggregation condition of a rule. If the aggregate conditions calculation matches the value defined in the rule, then incidents will be generated.

DO NOT REPRINT
© FORTINET

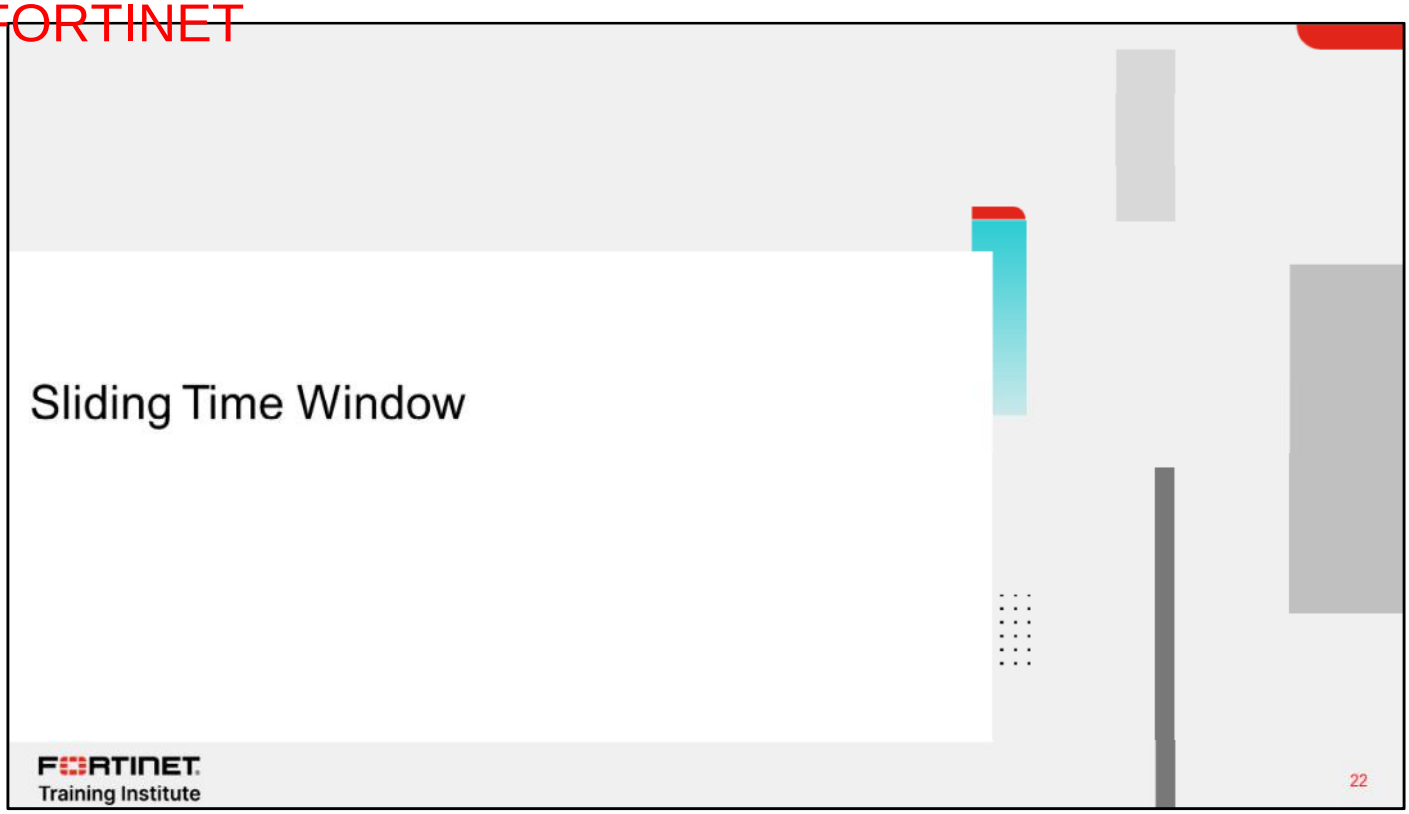


The 30-second timer for phRuleMaster is the reason why the **First Occurred** or **Last Occurred** incident timestamp is either on the minute, or 30 seconds past the minute on the incident dashboard.

Thirty seconds is an approximation because some rules may take longer to evaluate than others. Some are looking for occurrences of a single event, some are looking for occurrences of many events, and some are performing expensive calculations such as SUM across many thousands of matching events.

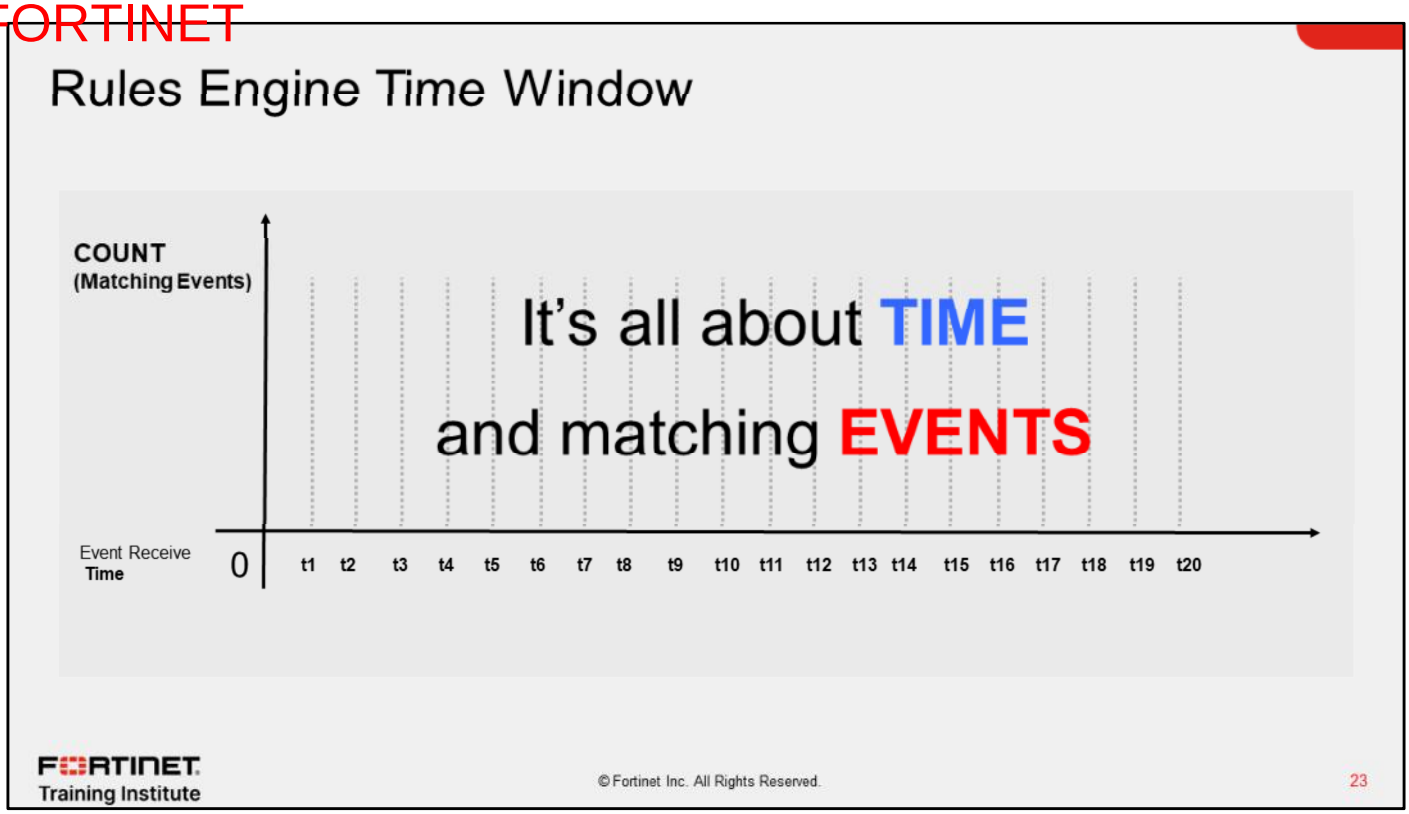
How much load is on the FortiSIEM is another factor. Every customer is different, and particular rules may or may not be in play in other customer environments. Even though the rule evaluation is multithreaded, events may spill across evaluation periods for complex rules.

DO NOT REPRINT
© FORTINET



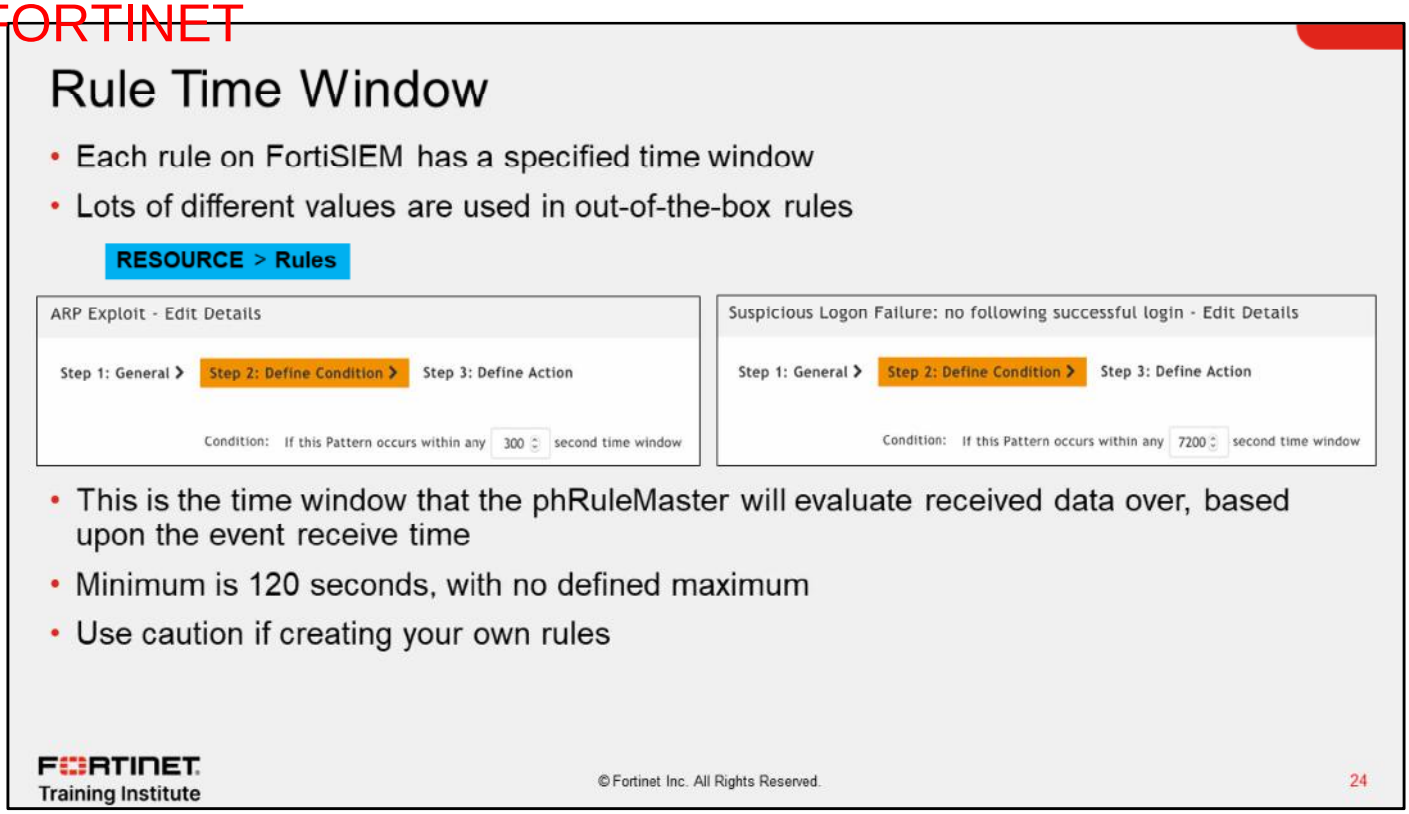
In this section, you will learn about the sliding time window that the rule engine uses to evaluate events.

DO NOT REPRINT
© FORTINET



Two main dimensions of the rule engine time window are time and events. Events are evaluated over a period of time and this general condition applies to the entire FortiSIEM. It's all about time and matching events. Time here does not refer to the timestamp of the event. It's the time when FortiSIEM received the event.

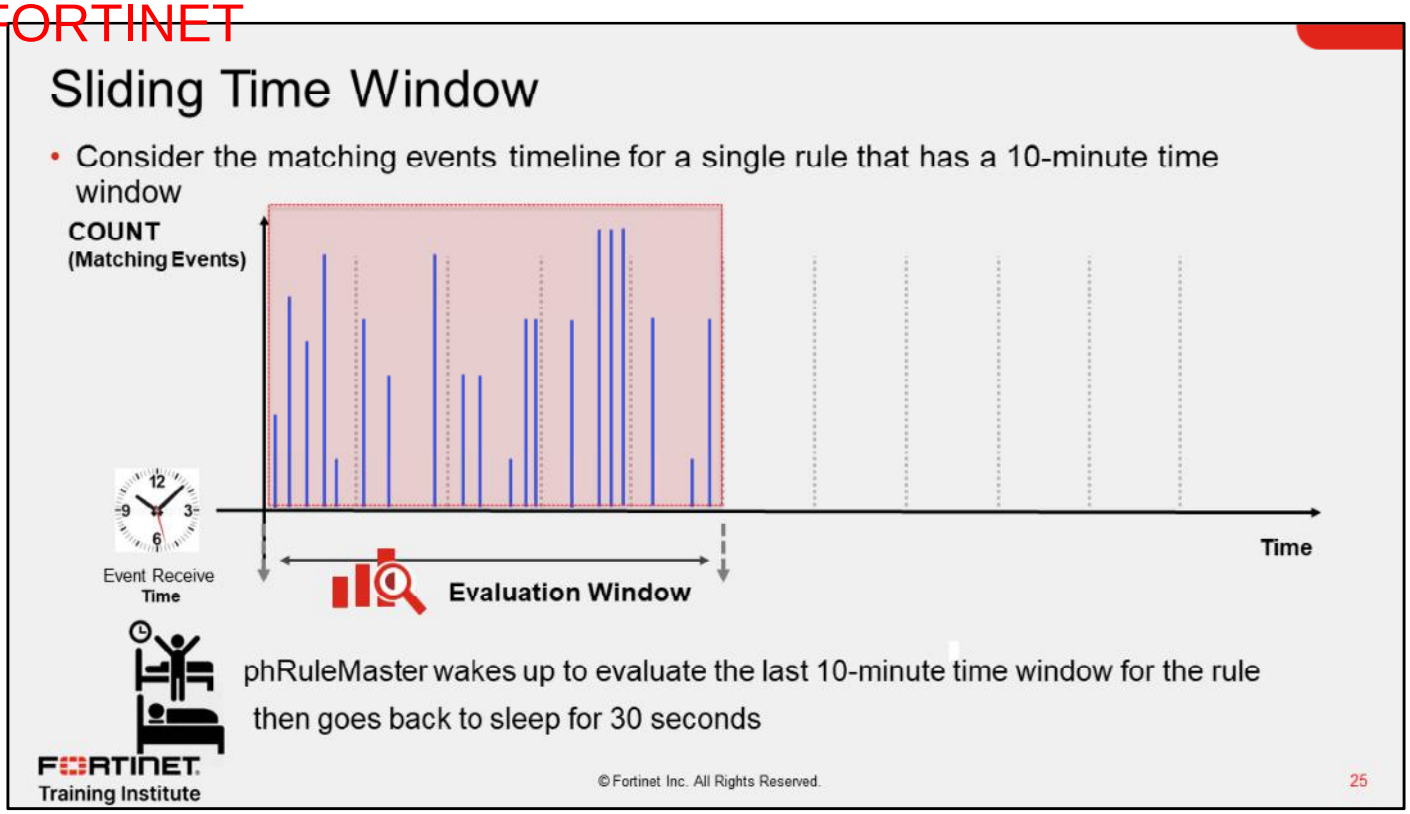
DO NOT REPRINT
© FORTINET



Each rule on FortiSIEM has a specified time window and different values are used in the out-of-the-box rules. This is the time window that the phRuleMaster will evaluate received data over, based upon the event receive time. The minimum value for the time window is 120 seconds, with no defined maximum.

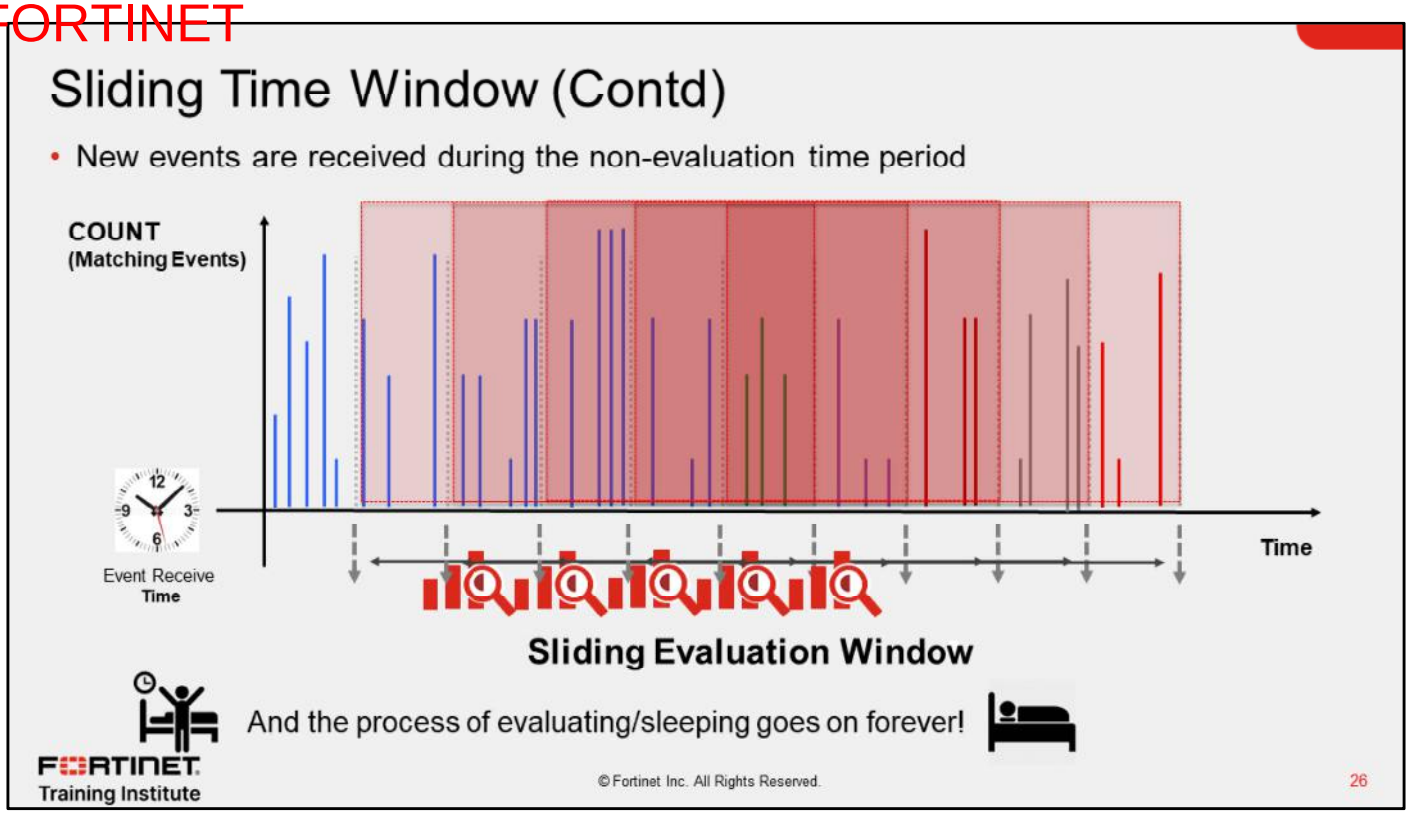
If you are creating your own rules, you must set the time window accordingly. You do not want the time window to be too large or, by the time an incident is generated for your rule it might be too late. At the same time, you should not make the time window too small because this will consume too much memory unnecessarily.

DO NOT REPRINT
© FORTINET



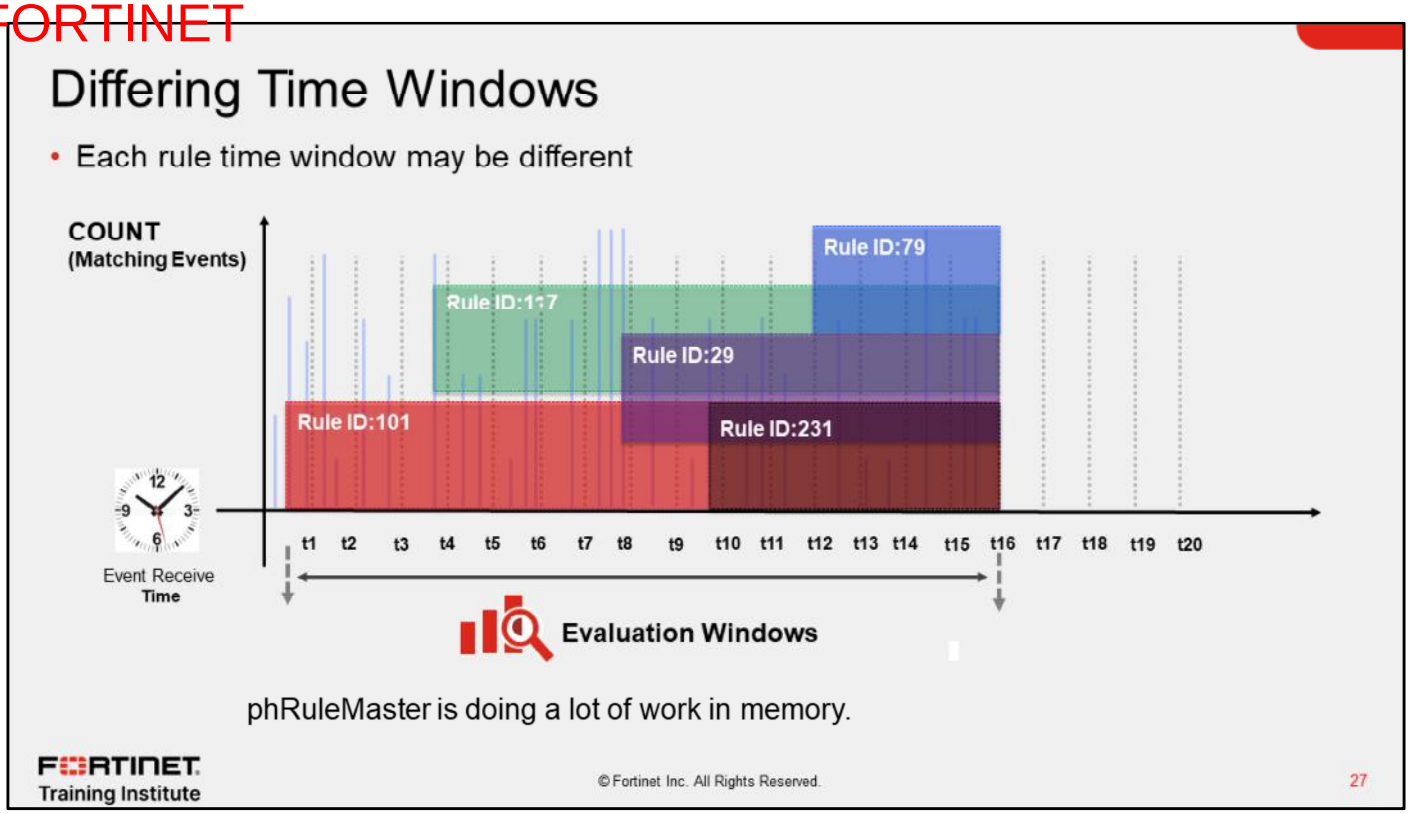
Consider the matching events timeline for a single rule that has a 10-minute time window. The phRuleMaster wakes up every 30 seconds to evaluate the last 10-minute time window for the rule.

DO NOT REPRINT
© FORTINET



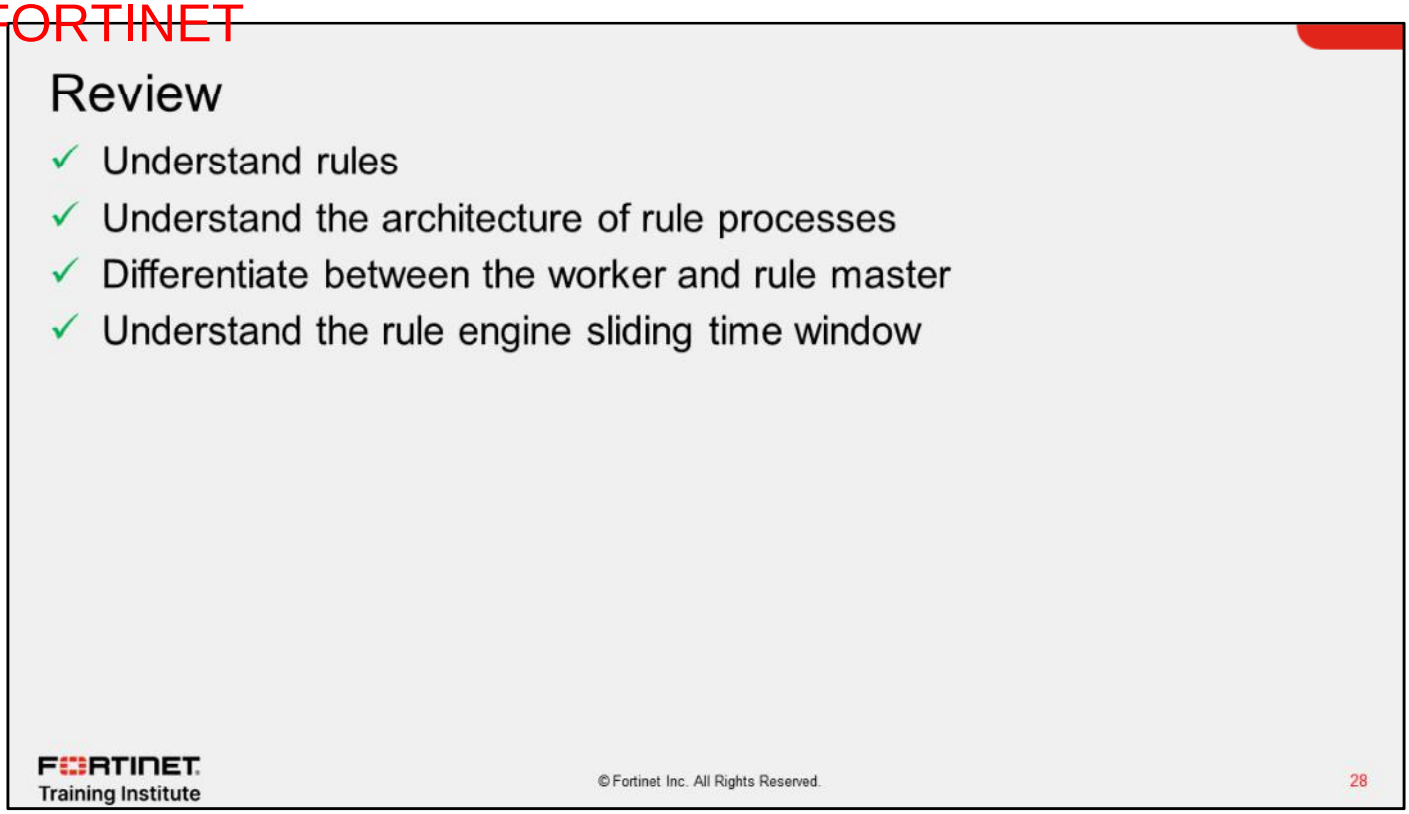
New events are received during the non-evaluation time period. When the phRuleMaster wakes up after 30 seconds it will evaluate all the events that were received during the non-evaluation period. The events that are outside the evaluation period, which have passed the evaluation sliding window, will not be considered in the aggregation calculation by the phRuleMaster.

DO NOT REPRINT
© FORTINET



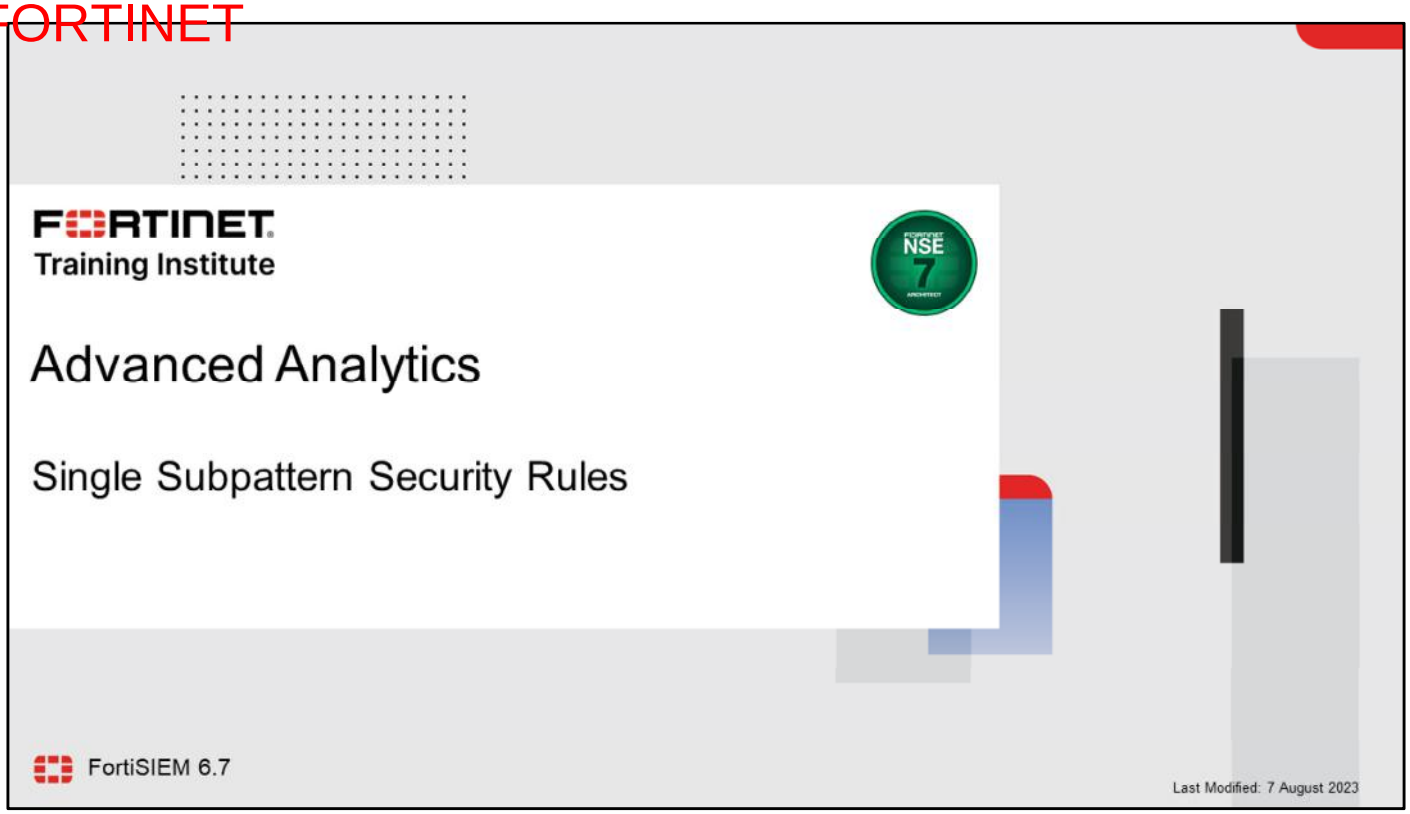
Every rule on FortiSIEM may have a different time window, which means that at any given time the phRuleMaster is evaluating multiple rules all at the same time over various time windows. Some rules might have a small-time window and fewer events to process while other rules might have a large window and many events to process. All these calculations and processing of events by the phRuleMaster are performed in the memory, hence it can be very resource-intensive.

DO NOT REPRINT
© FORTINET



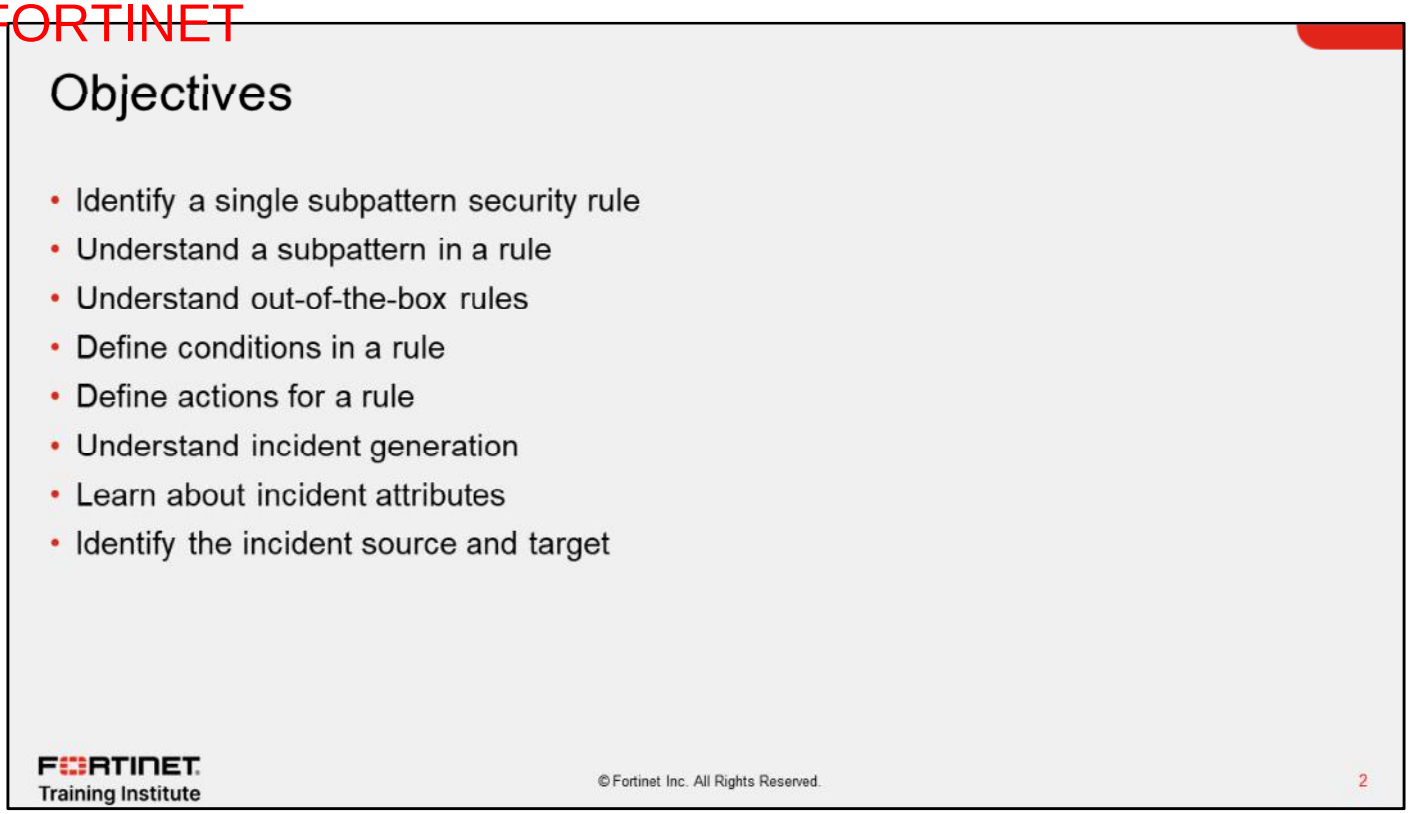
By mastering the objectives covered in this lesson, you learned about rules, including subpatterns in a rule and the back-end process that is involved in evaluating rules and triggering incidents. You also learned about the sliding time window in a rule, out-of-box rules, and how each rule time window may be different.

DO NOT REPRINT
© FORTINET



In this lesson, you will learn about the single subpattern security rule. You will also learn about incident generation and the attributes that are involved in incident generation.

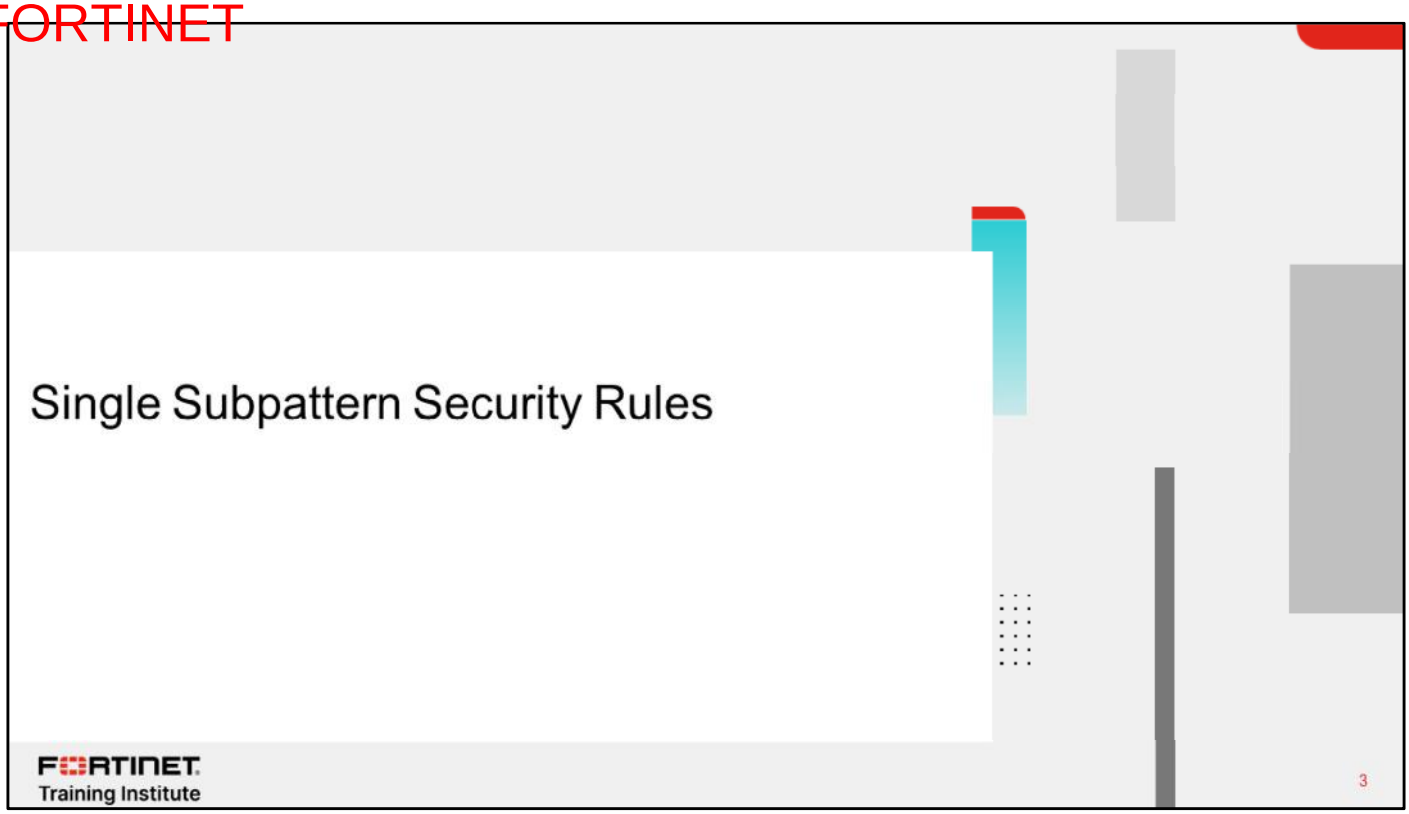
DO NOT REPRINT
© FORTINET



After completing this lesson, you should be able to achieve the objectives shown on this slide.

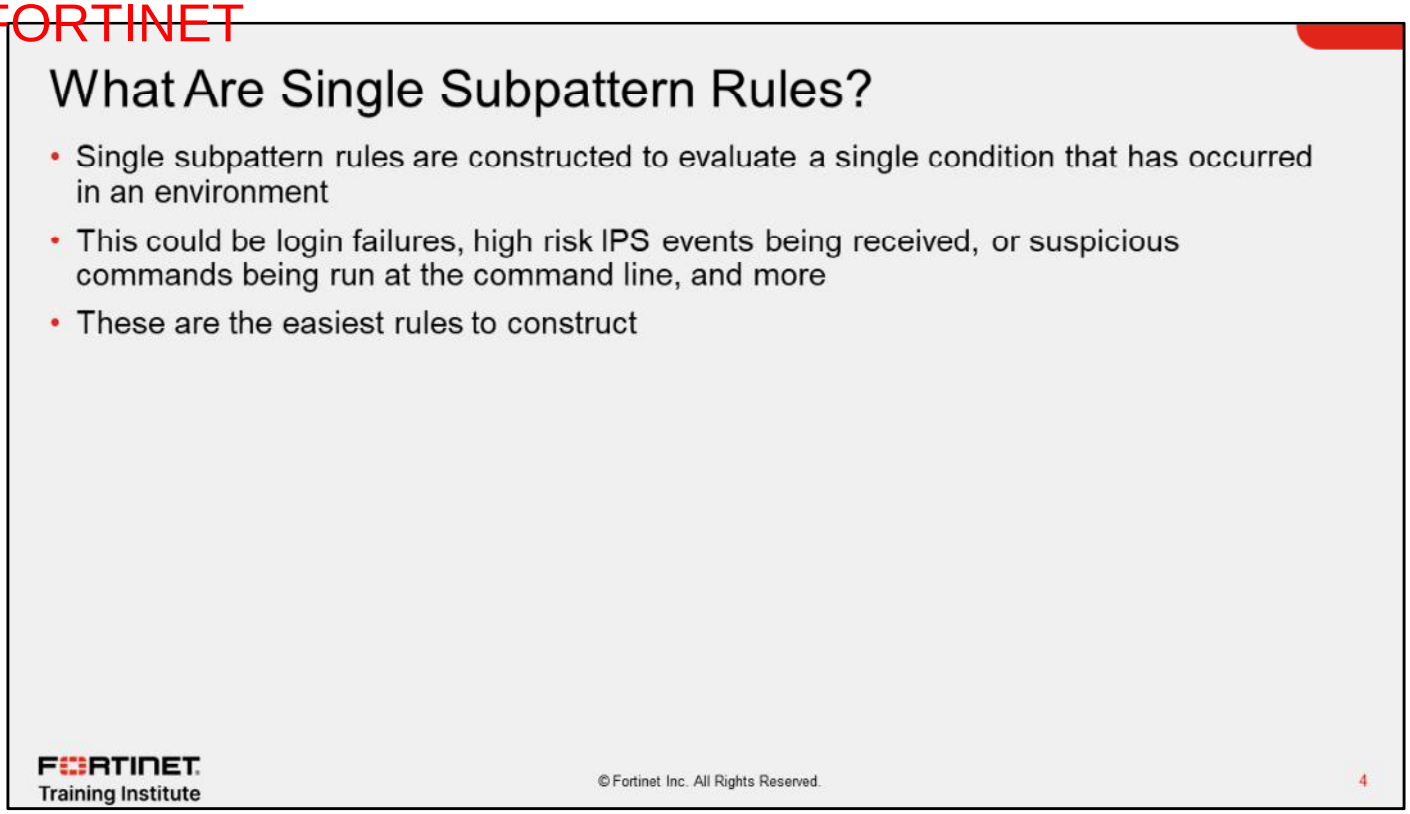
By demonstrating competence in the single subpattern security rule, you will be able to identify out-of-the-box security rules, which are single subpattern, and understand how incidents are generated by those rules.

DO NOT REPRINT
© FORTINET



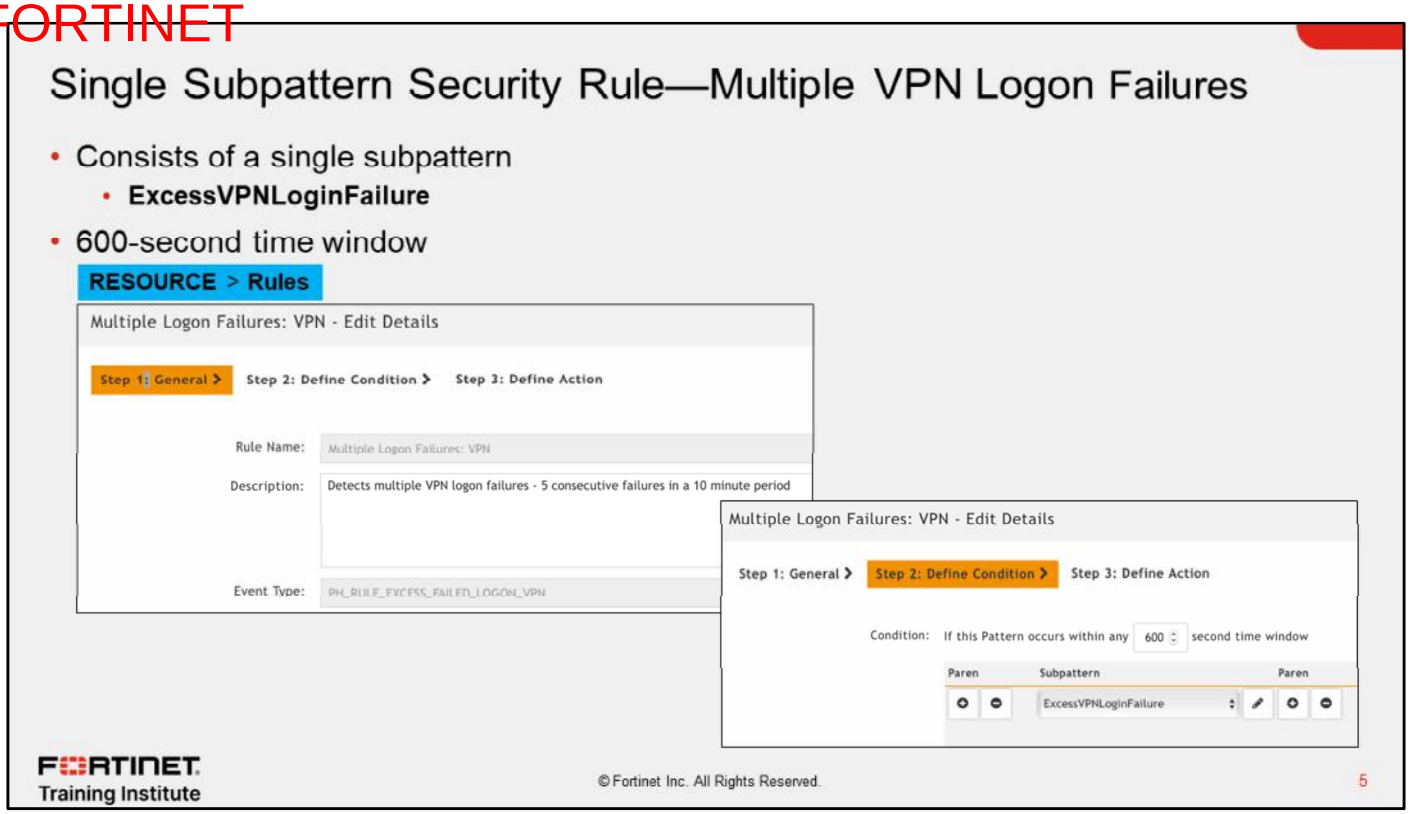
In this section, you will learn about event filters, group by, and aggregation conditions required in a single subpattern rule.

DO NOT REPRINT
© FORTINET



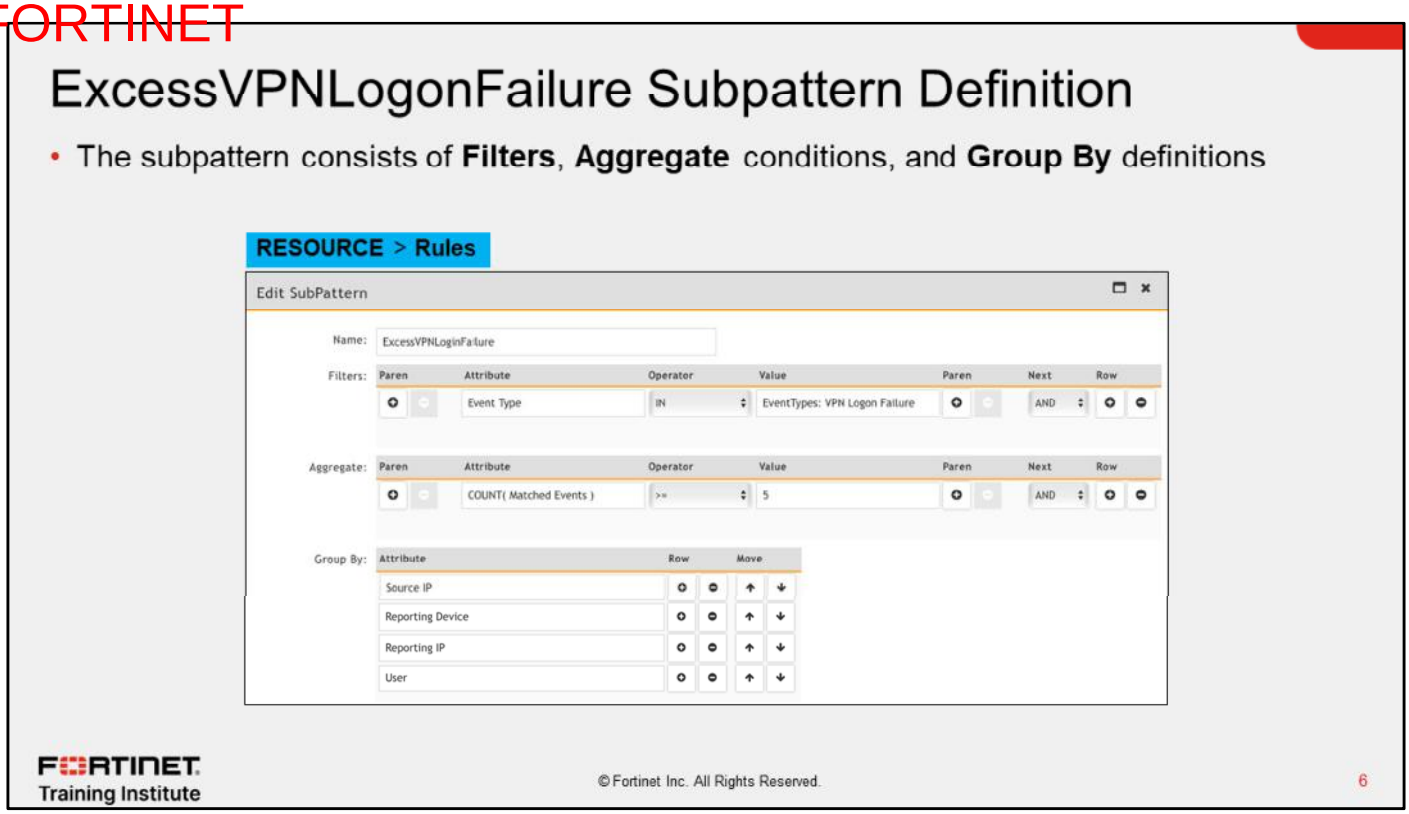
Single subpattern rules, as the name suggests, are constructed to evaluate a single condition that has occurred in an environment. This could be login failures, high risk IPS events being received, or suspicious commands being run at the command line, and more. These are the easiest rules to construct because, in most cases, you are simply configuring FortiSIEM to count the number of occurrences of a particular event being received during the time window.

DO NOT REPRINT
© FORTINET



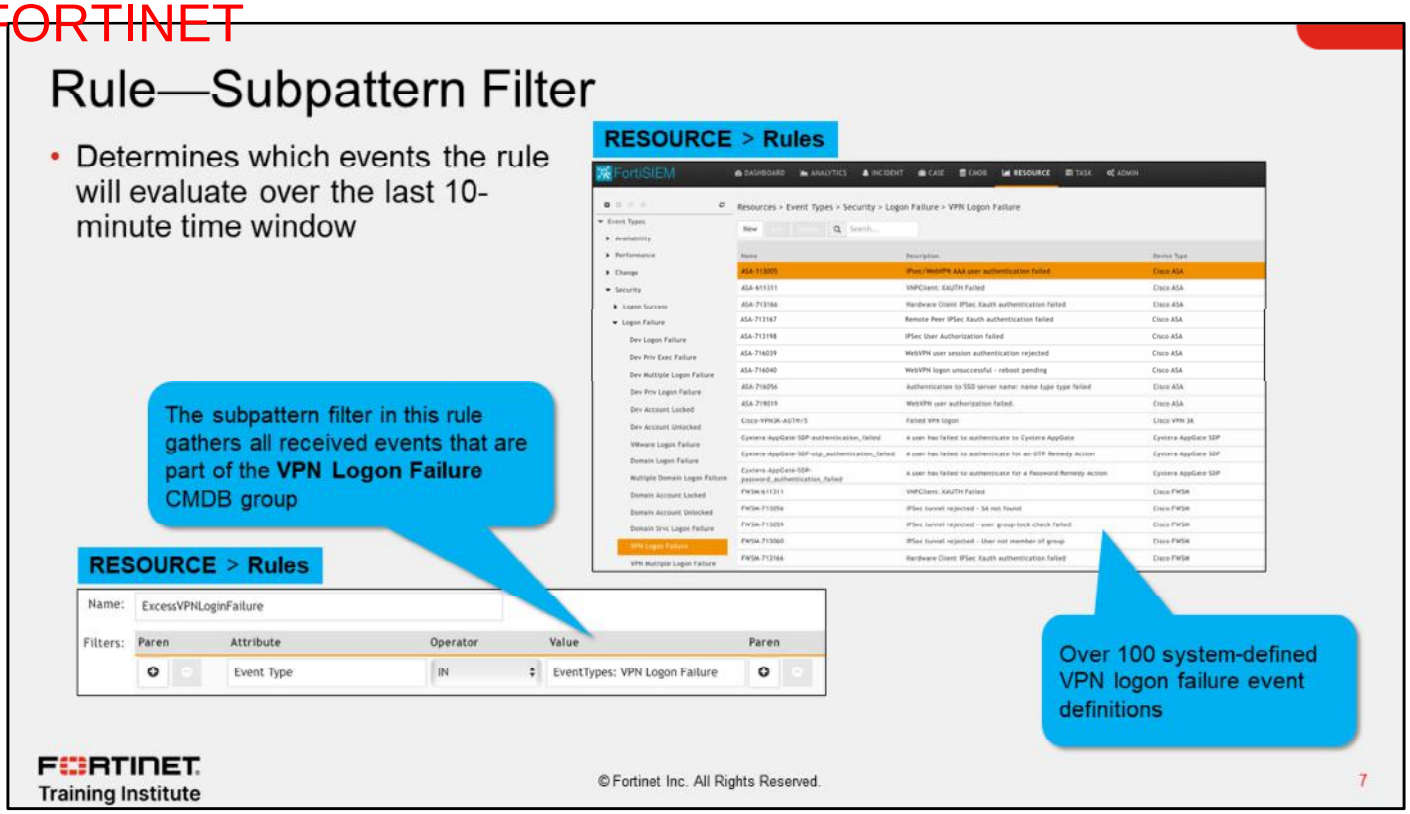
This rule is an example of multiple VPN logon failures, which detects five consecutive failures in a 10-minute period. This is an out-of-the-box rule that is active by default.

DO NOT REPRINT
© FORTINET



The **ExcessVPNLogonFailure** subpattern consists of **Filters**, **Aggregate** conditions, and **Group By** definitions.

DO NOT REPRINT
© FORTINET



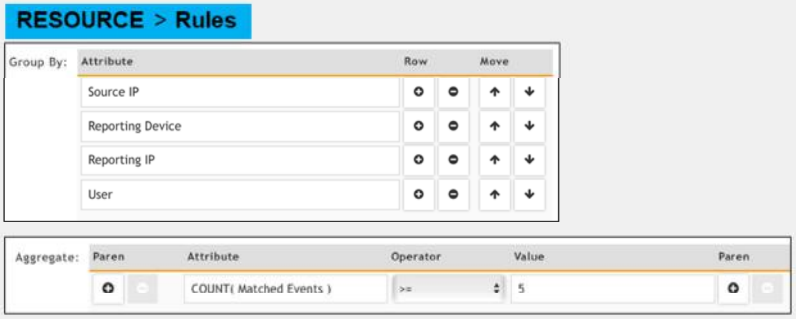
The subpattern filter determines which events the rule evaluates over the last 10-minute time window. It gathers all received events that are part of the **VPN Logon Failure** CMDB group. In the example shown on this slide, there are more than 100 different types of system-defined event definitions from various product types, such as Cisco ASA, FortiGate, and so on.

You cannot modify the out-of-the-box event types, but you can create your own VPN logon failure event type in the **VPN Logon Failure** CMDB group.

DO NOT REPRINT
© FORTINET

Rule Subpattern—Group By and Aggregation

- The **Group By** definition groups attributes from events that *match the rule filter* with the same values



- The **Aggregate** condition identifies a rule match when a resultant row from the query has five or more matching events within the sliding 10-minute time window
  - Same source IP, reporting device, reporting IP, and user
  - In SQL terms, a single row with a count of five or more

FORTINET Training Institute

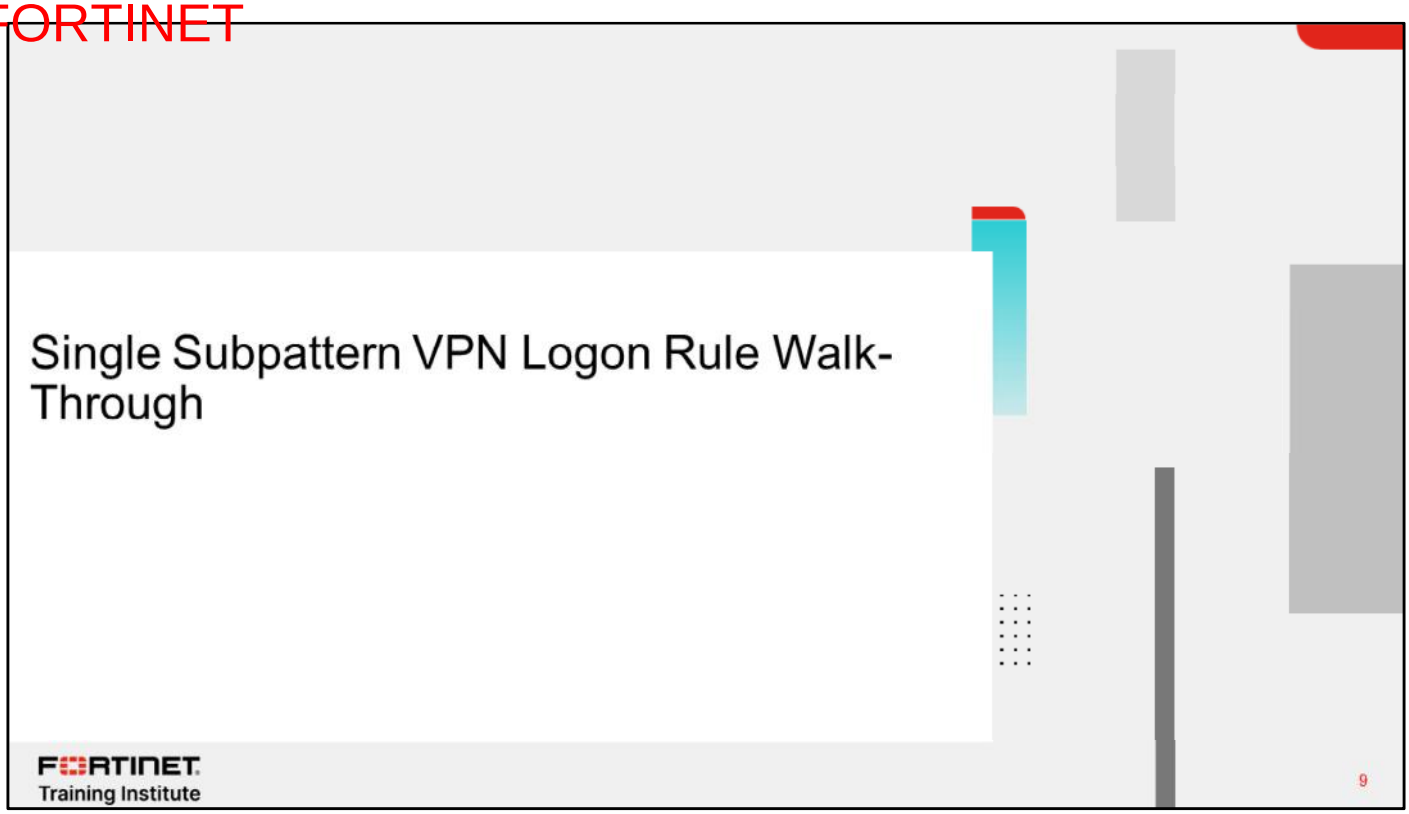
© Fortinet Inc. All Rights Reserved.

8

The **Group By** definition groups attributes from events that match the rule filter with exactly the same values. In the example shown on this slide, if multiple VPN logon failure events have the same source IP, reporting device, reporting IP, and user, they are grouped together in one row, and the count column tracks the number of events for each of those rows.

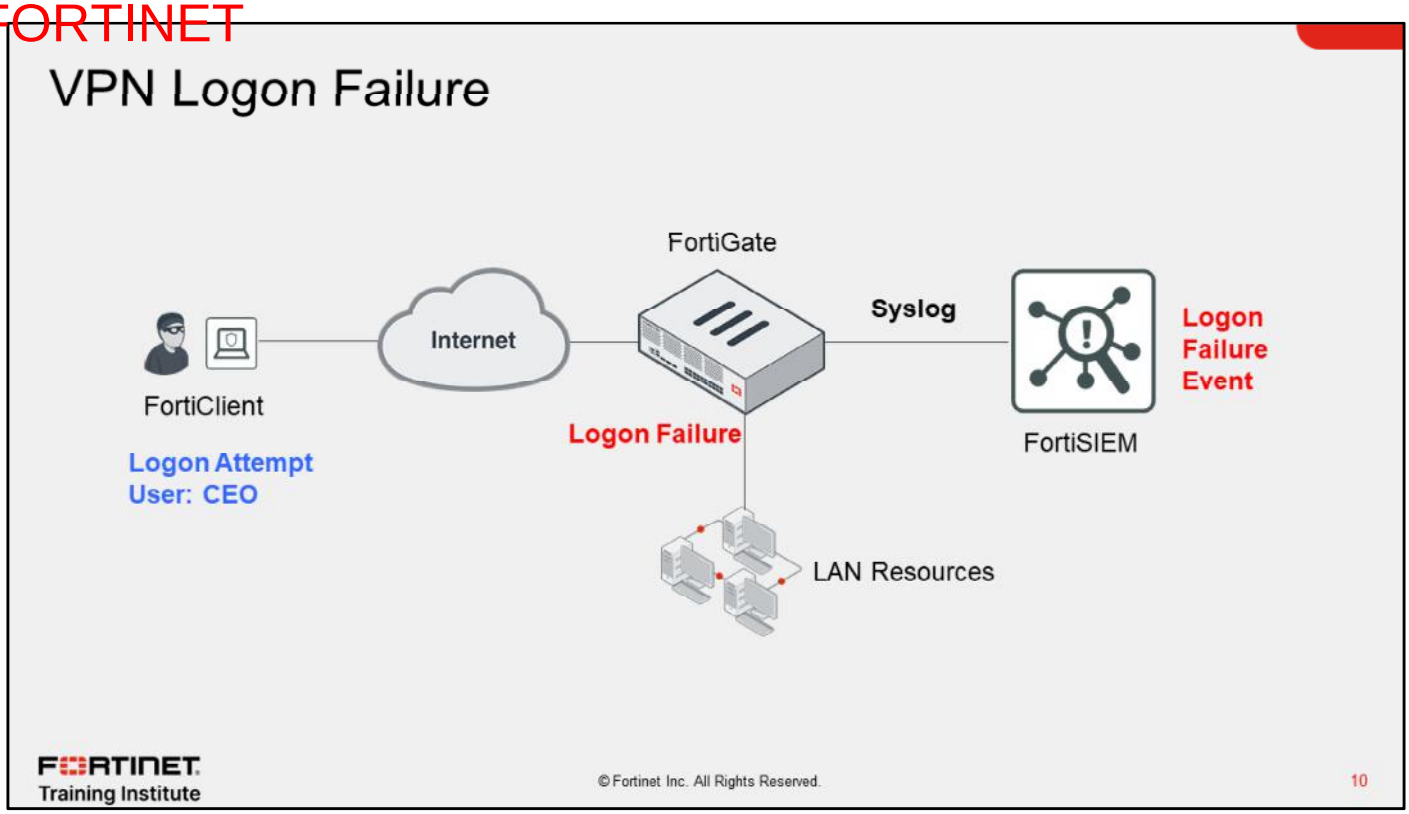
The **Aggregate** condition determines a rule match when a resultant row from the query has five or more matching events within the sliding 10-minute time window.

DO NOT REPRINT
© FORTINET



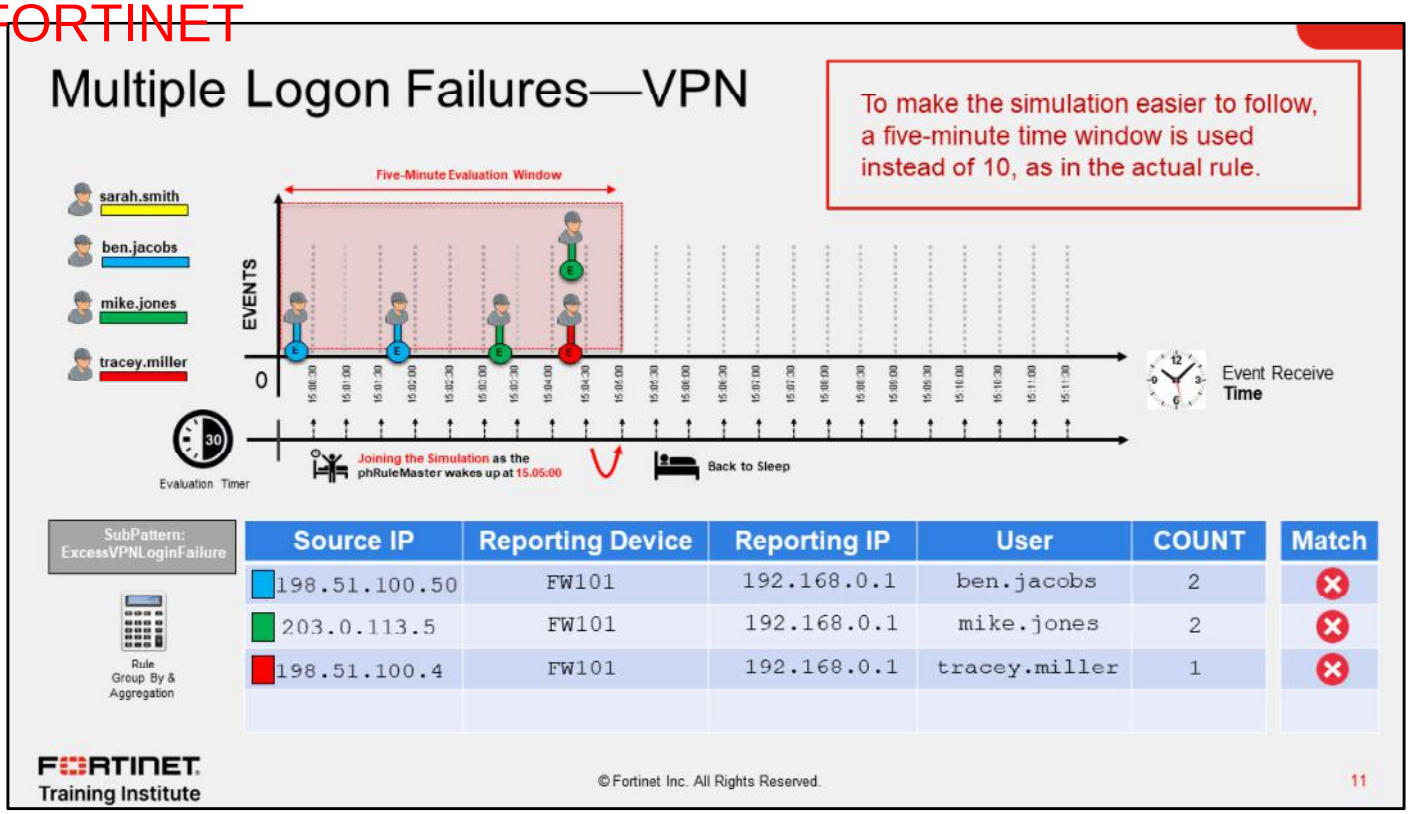
In this section, you will walk through a single subpattern VPN logon failure.

DO NOT REPRINT
© FORTINET



When someone enters incorrect credentials for a VPN logon, a logon failure event is sent to FortiSIEM. If the number of logon failures within a specified time period is higher than the parameters configured in the rule, an incident is triggered.

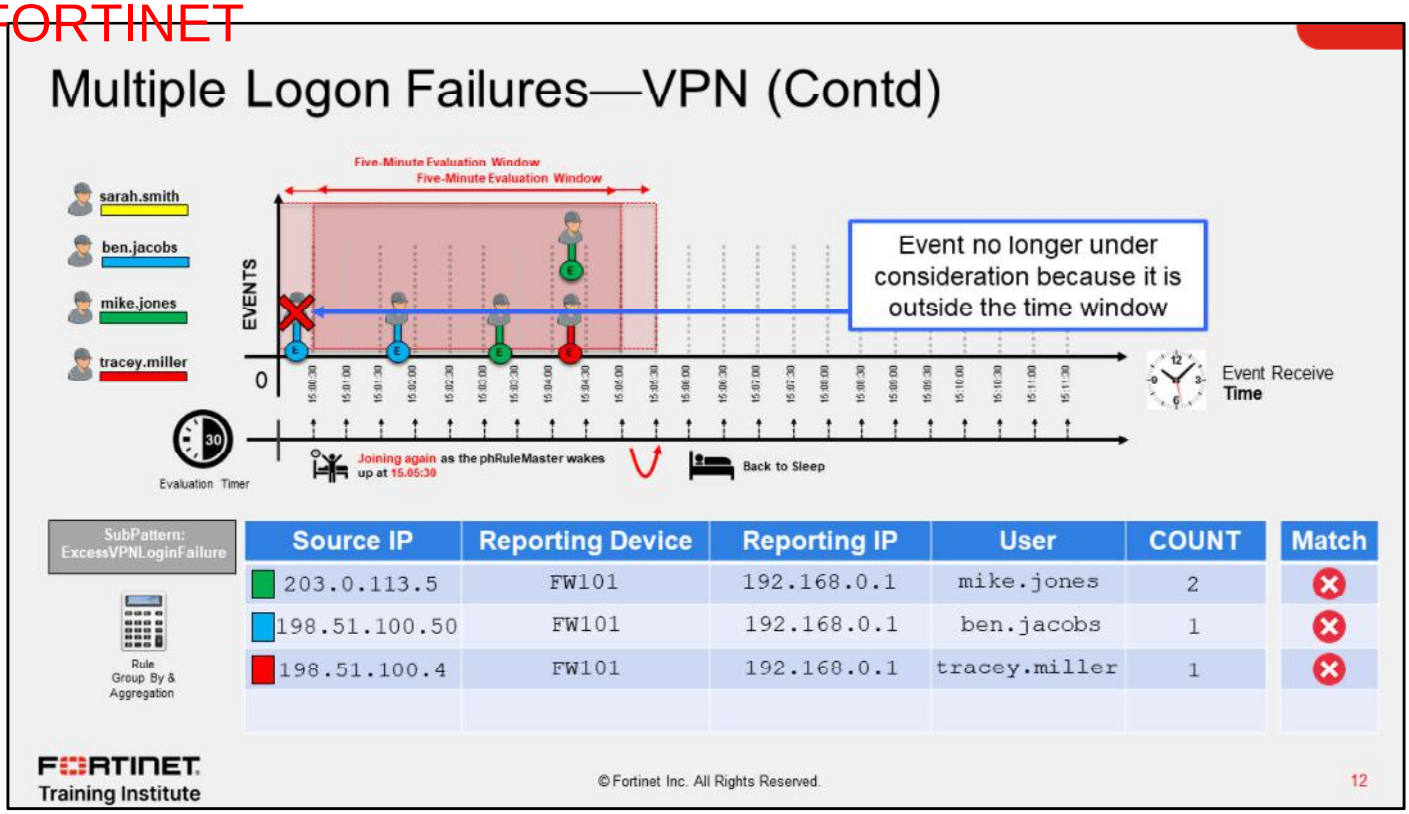
DO NOT REPRINT
© FORTINET



To make the simulation easier to follow, the example shown on this slide uses a five-minute time window instead of the actual 10-minute window. There are four users—Sarah, Ben, Mike, and Tracey—who are part of the VPN user group.

The phRuleMaster wakes up at 15:05:00 to evaluate the events that have been queued. It evaluates all the events related to VPN logon failures for the past five minutes. Ben and Mike had two logon failures and Tracey had one logon failure**.** The group by definition is configured for source IP, reporting device, reporting IP, and user. The aggregation count tracks the number of events against each row. In this evaluation period, the count is less than five for the three users, therefore no incident is triggered and the phRuleMaster goes back to sleep.

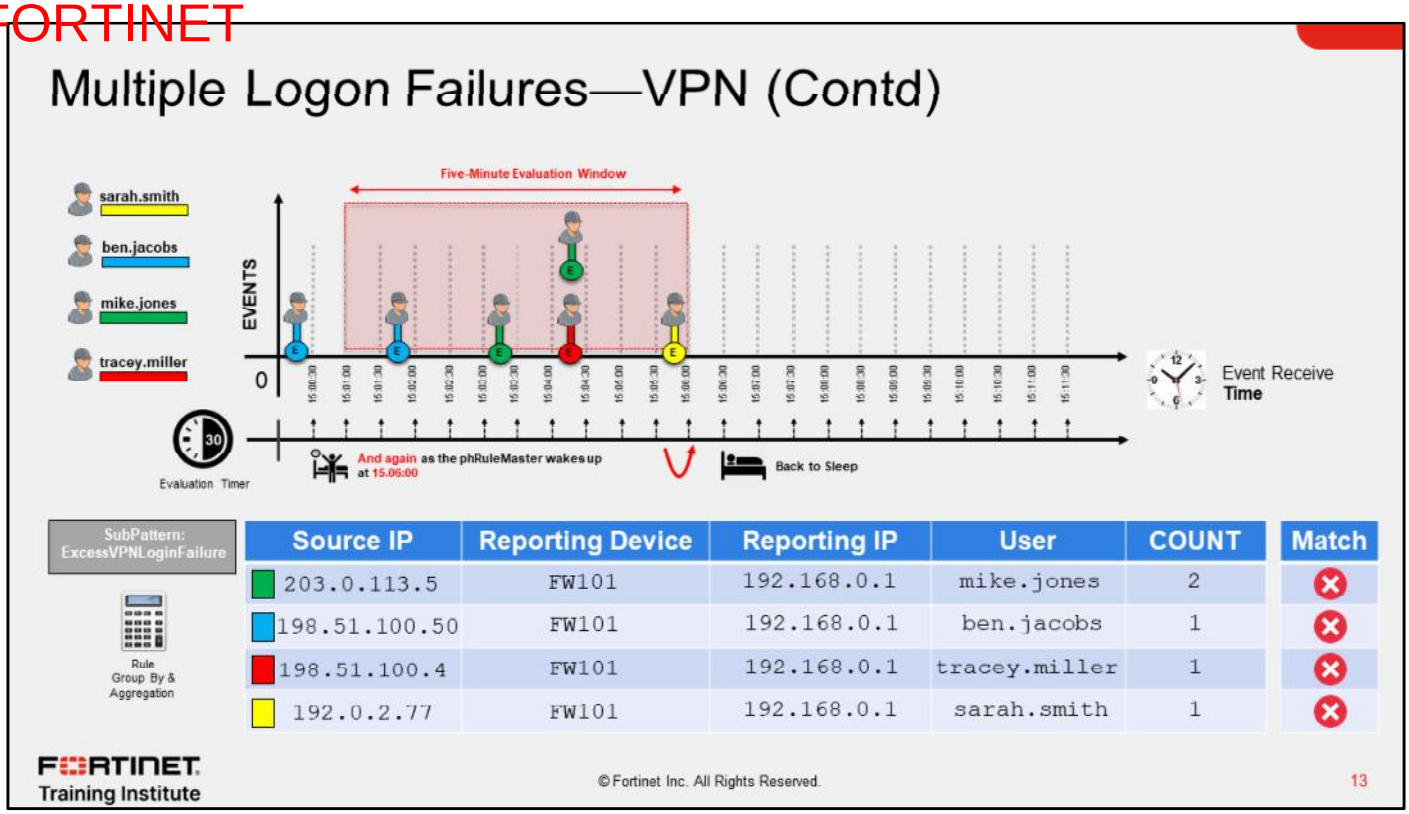
DO NOT REPRINT
© FORTINET



After 30 seconds, the phRuleMaster wakes up again to evaluate the events that have been queued. It again evaluates all the events related to VPN logon failures for the past five minutes. Any event outside the evaluation time period is no longer under consideration.

The COUNT column is updated with two for Mike, one for Ben, and one for Tracey. The COUNT value for each row is still less than five, therefore no incident is triggered and the phRuleMaster goes back to sleep.

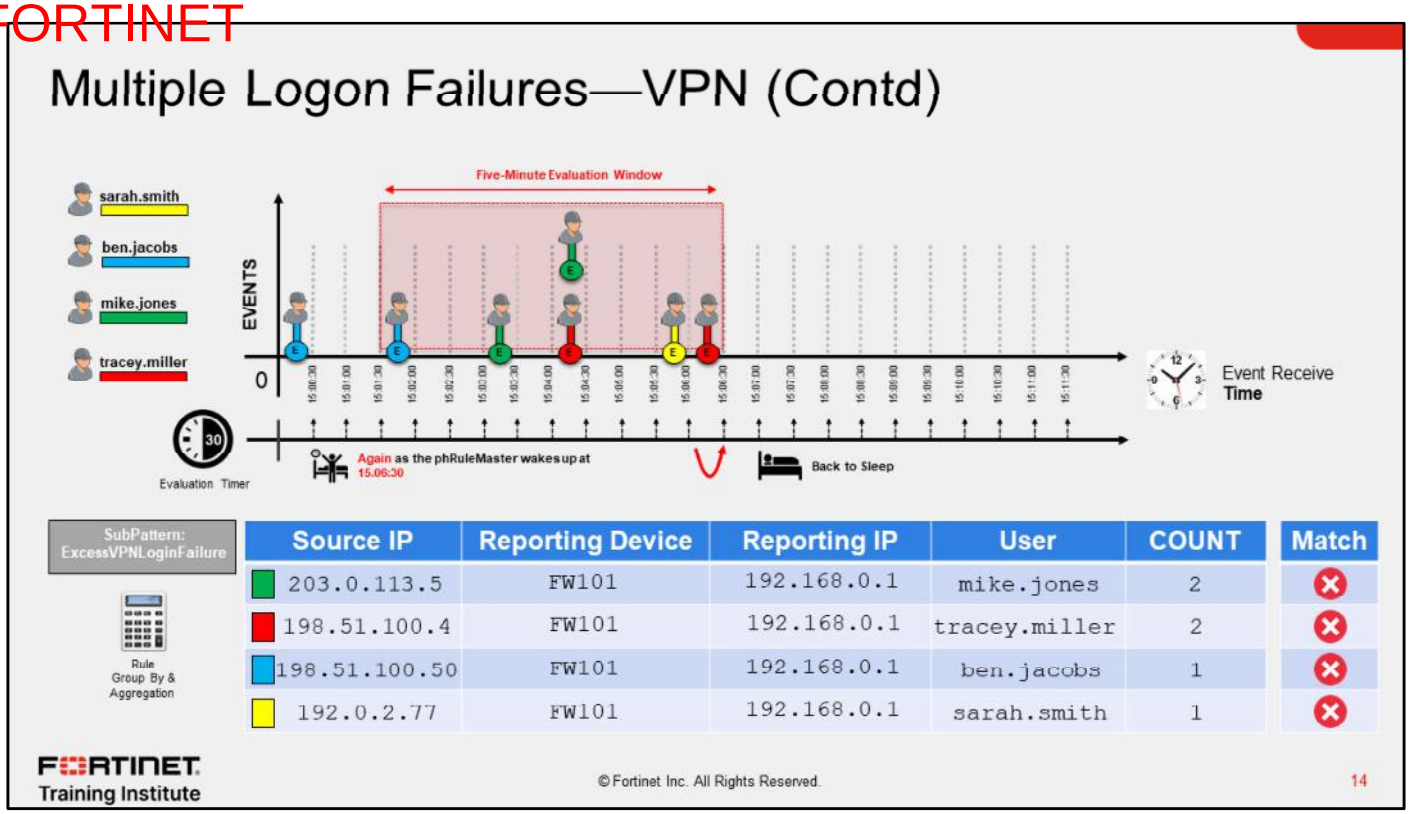
DO NOT REPRINT
© FORTINET



After 30 seconds, the phRuleMaster wakes up again to evaluate the events that have been queued. It again evaluates all the events related to VPN logon failures for the past five minutes.

This time there is one failure event for Sarah, so a new row is added based on the group by condition, and the COUNT column is updated with two for Mike, one for Ben, one for Tracey, and one for Sarah. The COUNT value for each row is still less than five, so no incident is triggered and the phRuleMaster goes back to sleep.

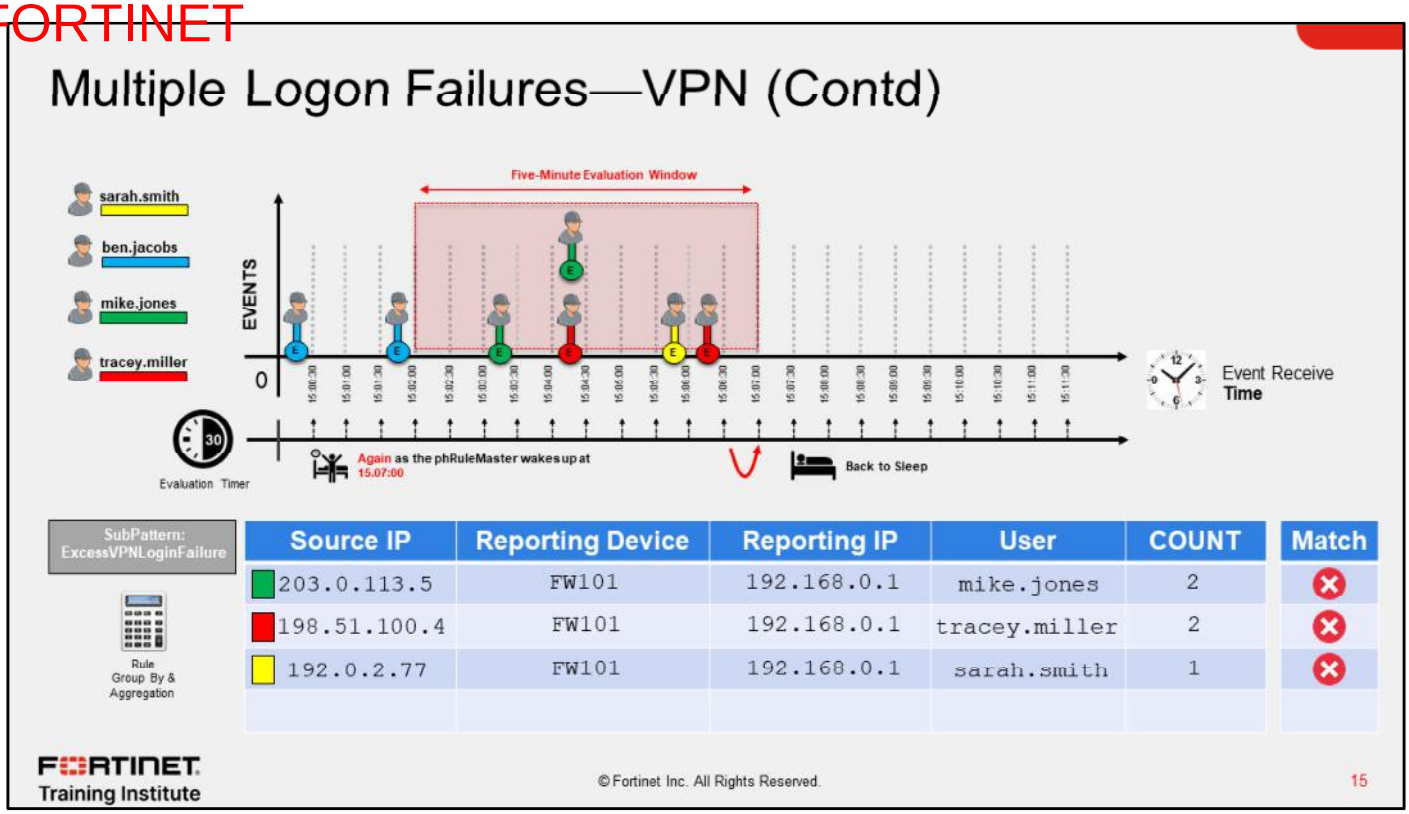
DO NOT REPRINT
© FORTINET



After 30 seconds, the phRuleMaster wakes up again to evaluate the events that have been queued. It again evaluates all the events related to VPN logon failures for the past five minutes.

The COUNT column is updated with two for Mike, one for Ben, two for Tracey, and one for Sarah. The COUNT value for each row is still less than five, so no incident is triggered and the phRuleMaster goes back to sleep.

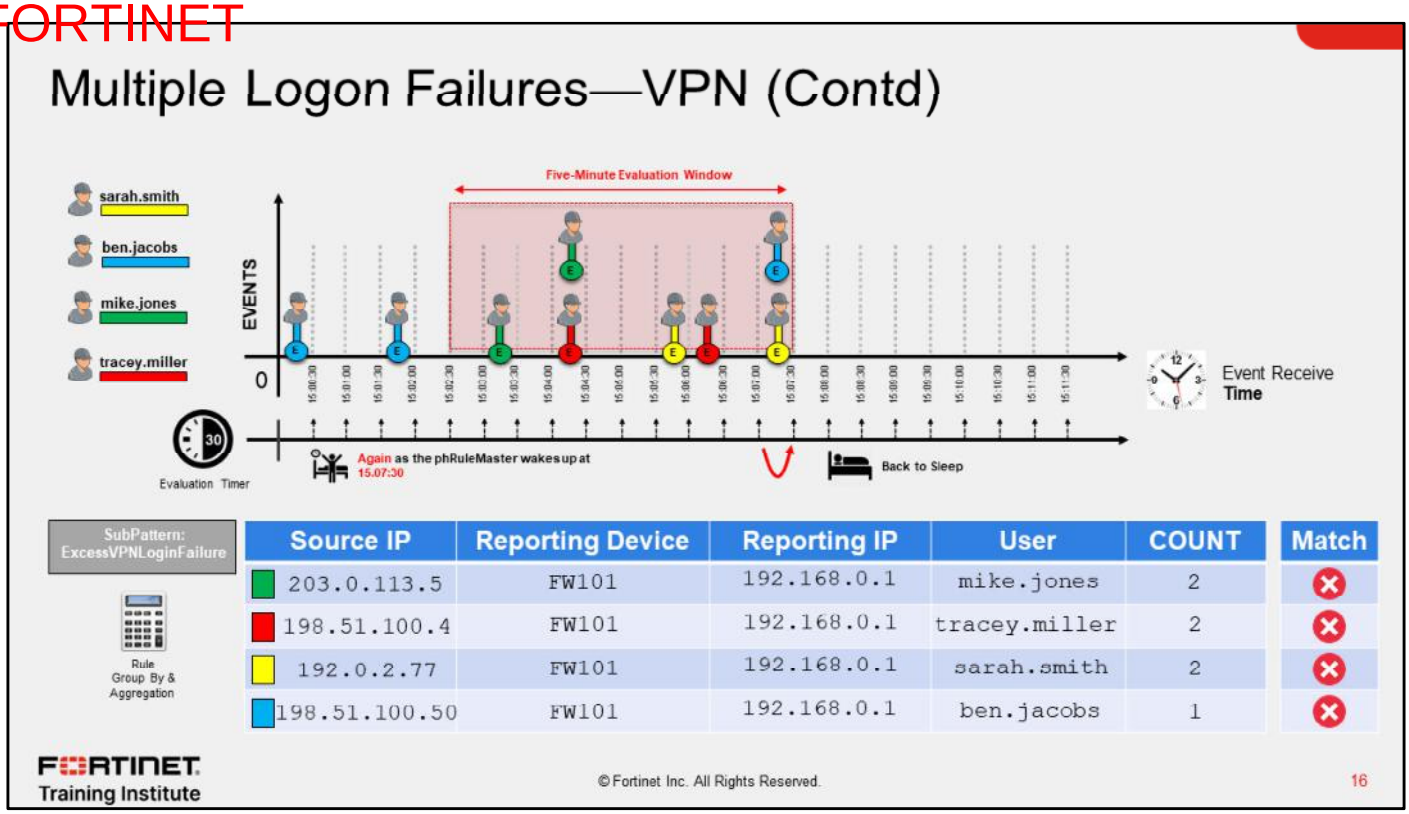
DO NOT REPRINT
© FORTINET



After 30 seconds, the phRuleMaster wakes up again to evaluate the events that have been queued. It again evaluates all the events related to VPN logon failures for the past five minutes.

The COUNT column is updated with two for Mike, two for Tracey, and one for Sarah. The row for Ben is removed because it is not in the evaluation period. The COUNT value for each row is still less than five, so no incident is triggered and the phRuleMaster goes back to sleep.

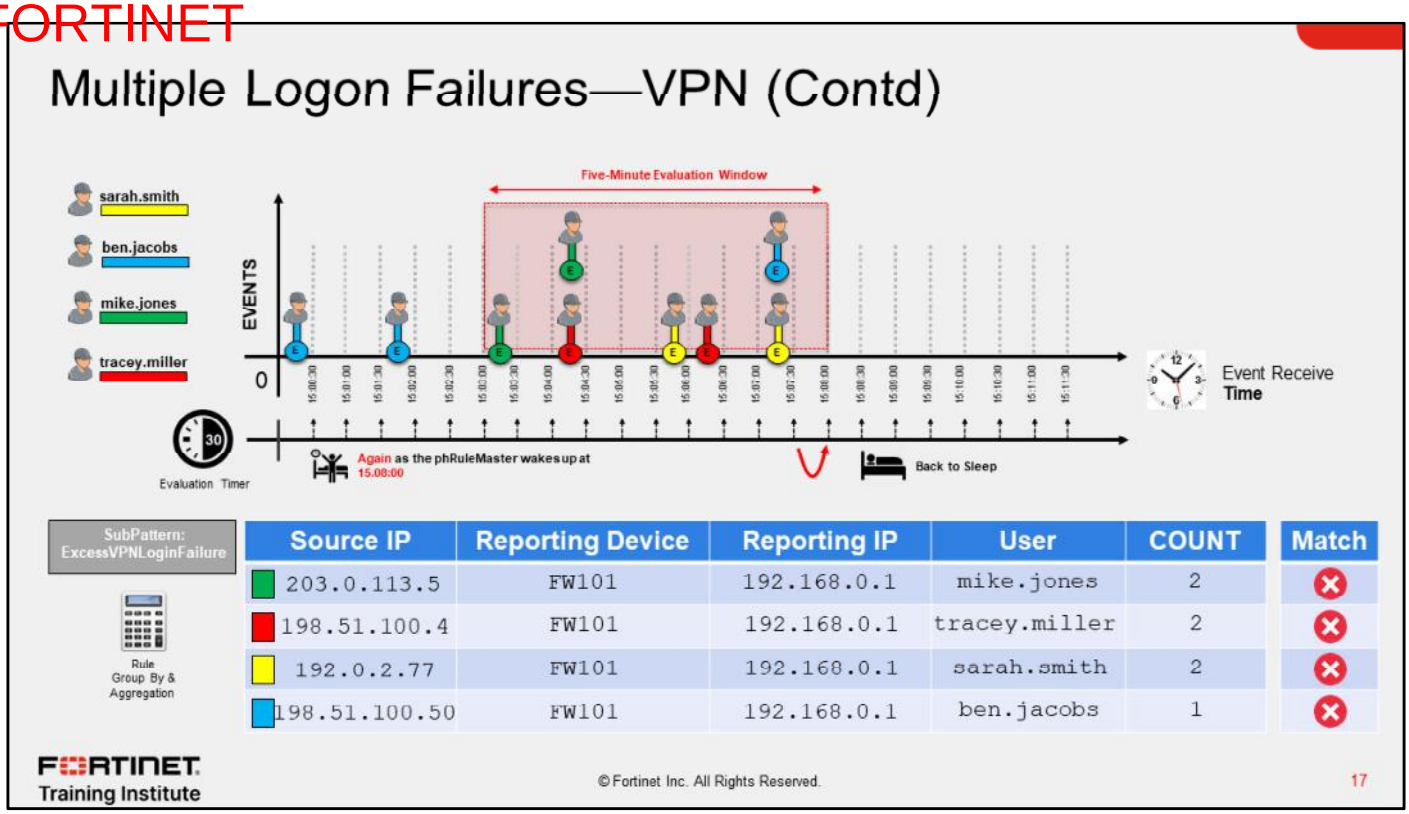
DO NOT REPRINT
© FORTINET



After 30 seconds, the phRuleMaster wakes up again to evaluate the events that have been queued. It again evaluates all the events related to VPN logon failures for the past five minutes.

The COUNT column is updated with two for Mike, two for Tracey, two for Sarah, and one for Ben. A new row for Ben is added again because there is a new logon failure event within the new evaluation period. The COUNT value for each row is still less than five, so no incident is triggered and the phRuleMaster goes back to sleep.

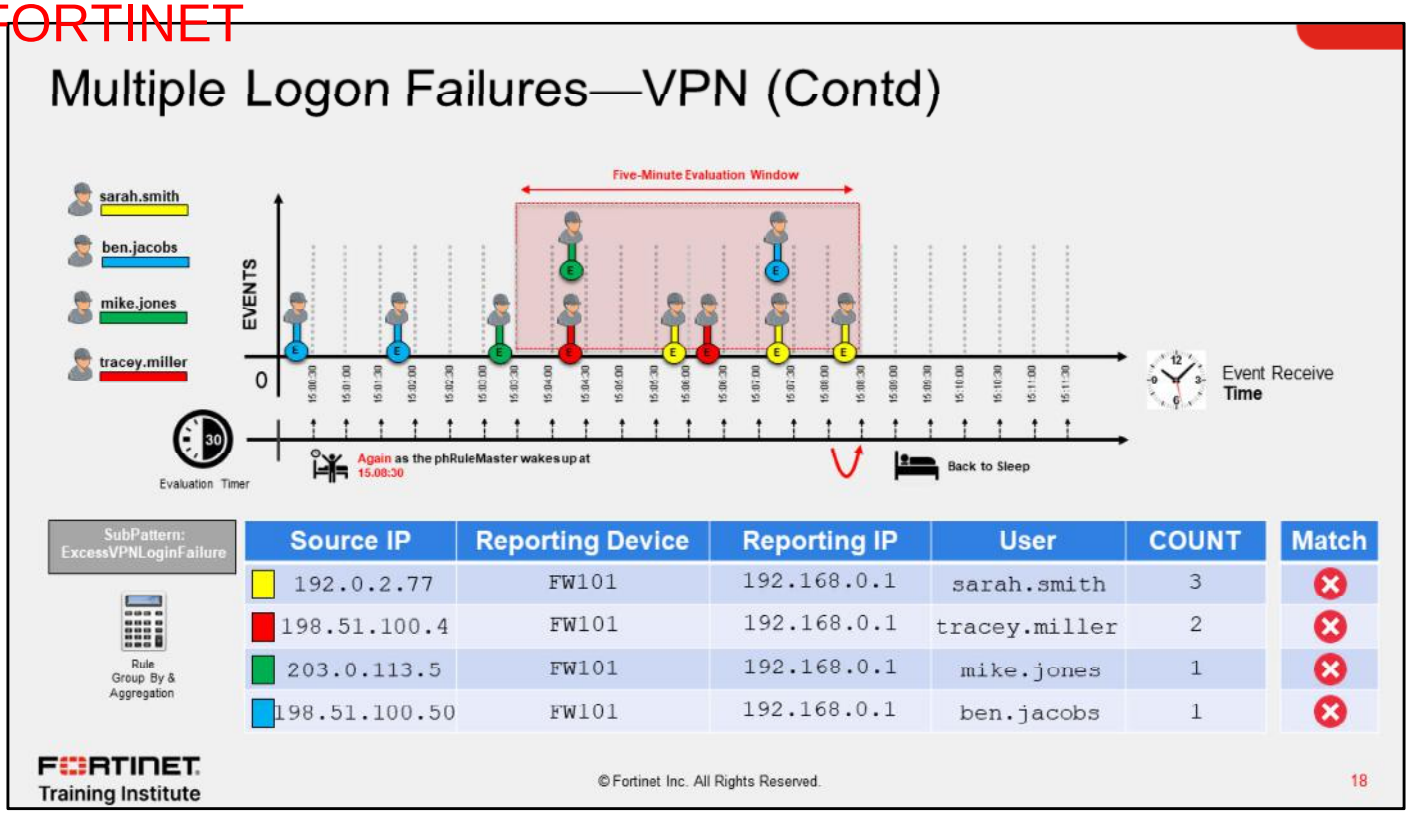
DO NOT REPRINT
© FORTINET



After 30 seconds, the phRuleMaster wakes up again to evaluate the events that have been queued. It again evaluates all the events related to VPN logon failures for the past five minutes.

The COUNT column is updated with two for Mike, two for Tracey, two for Sarah, and one for Ben. The COUNT value for each row is still less than five, therefore no incident is triggered and the phRuleMaster goes back to sleep.

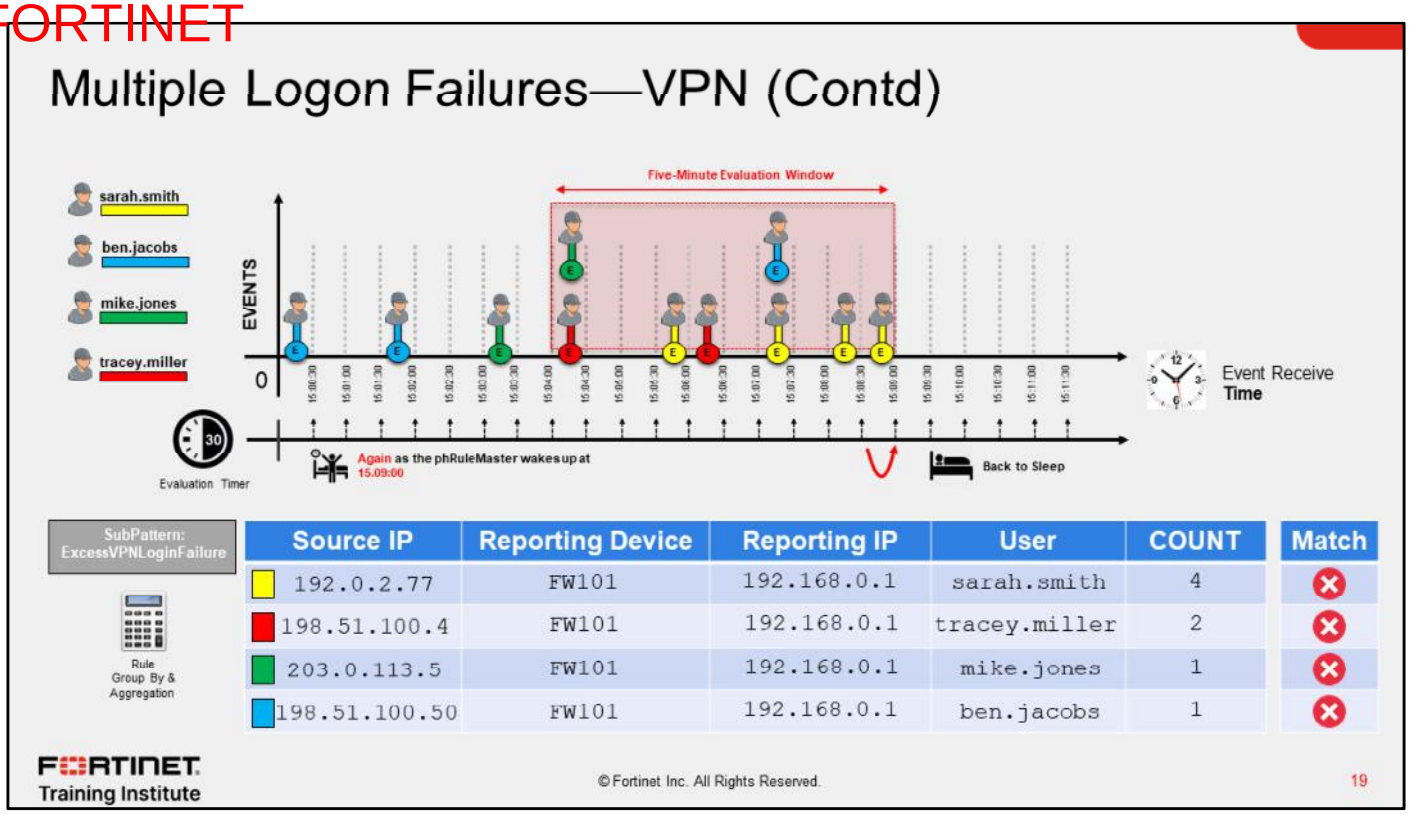
DO NOT REPRINT
© FORTINET



After 30 seconds, the phRuleMaster wakes up again to evaluate the events that have been queued. It again evaluates all the events related to VPN logon failures for the past five minutes.

The COUNT column is updated with one for Mike, two for Tracey, three for Sarah, and one for Ben. The COUNT value for each row is still less than five, so no incident is triggered and the phRuleMaster goes back to sleep.

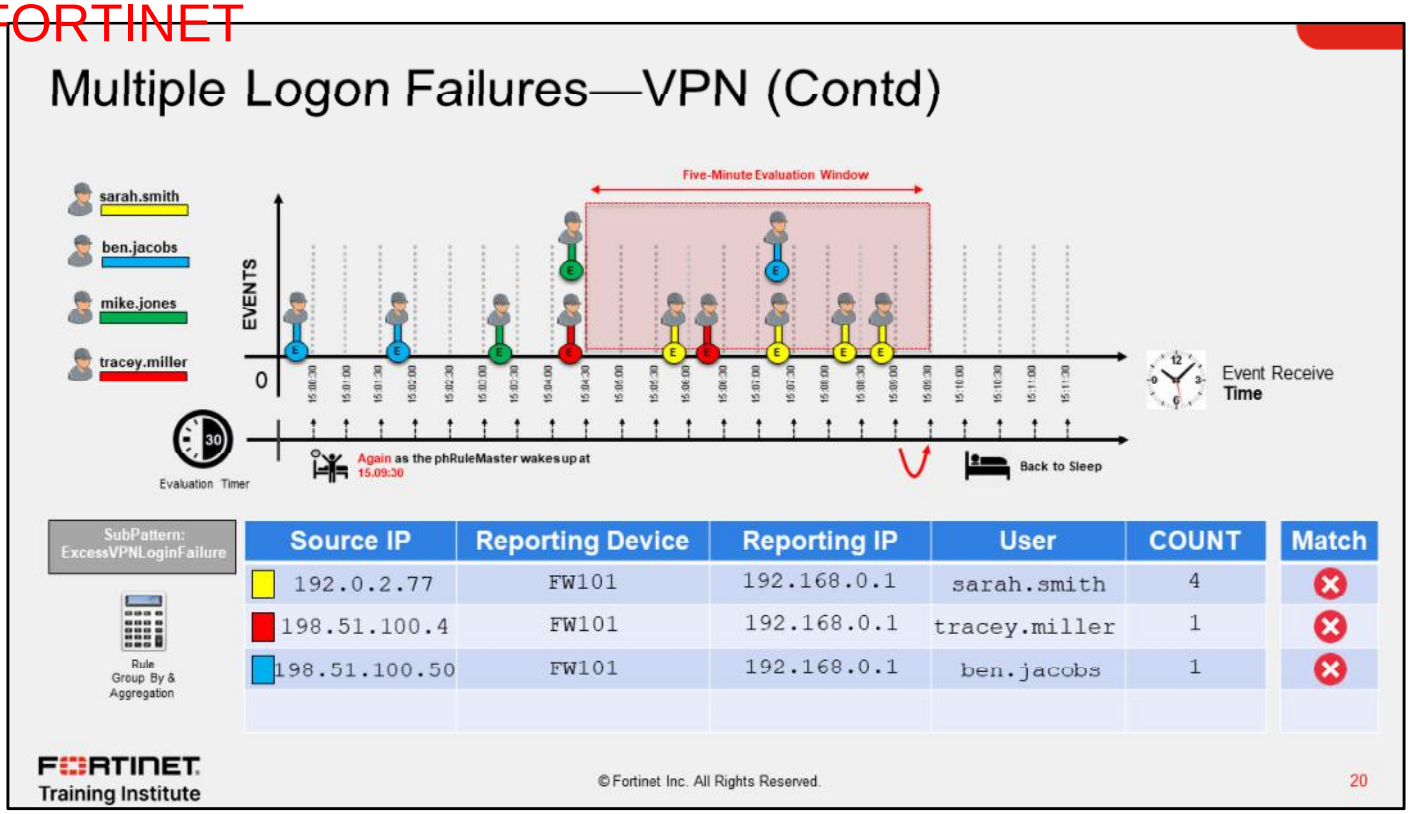
DO NOT REPRINT
© FORTINET



After 30 seconds, the phRuleMaster wakes up again to evaluate the events that have been queued. It again evaluates all the events related to VPN logon failures for the past five minutes.

The COUNT column is updated with one for Mike, two for Tracey, four for Sarah, and one for Ben. The COUNT value for each row is still less than five, so no incident is triggered and the phRuleMaster goes back to sleep.

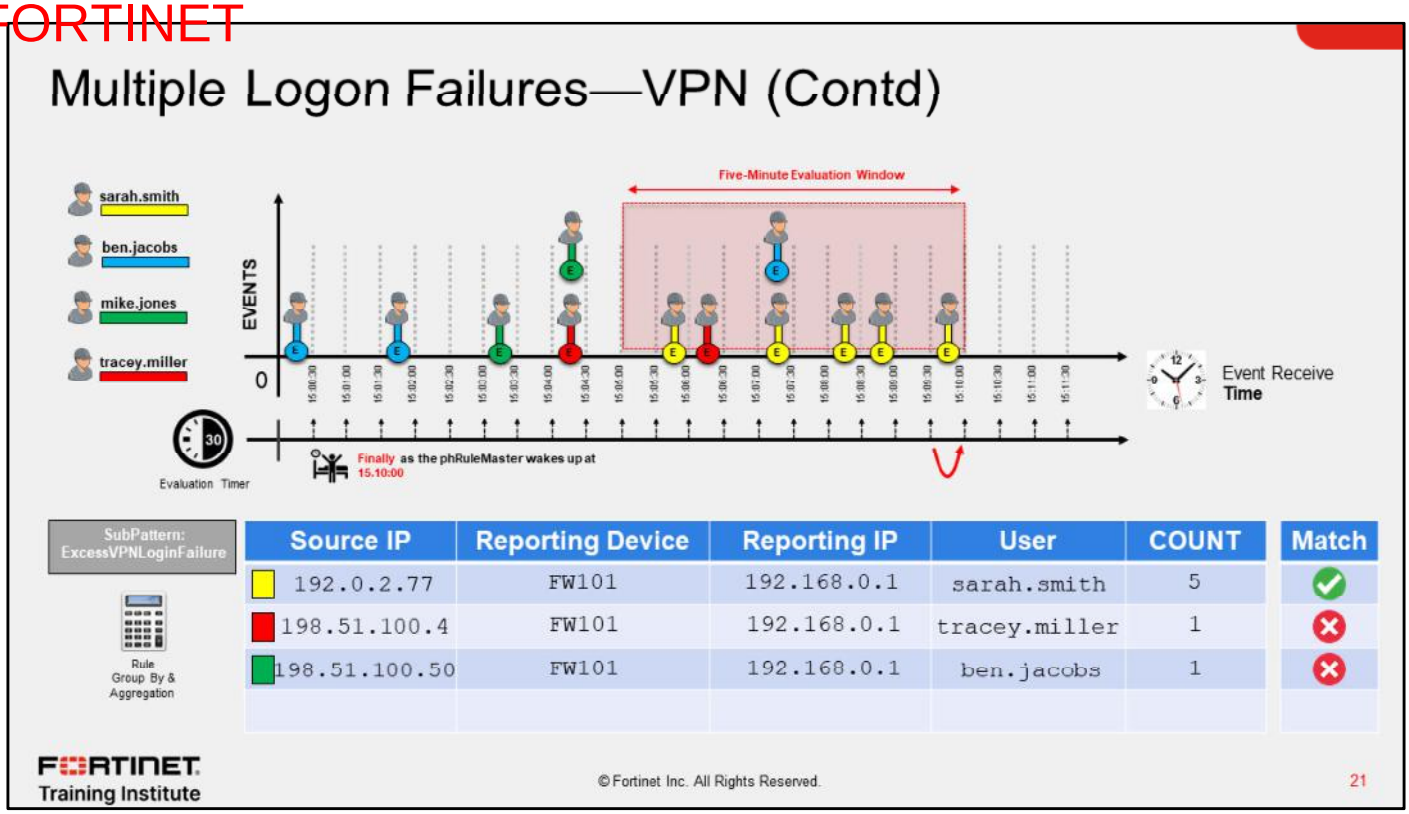
DO NOT REPRINT
© FORTINET



After 30 seconds, the phRuleMaster wakes up again to evaluate the events that have been queued. It again evaluates all the events related to VPN logon failures for the past five minutes.

The COUNT column is updated with one for Tracey, four for Sarah, and one for Ben. The row for Mike is removed because it is not in the evaluation period. The COUNT value for each row is still less than five, so no incident is triggered and the phRuleMaster goes back to sleep.

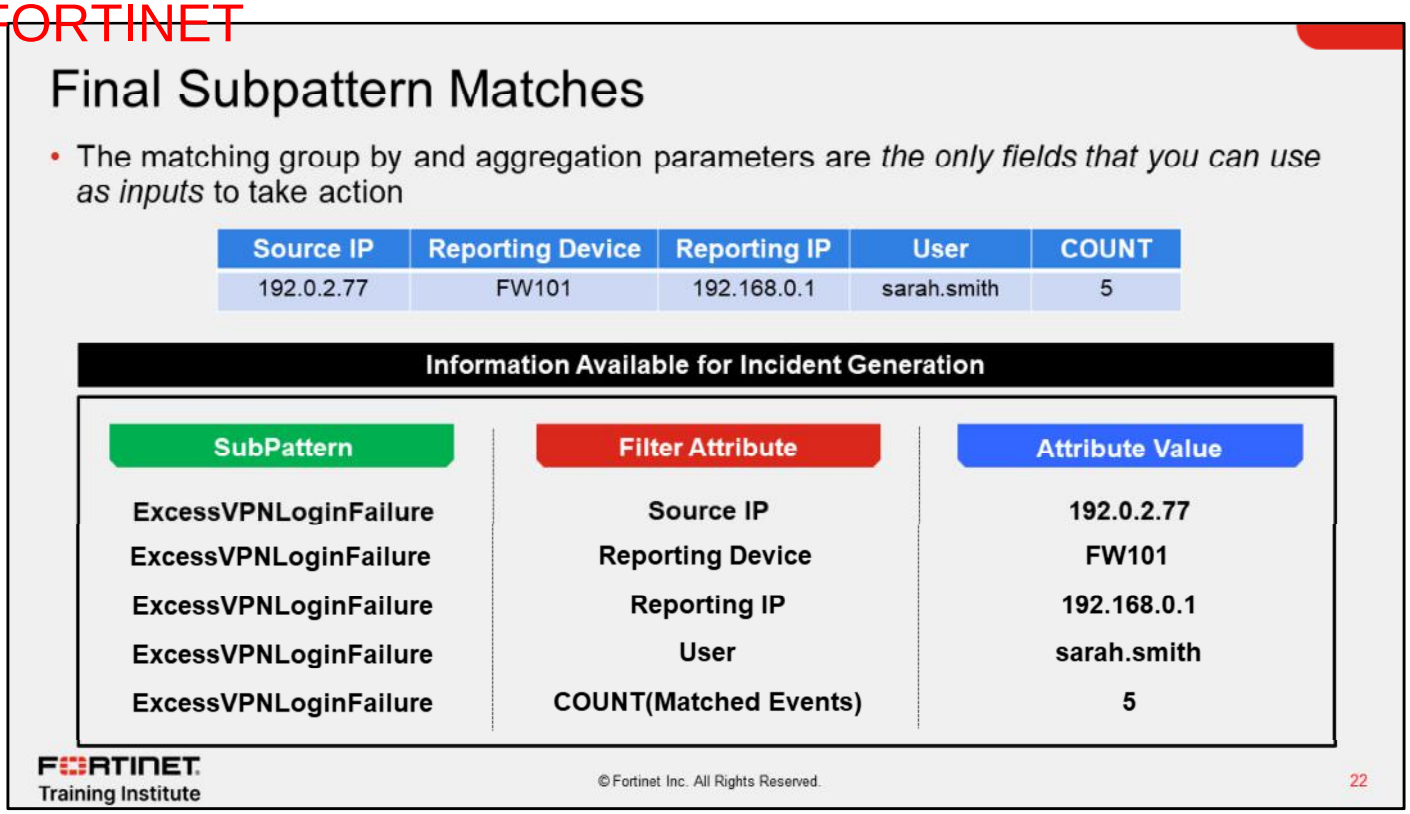
DO NOT REPRINT
© FORTINET



After 30 seconds, the phRuleMaster wakes up again to evaluate the events that have been queued. It again evaluates all the events related to VPN logon failures for the past five minutes.

The COUNT column is updated with five for Sarah, one for Tracey, and one for Ben. This time there is a match for Sarah because of the five failed logon attempts.

DO NOT REPRINT
© FORTINET

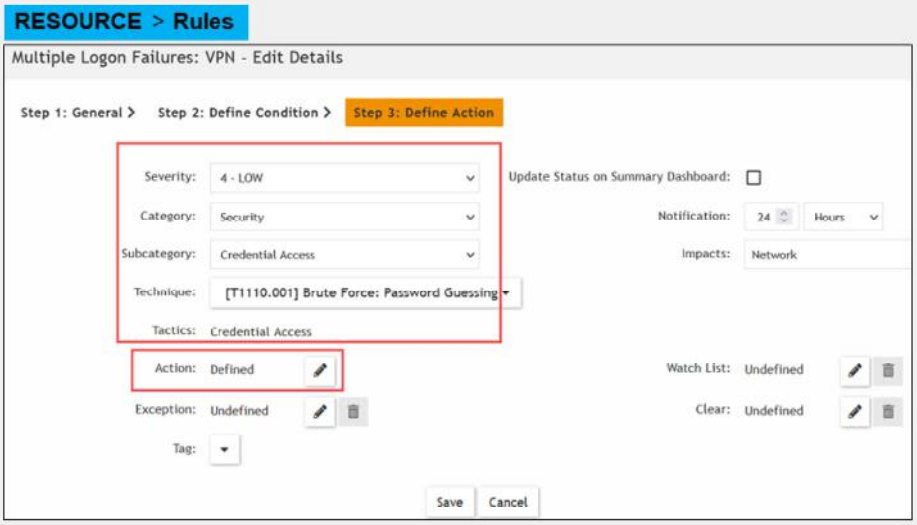


The matching group by and aggregation parameters are the only fields that you can use as inputs to take action.

DO NOT REPRINT
© FORTINET

Rule Options—Multiple VPN Logon Failures

- Once the incident parameters are identified, you must configure an action, and set the **Severity**, **Category**, and **Subcategory** of the resulting incident



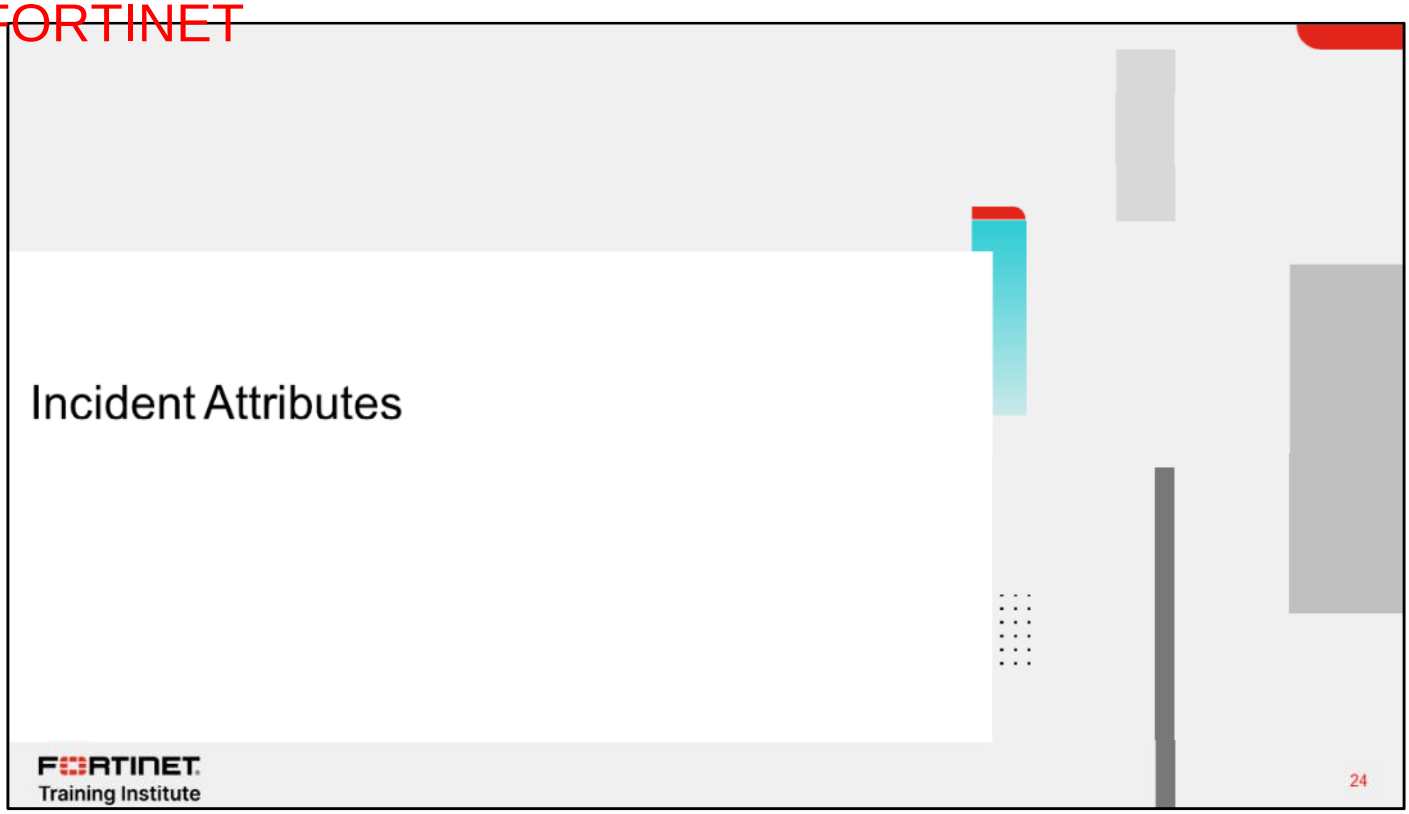
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

23

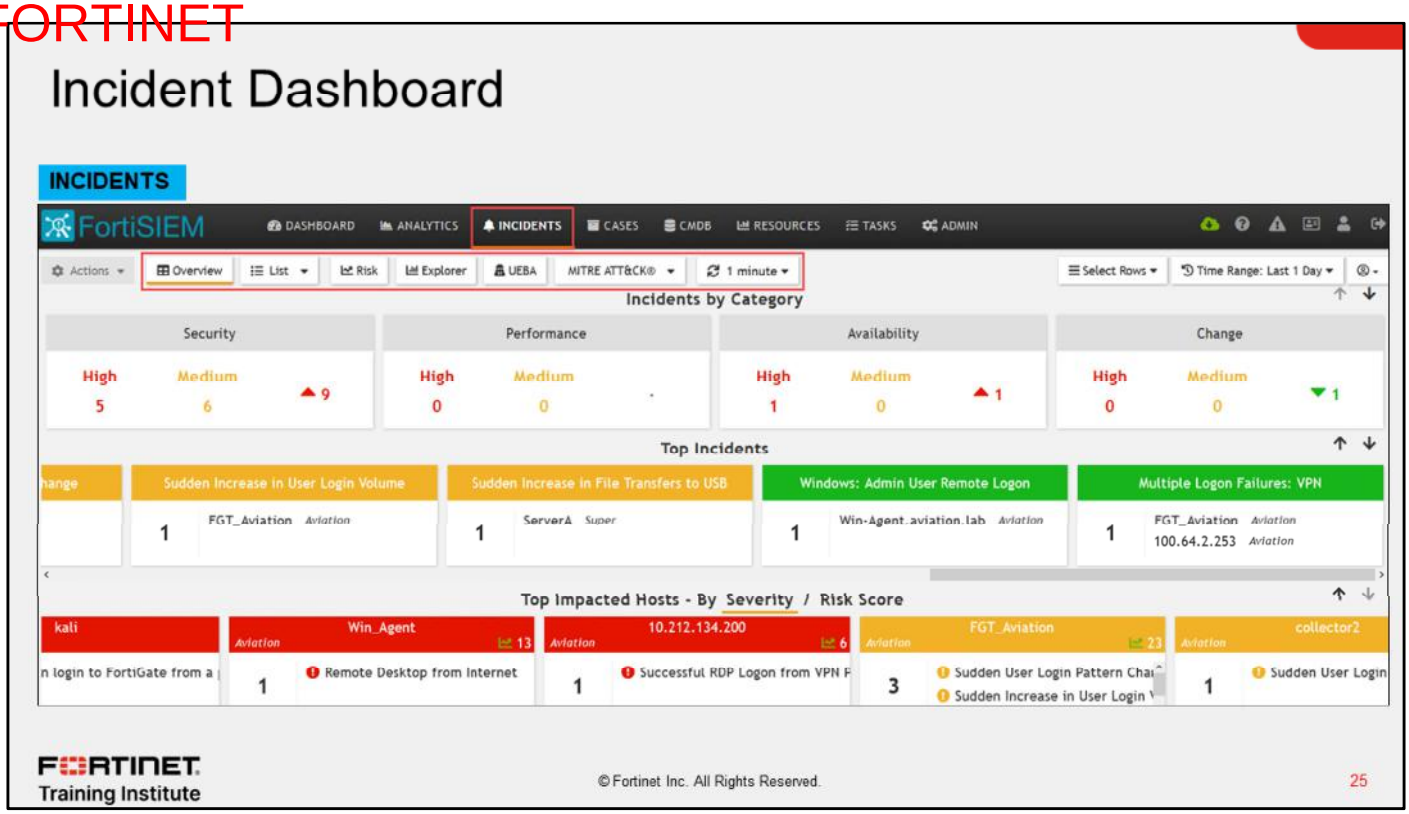
Once the incident parameters are determined, you must configure an **Action**, along with the **Severity**, **Category**, and **Subcategory** of the resulting incident.

DO NOT REPRINT
© FORTINET



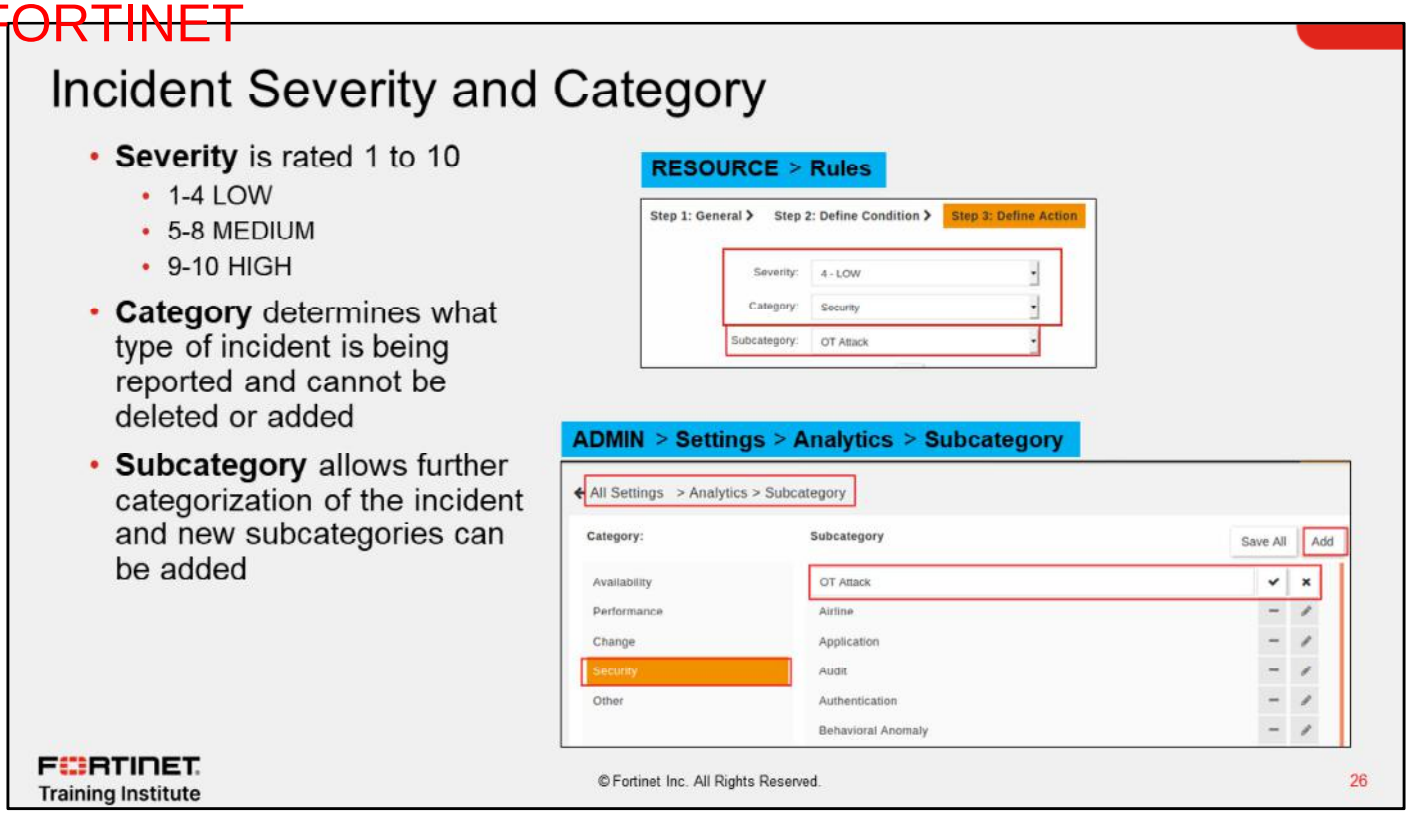
In this section, you will learn about incident attributes, severity, and category.

DO NOT REPRINT
© FORTINET



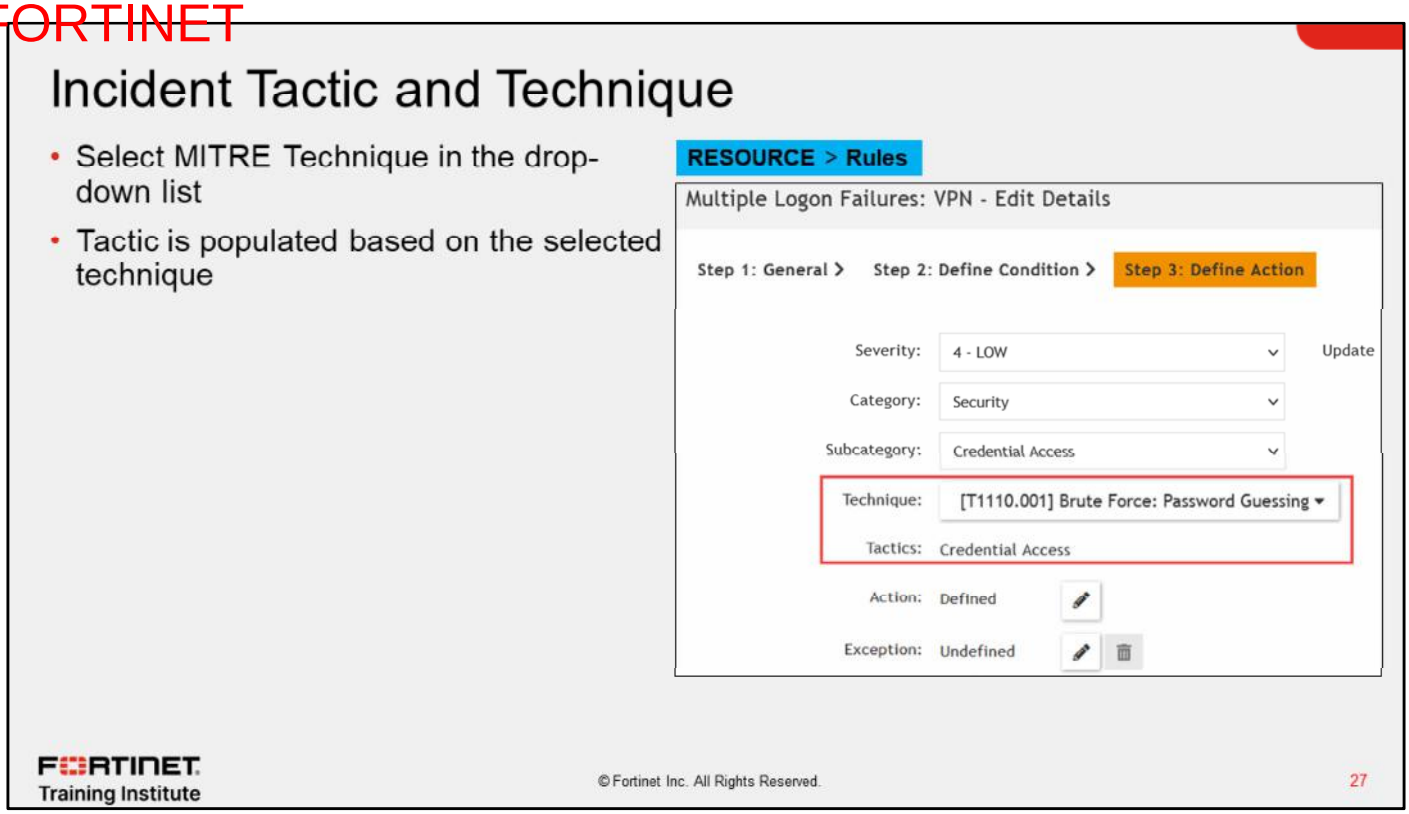
The incident dashboard displays incident information for your IT infrastructure based on the filter conditions you configure. You can also view incidents grouped by incident attributes, use values in incident attributes to refine your searches, view information about rules that triggered incidents, and use incident information to create rule exceptions and event dropping rules.

DO NOT REPRINT
© FORTINET



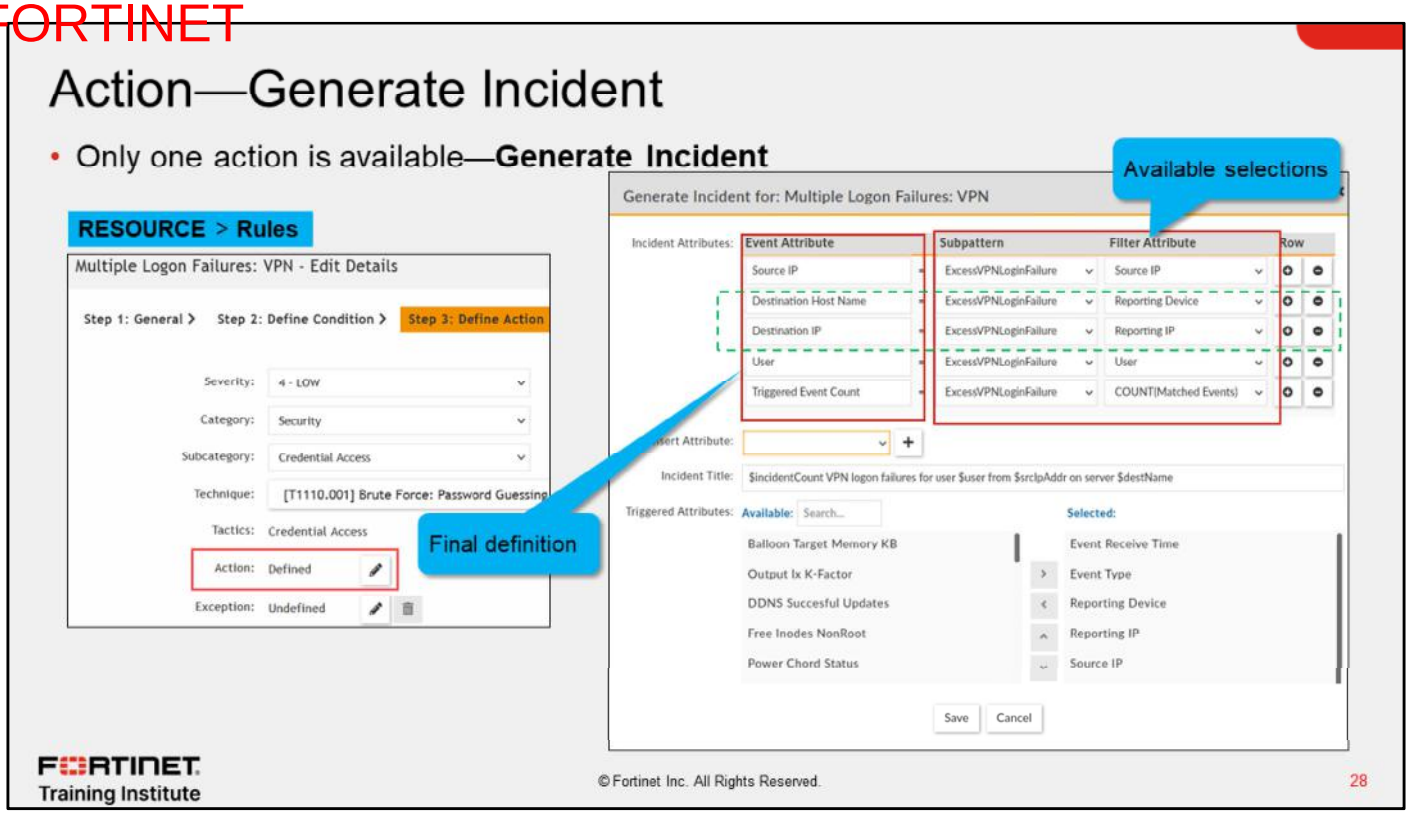
When you configure a rule, you can define the action that must be taken when the rule matches the aggregate conditions. Before you define an action, you must specify the severity, category, and subcategory for the incident. Severity is rated 1-10, where 1-4 is low, 5-8 is medium, 9-10 is high. The category determines the type of incident that is being reported, which is useful for filtering on the incident dashboard. The subcategory allows a further categorization of the incident.

DO NOT REPRINT
© FORTINET



In the **Technique** drop-down list, select a technique that describes the incident that the rule created. The tactic is automatically populated based on the selected technique.

DO NOT REPRINT
© FORTINET



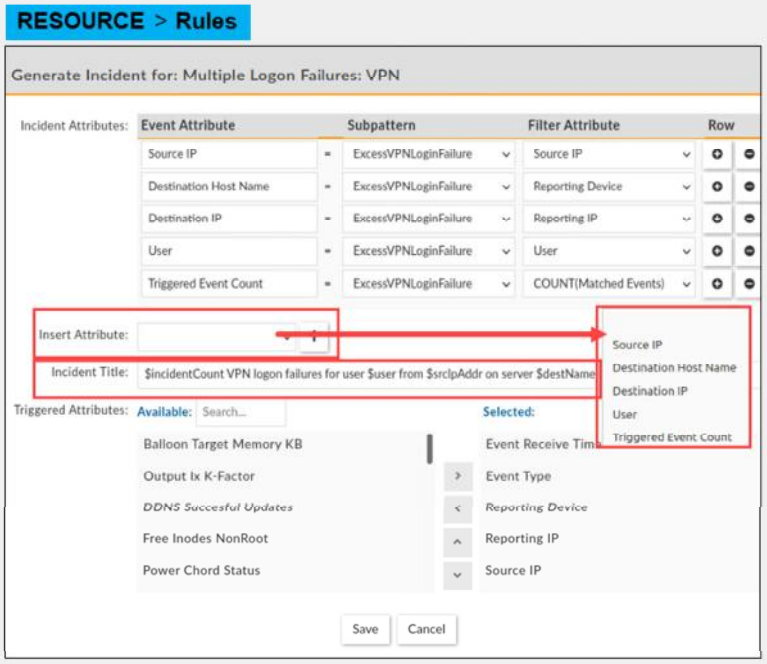
When you edit the **Action** field, the only fields that you can select are from the matching subpattern aggregate and group by conditions. Most often for out-of-the-box rules, the event attribute is the filter attribute, but the two are different.

In the example shown on this slide, the reporting device is the destination host name and the reporting IP is the destination IP because the firewall that is reporting the events to FortiSIEM is actually the device where the VPN terminates, so it is considered the final destination.

DO NOT REPRINT
© FORTINET

Custom Incident Title

- Title of the incident can be customized
- Add attribute variables, strings, numbers, and characters to the incident title
- Every attribute variable should be preceded by a **$** sign
- Select the attribute variable in the drop-down list in the **Insert Attribute** field



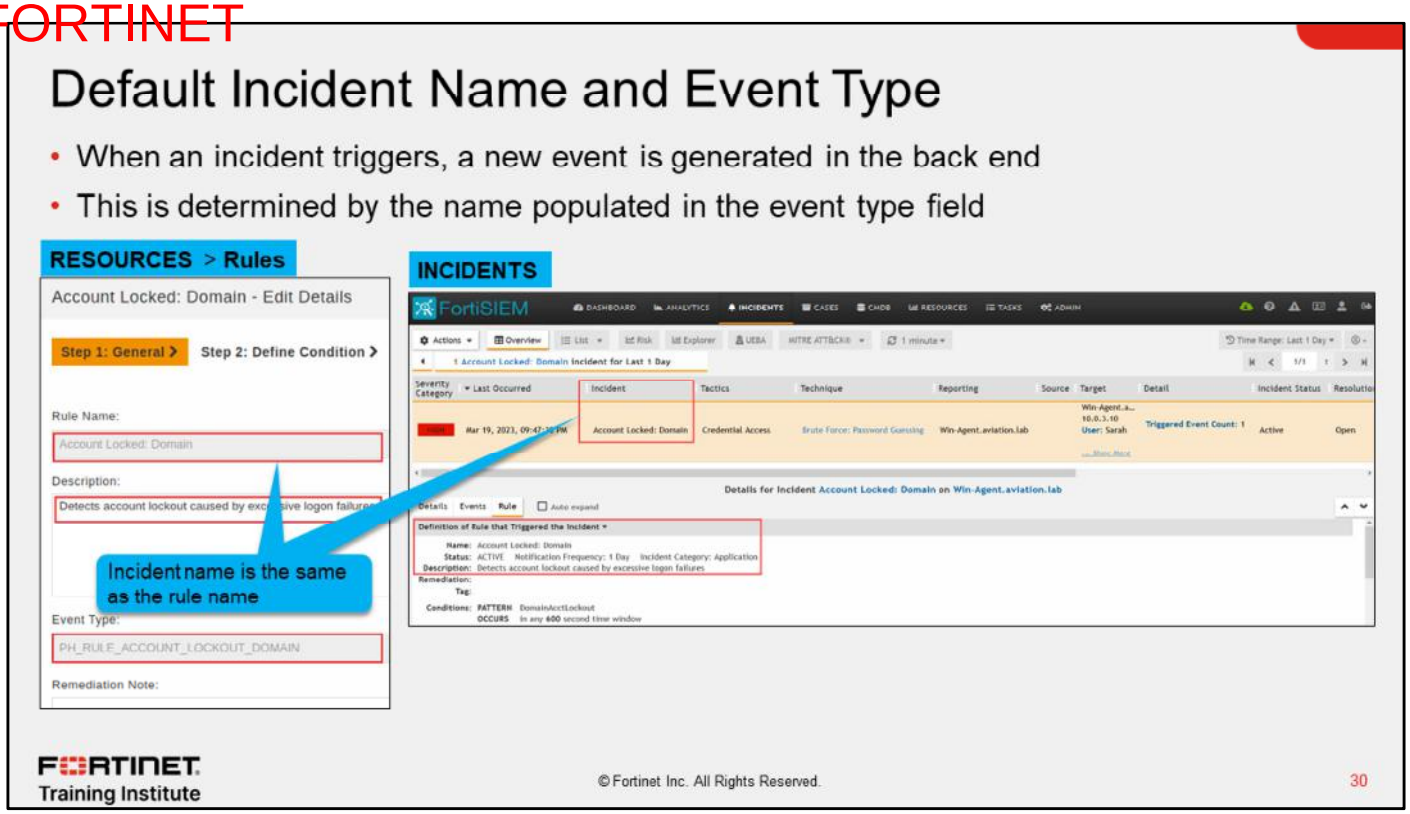
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

29

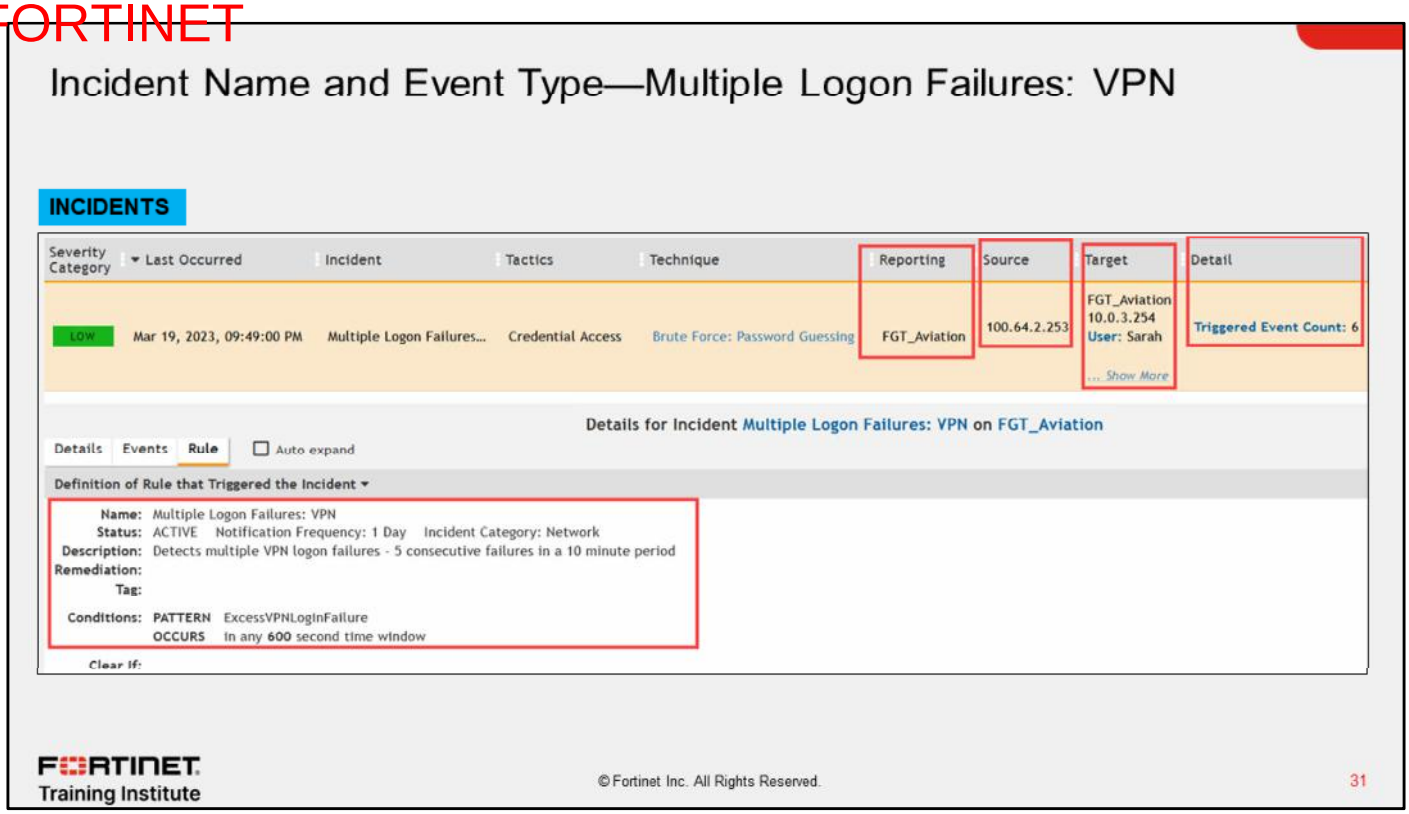
You can customize the default incident title by adding attribute variables related to the incident, strings, numbers, and characters. You should precede each attribute variable that you type in the incident title with a **$**. When the incident is generated, the title of the incident has the value associated with the attribute variable that you specified in the incident title.

DO NOT REPRINT
© FORTINET



When an incident triggers, a new event is actually generated in the back end, which is not visible by default. The name of the incident is determined by the rule name, and the description of the incident is also borrowed from the description in the rule.

DO NOT REPRINT
© FORTINET



This slide shows another incident example—**Multiple Logon Failures: VPN**. The **Source**, **Target**, and **Detail** columns show a single line summary of the incident, which is automatically determined by FortiSIEM by analyzing the incident attributes that you defined in the rule.

What was considered the source of the attack and the target in your environment against which the attack was made, and any other quick supporting details, are populated by extracting the information from the events. Additionally, other incident attributes that have been generated that are not related to the source or target are collated under the **Detail** section of the incident. This information is important for notification policies and, especially, remediation actions.

DO NOT REPRINT
© FORTINET

Incident Source Attributes

- If an incident attribute is in this table (hardcoded), it appears under the **Incident Source** field on the **Incident** tab.

Incident Source Attributes		
Source IP	Source Interface SNMP index	Source State
Source MAC	Sent TCP flags	Source City
Source TCP/UDP Port	Source Host Name	Source Organization
Source Autonomous System Number	Source Firewall Zone	Source Latitude
Source VLAN	Source FibreChannel WWN Id	Source Longitude
Source Interface Name	Source TOS	Source App Name
Source Device Port	Source Interface Alias	Source Application Protocol
Source DSCP	Source Physical Machine Name	Source User
Src Terminal Name	Source Domain	
Source VM Name	Source Country	

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

32

The incident source attributes shown on this slide are hardcoded on FortiSIEM. These attributes will appear under the incident *source* field on the **Incident** tab. The most common source attributes are source IP and source user.

DO NOT REPRINT
© FORTINET

Incident Target Attributes

Incident Target Attributes				
Application Name	Destination FibreChannel WWN Id	Destination Longitude	Target Name	Host State
Destination IP	Destination Interface Alias	Destination Application Protocol	Target Type	Host City
Destination Host Name	Destination DSCP	Destination User	Target Hash Code	Host Organization
Destination MAC	Destination VM Name	Target Host MAC	Host Vendor	Host Latitude
Destination TCP/UDP Port	Destination Physical Machine Name	Target Host Name	Host Model	Host Longitude
Destination Autonomous System Number	Destination Resource Pool	Target Computer	Host IP	Host Name

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

33

The incident target attributes shown on this slide are hardcoded on FortiSIEM. These attributes will appear under the incident **Target** field on the **Incident** tab. The most common target attributes are destination IP, destination host name, host IP, and host name.

DO NOT REPRINT
© FORTINET

Incident Target Attributes (Contd)

- Reporting IP and reporting device are not incident target attributes

Incident Target Attributes			
Destination VLAN	Destination Domain	Target Domain	User
Destination Interface Name	Destination Country	Target Object Type	Host MAC
Destination Device Port	Destination State	Target Object Name	Host VLAN
Destination Interface SNMP Index	Destination City	Target Cluster Balance	Host Location
Received TCP flags	Destination Organization	Target User	Host Virtual IP
Destination Firewall Zone	Destination Latitude	Target User Group	Host Country

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

34

This slide shows a few more examples of incident target attributes. Note that the reporting IP and reporting device are not incident target fields. So, if you have the reporting IP or reporting device as filter attributes while defining incident attributes, you must select a different event attribute.

DO NOT REPRINT
© FORTINET

Incident Attributes Special Case

- **User** and **Target User** are both incident target fields by default

Incident Target Attributes	
User	Target User



- If a list of incident attributes contains both **User** and **Target User**, the user attribute is placed in the incident source field

Incident Source Attributes
User

Incident Target Attributes
Target User

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

35

This is one of the special cases where **User** and **Target User** are both incident target fields by default. If a list of incident attributes contains both **User** and **Target User**, the user attribute is placed in the incident source field.

DO NOT REPRINT
© FORTINET

User vs. Target User

- What is a user and target user?
- Think about a login to the network or a device
  - There is only one field in the log message related to the user who tried to log in
- Now, think about an administrative change to a user account in Active Directory—maybe an account was disabled
  - The log contains two user-related fields
  - Who made the change, and which account was modified
- In this case, FortiSIEM parses the account that was changed to be the target user, and who made the change as the user


User


Target User

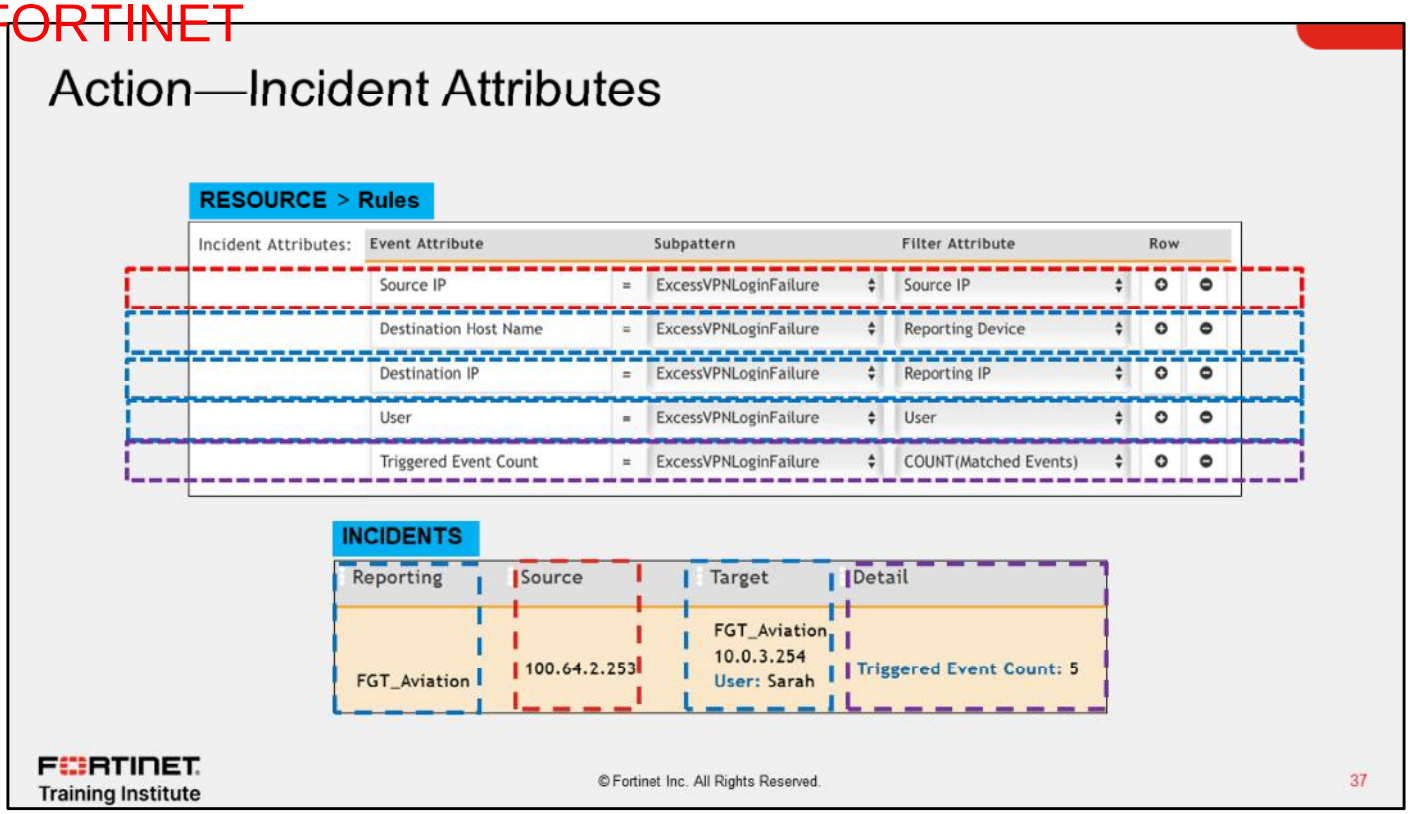
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

36

To differentiate between the user and target user, think about a login to the network or a device. There is only one field in the log message related to which user tried to log on. Now, think about an administrative change to a user account in Active Directory—maybe an account was disabled. The log contains two user-related fields—the one who made the change, and whose account was modified. In this case, FortiSIEM parses the account that was changed to be the target user, and who made the change as the user.

DO NOT REPRINT
© FORTINET



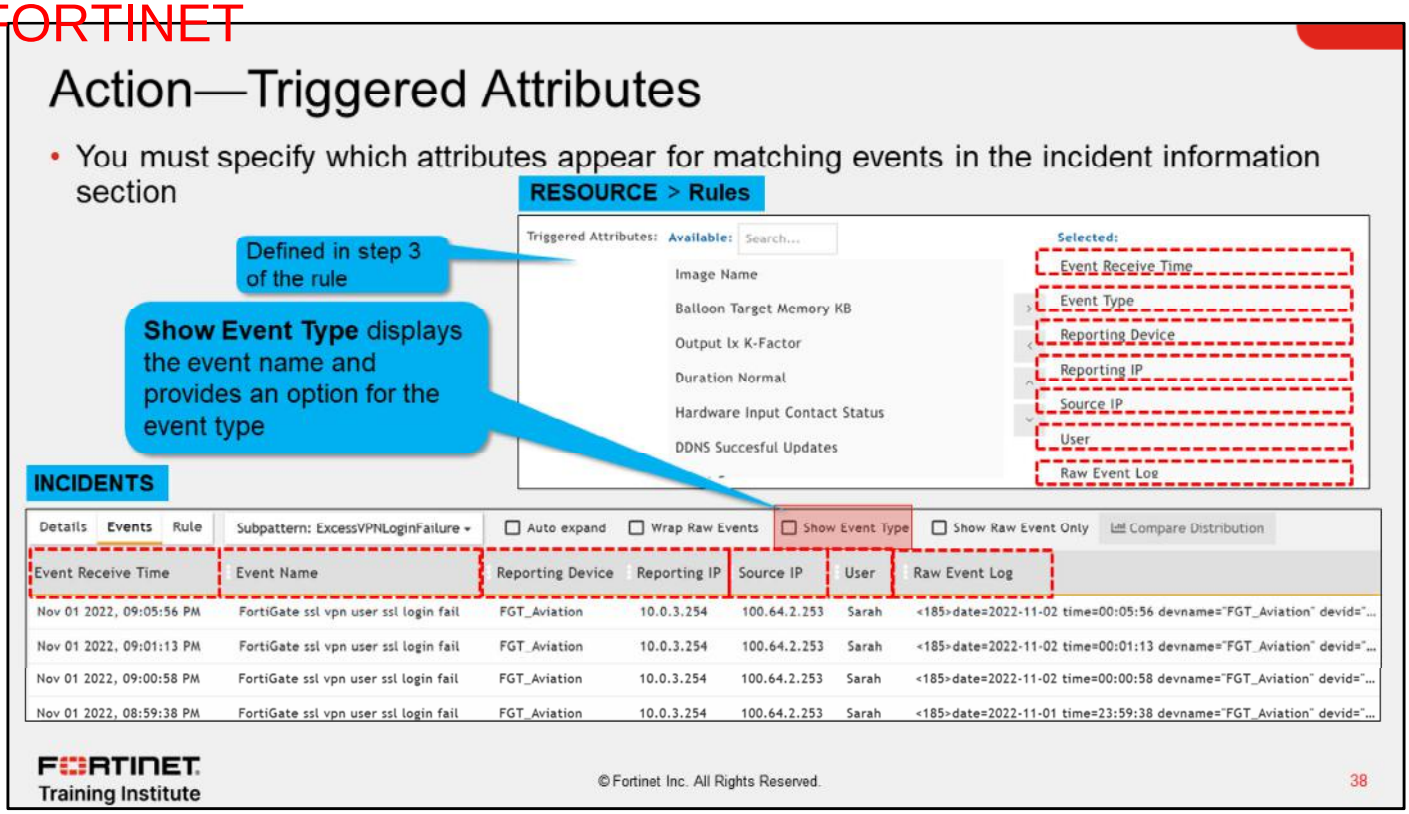
Take a look at the incident attributes for the **Multiple Logon Failures: VPN** rule shown on this slide.

The incident **Source** is the **Source IP** event attribute.

The incident **Target** is the **Reporting IP**, **Reporting Device**, and **User** attributes.

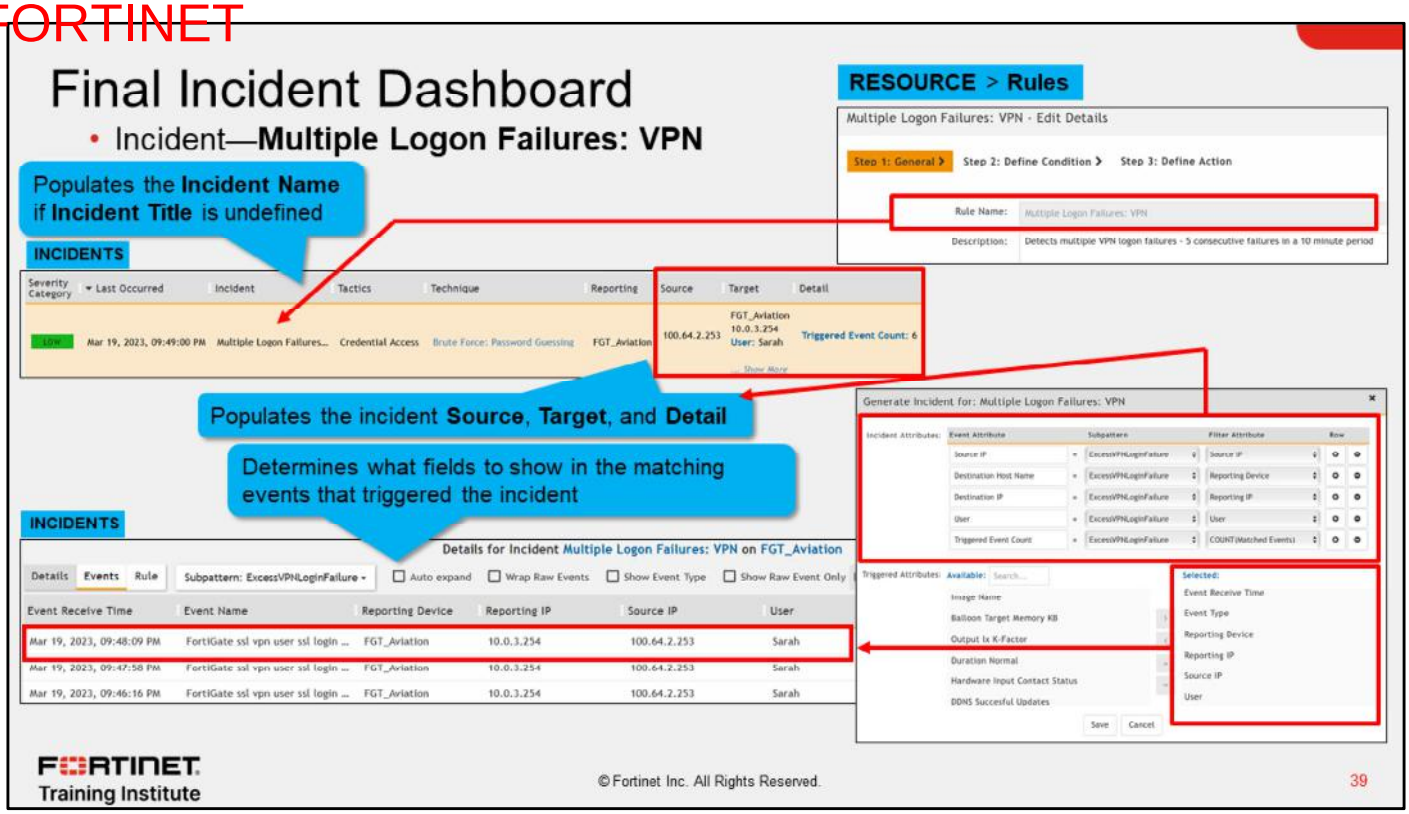
Any other event attributes, in this case **Triggered Event Count**, are displayed in the incident **Detail.**

DO NOT REPRINT
© FORTINET



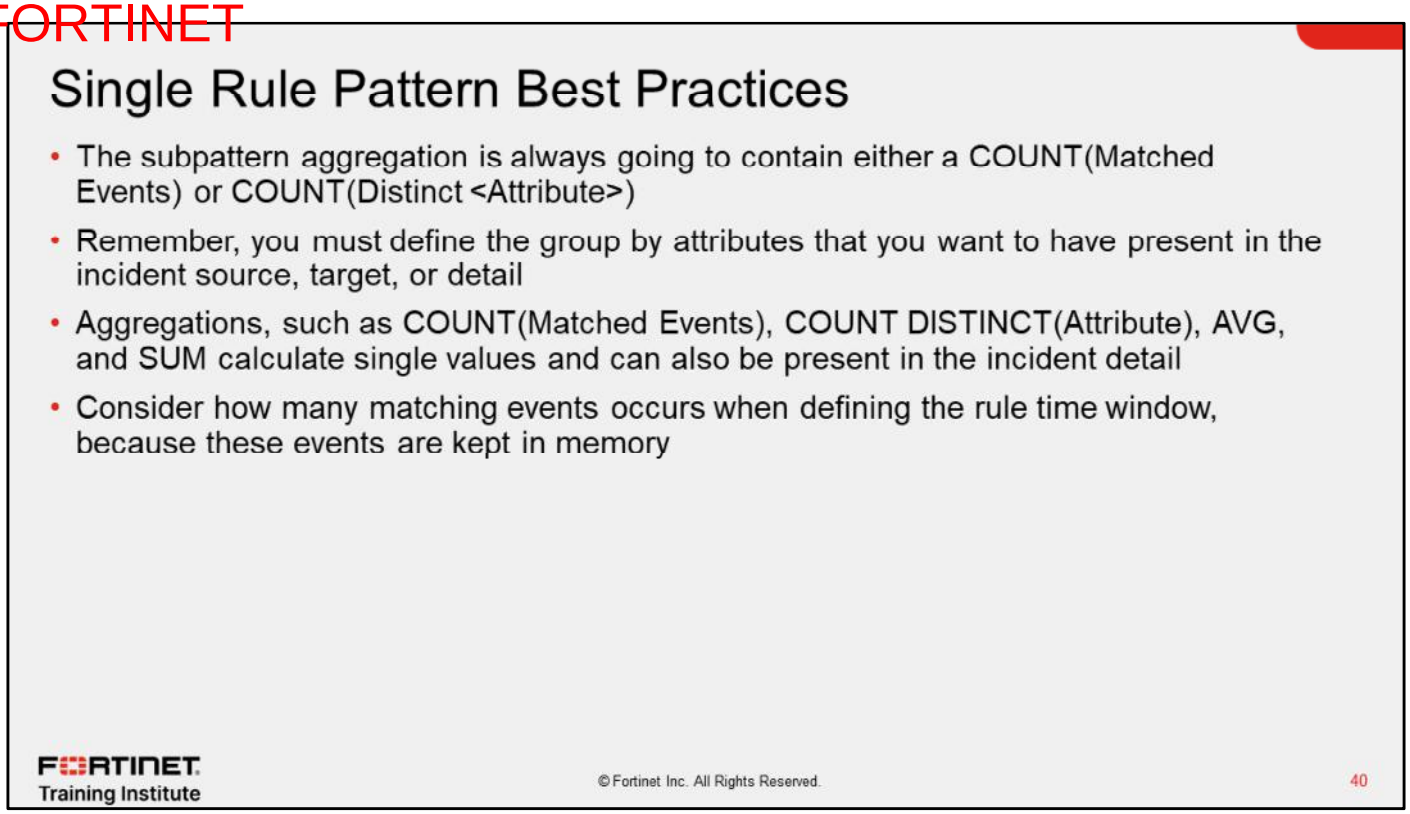
When you click the **Events** tab in the incident, you can see the attributes that will appear for matching events for that incident. These are the attributes you selected in the rule when you configured the action. **Event Type** displays **Event Name** and provides an option for **Event Type** called **Show Event Type**.

DO NOT REPRINT
© FORTINET



This slide summarizes the multiple VPN logon failures incident. The incident name comes from the rule name. The **Incident Attributes** populates the incident **Source**, **Target**, and **Detail** information. The **Triggered Attributes** determine the fields to show in the matching events that triggered the incident. By default the incident name is the name of the rule that triggered that incident, but that can be overwritten by configuring a name for the incident through the **Incident Title** field on the rule.

DO NOT REPRINT
© FORTINET

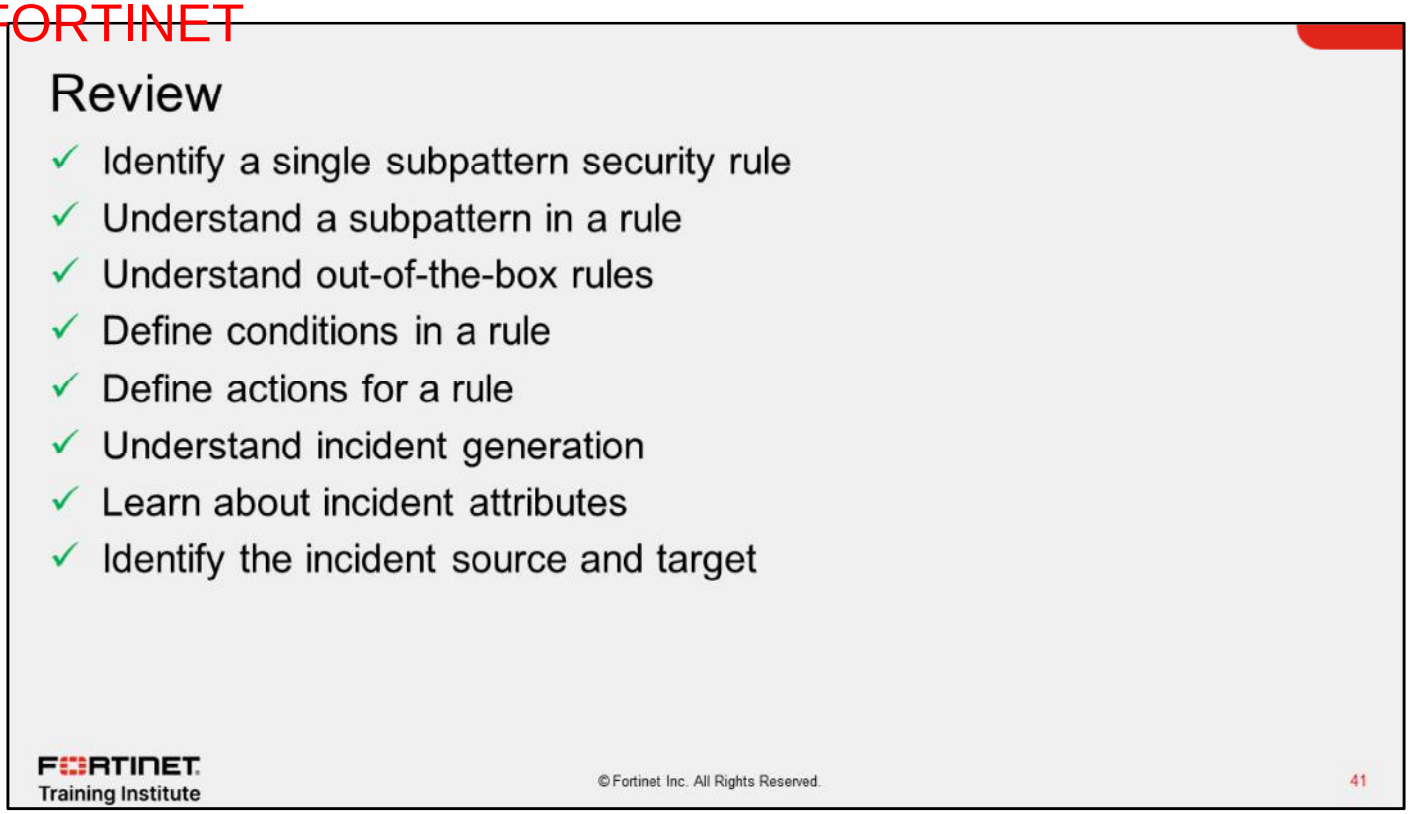


This slide lists some single rule pattern best practices.

The subpattern aggregation contains either a matched event count or distinct attribute count. Remember that you must define the group by attributes that you want to see present in the incident source, target, or detail. Aggregations, such as matched event counts, distinct attribute counts, average, sum, and so on, calculate single values and can also be present in the incident detail.

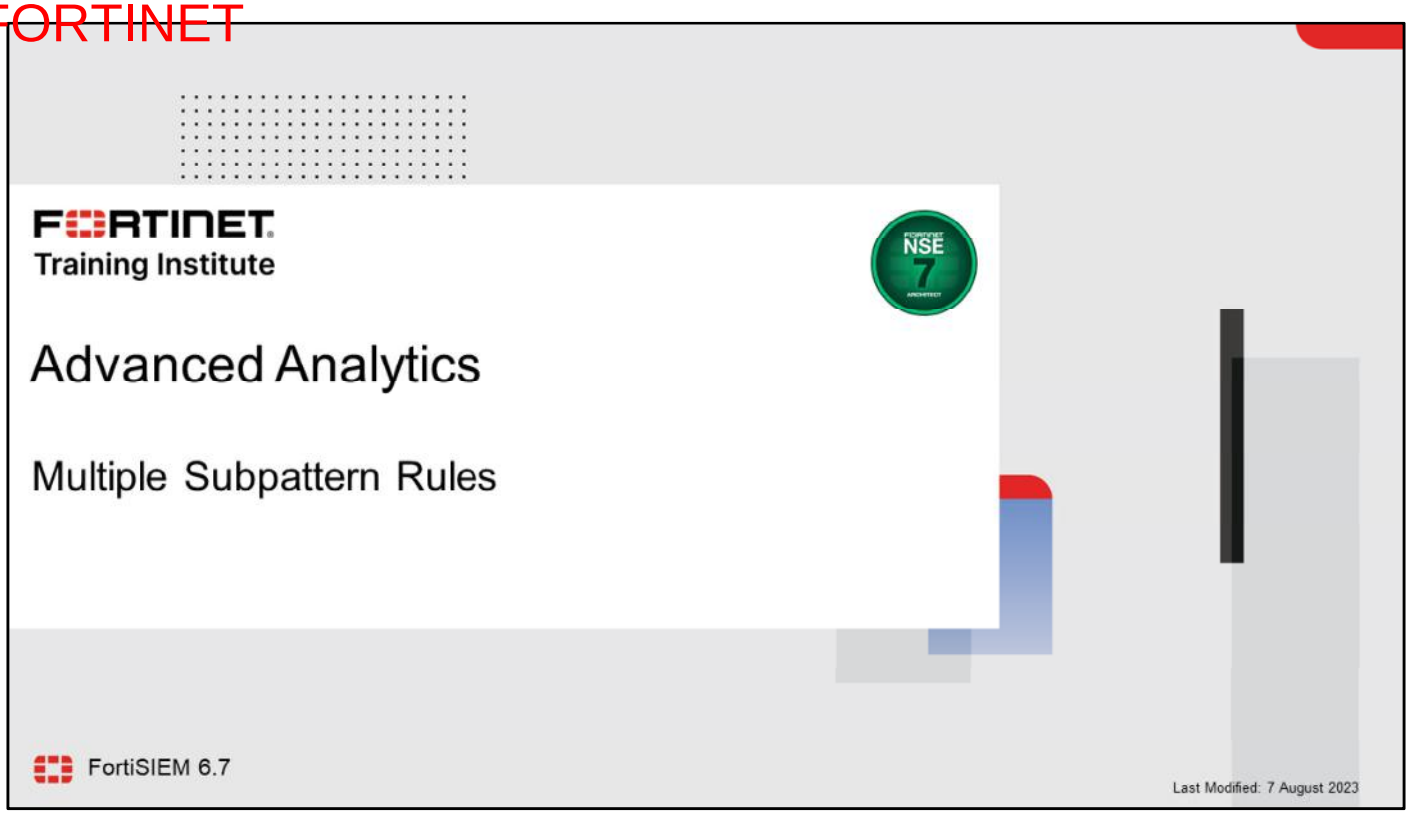
Consider how many matching events occur when defining the rule time window, because these events are stored in memory.

DO NOT REPRINT
© FORTINET



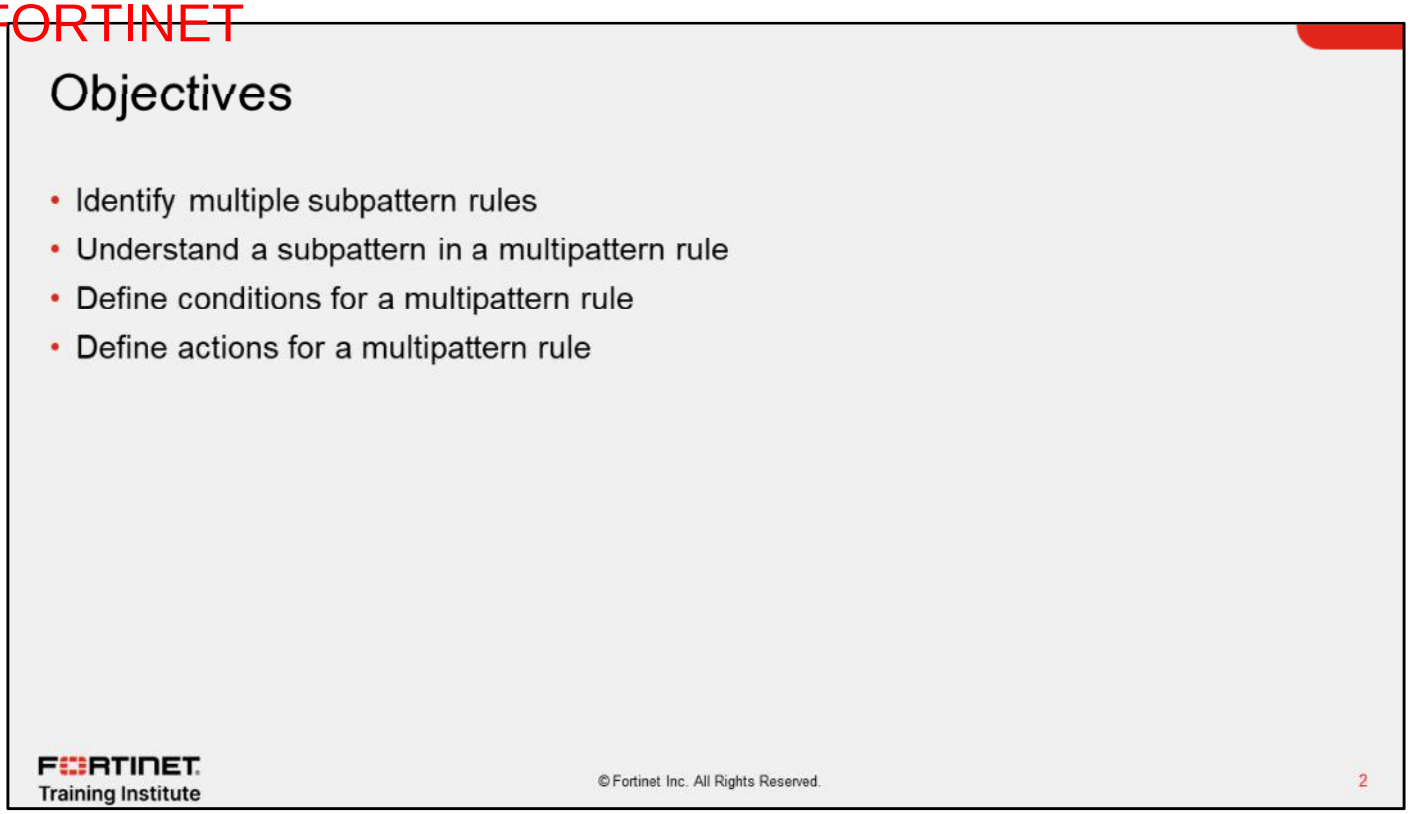
By mastering the objectives covered in this lesson, you learned about the single subpattern security rule. You also learned about incident generation and the attributes that are involved in incident generation.

DO NOT REPRINT
© FORTINET



In this lesson, you will learn about multiple subpattern rules, and the procedure to construct a multiple subpattern rule.

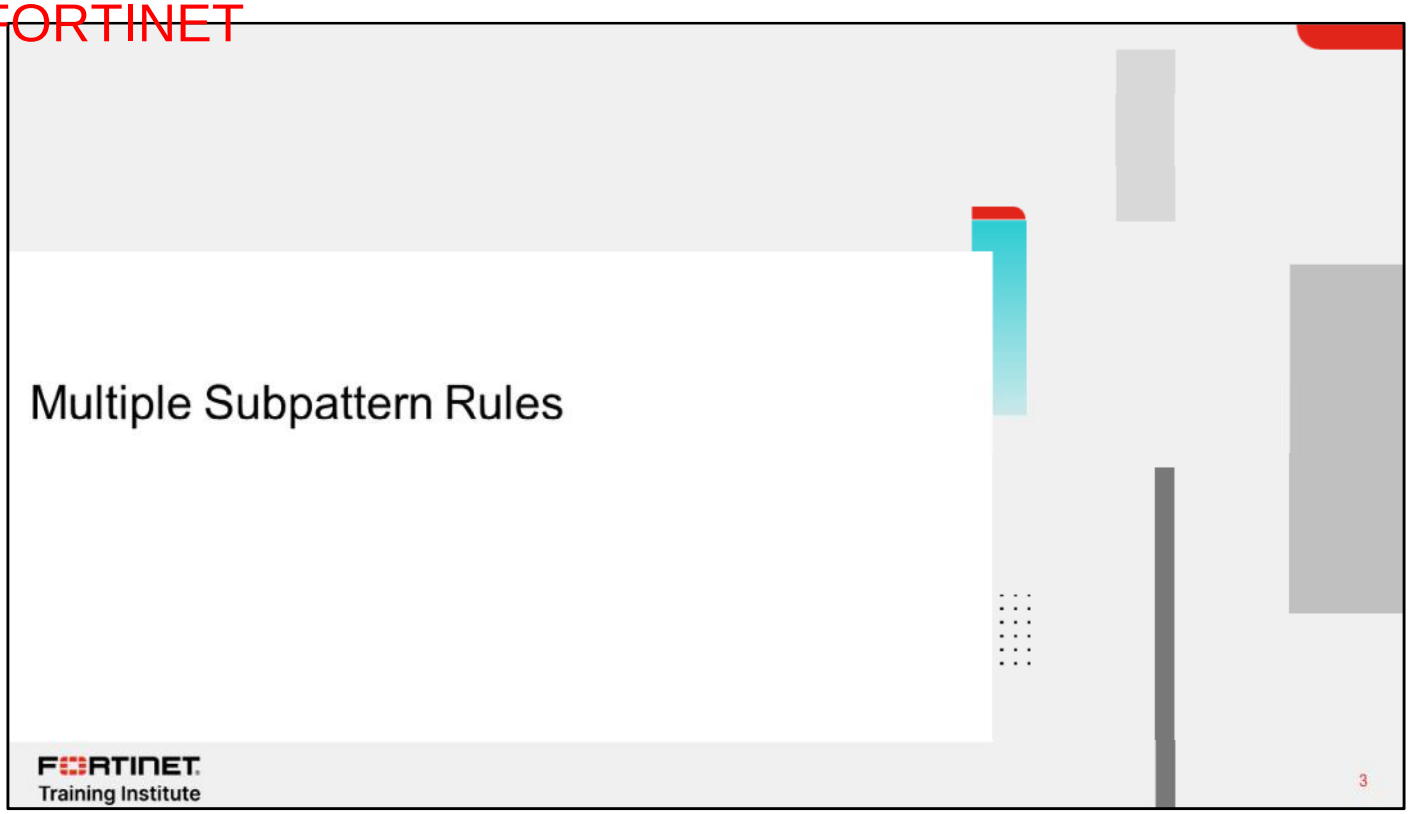
DO NOT REPRINT
© FORTINET



After completing this lesson, you should be able to achieve the objectives shown on this slide.

By demonstrating competence in identifying multiple subpattern rules, you will be able to define conditions and take action in a multipattern rule.

DO NOT REPRINT
© FORTINET



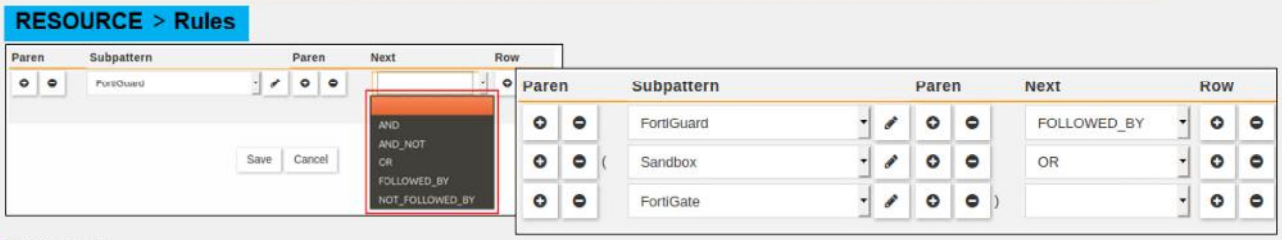
In this section, you will learn about event filters, group by, and aggregation conditions required in a multiple subpattern rule.

DO NOT REPRINT
© FORTINET

Next Operator for Multiple Subpattern Rules

- Multiple subpattern rules cover conditions where two patterns may need to occur within a specific time period, or one of a selection of patterns needs to occur to prove an incident condition exists

Operator	Description
AND	Subpattern X AND SubPattern Y occurred within the Time Window
AND_NOT	Subpattern X AND NOT SubPattern Y occurred within the Time Window
OR	Subpattern X OR SubPattern Y occurred within the Time Window
FOLLOWED_BY	Subpattern X was FOLLOWED BY SubPattern Y within the Time Window
NOT_FOLLOWED_BY	Subpattern X was NOT FOLLOWED BY SubPattern Y within the Time Window



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

4

FortiSIEM supports rules with multiple subpatterns. These cover conditions where two patterns may need to occur within a specific time period, or one of a selection of patterns needs to occur to prove an incident condition exists. If multiple patterns are used, then just like the analytics, you must specify a next operator. FortiSIEM supports five types of next operators, shown on this slide.

DO NOT REPRINT
© FORTINET

Out-of-the-Box Rules With Multiple Conditions

Rule	Rule
Multiple Logon Failures: Server	Transient Account Usage
Server Network Interface Staying Down	Brute Force Host Login Success
Network Device Interface Flapping	P2P Traffic Consuming High Network Bandwidth
Important Process Staying Down	Unusually Large ICMP Echo Packets
Multiple Logon Failures: Database	Excessive End User Mail to Unauthorized Mail Gateways
Suspicious Logon Failure: No Following Successful Login	Excessive End User DNS Queries
Successful Account Activity on a Disabled Account	Startup Config Change (With Login Info)
Failed Account Activity on Disabled Account	
Dormant Account Successful Logon	

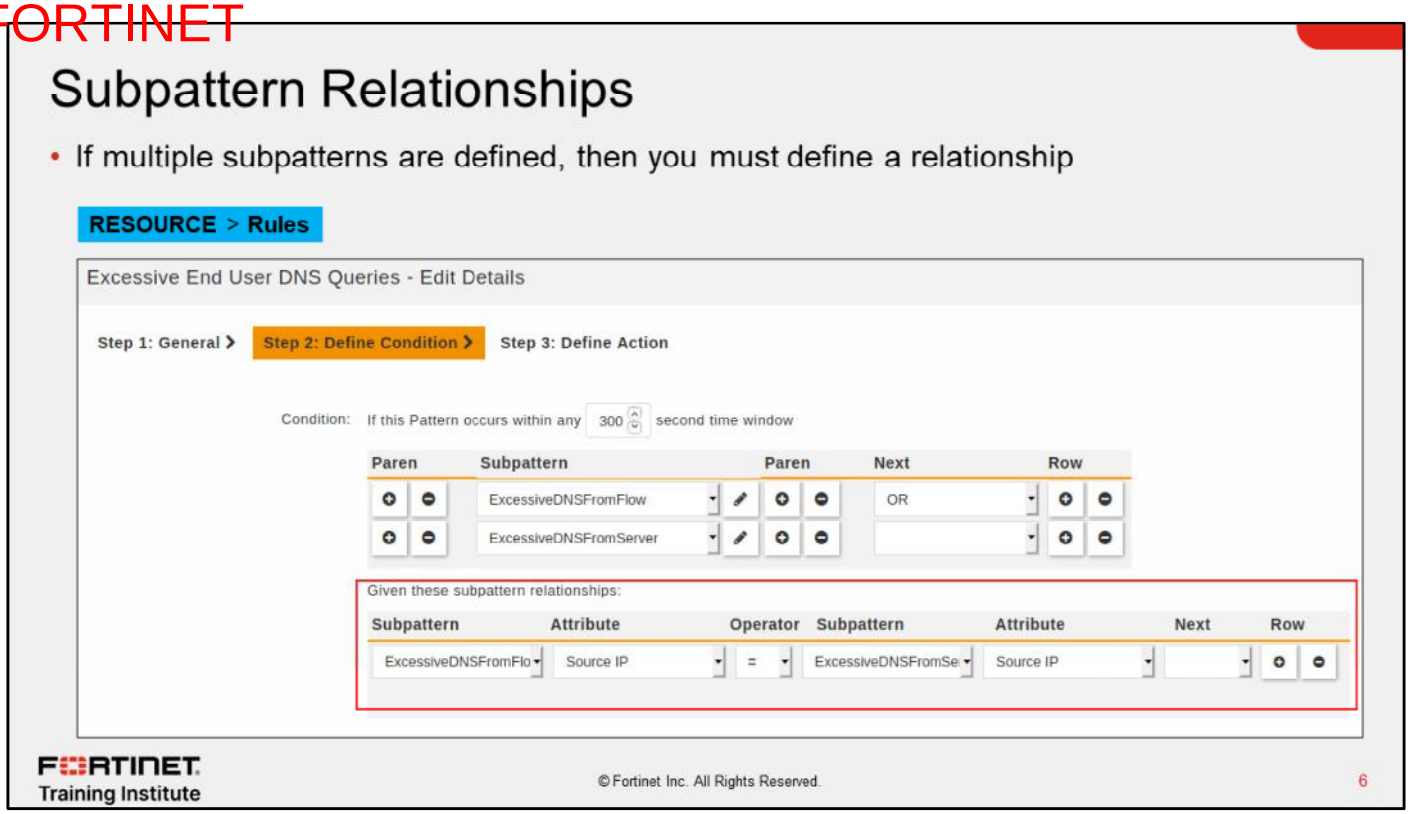
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

5

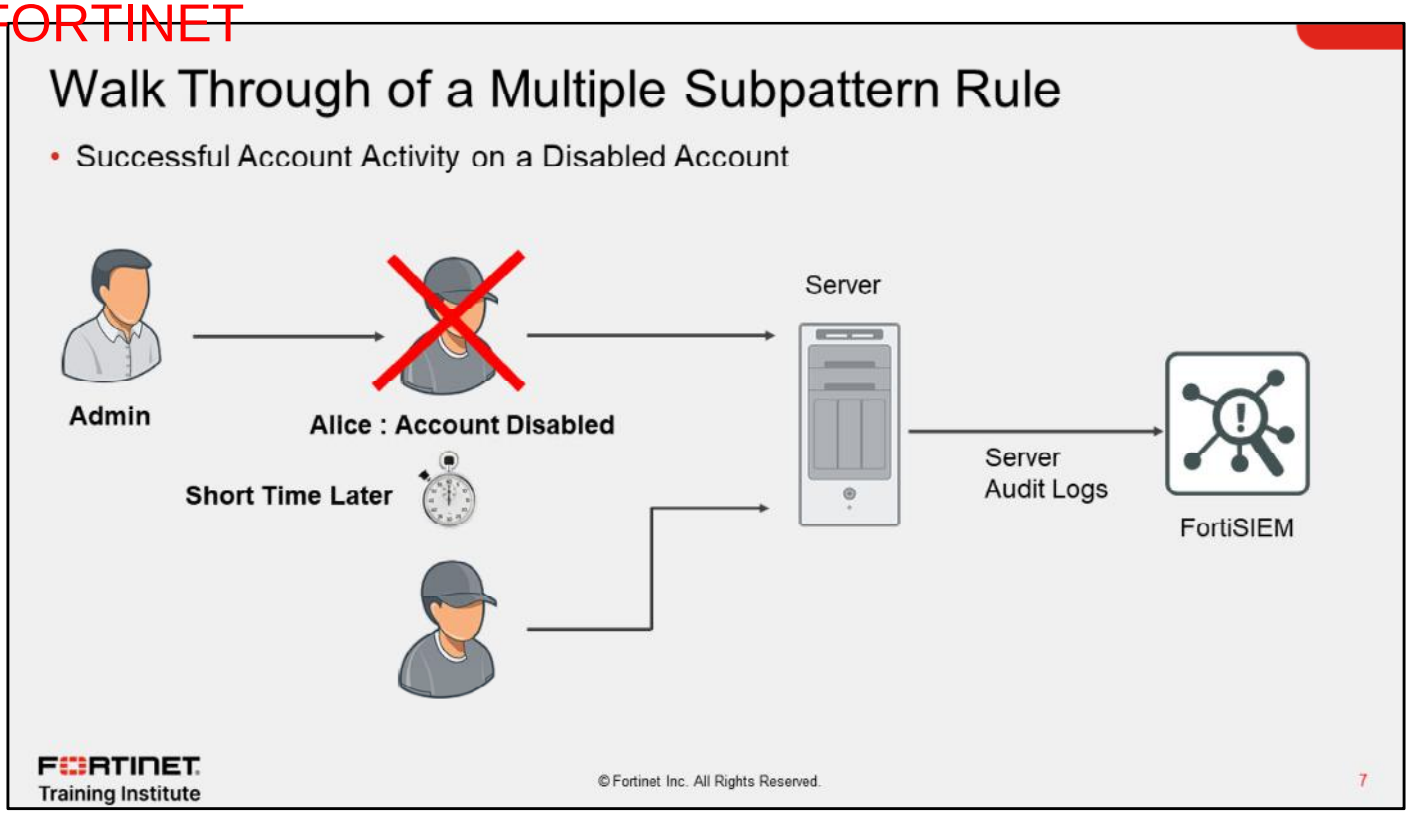
This slide highlights some out-of-the-box rules that have multiple conditions.

DO NOT REPRINT
© FORTINET



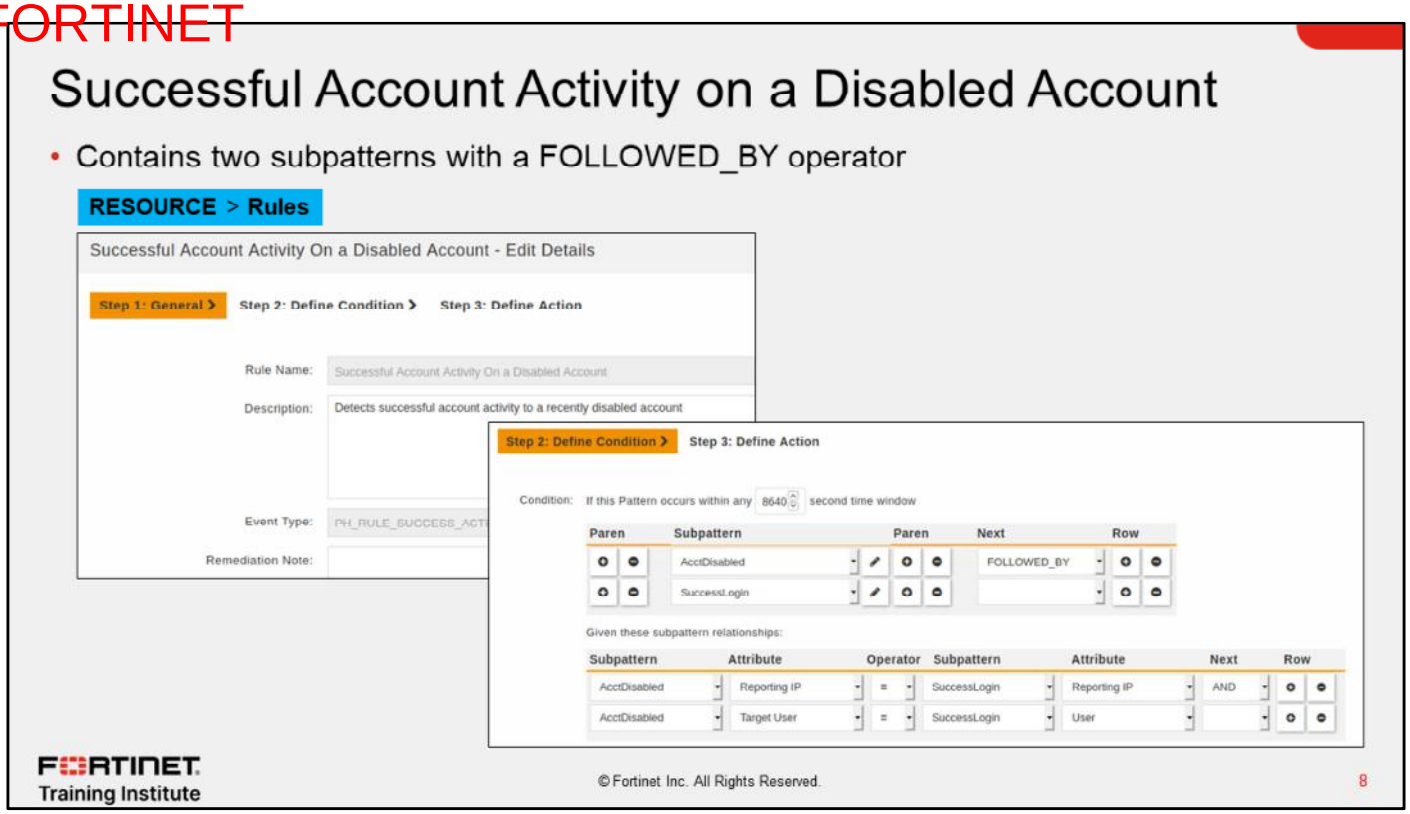
This slide shows the Excessive End User DNS Queries rule, which detects a scenario where a host that is not a DNS server, is sending excessive DNS requests. Authorized DNS servers are represented by the DNS server group. In a typical scenario, the frequency of end host DNS requests is not high unless there is a script running. This might indicate the presence of malware on the end host.

DO NOT REPRINT
© FORTINET



Successful Account Activity on a Disabled Account is a built-in rule that tracks user account login activity for a user whose account has recently been disabled by the administrator.

DO NOT REPRINT
© FORTINET

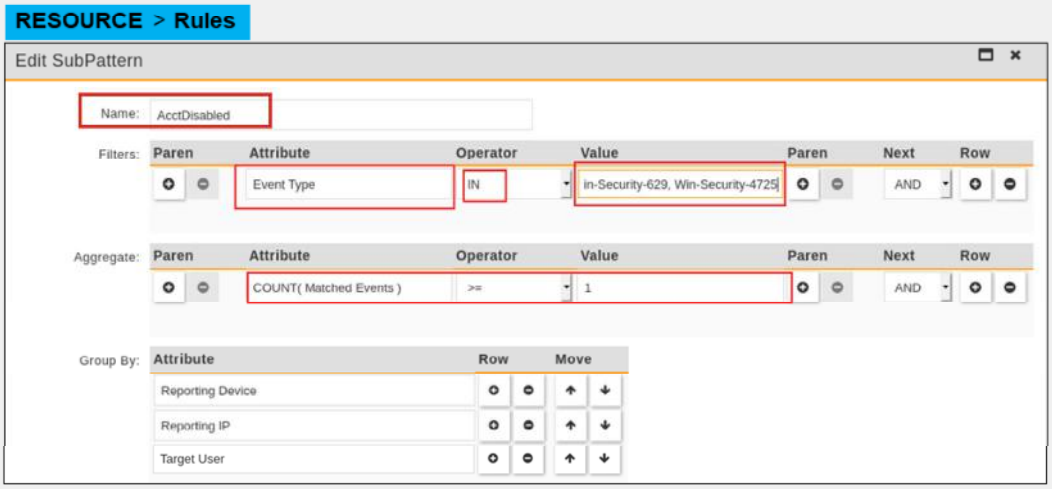


The rule contains two subpatterns—**AcctDisabled** and **SuccessLogin**—with a **FOLLOWED_BY** operator. The rule is tracked over a much larger time window of 86400 seconds, which is 24 hours.

DO NOT REPRINT
© FORTINET

Successful Account Activity on a Disabled Account (Contd)

- Tracks Windows events that relate to user accounts being disabled



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

9

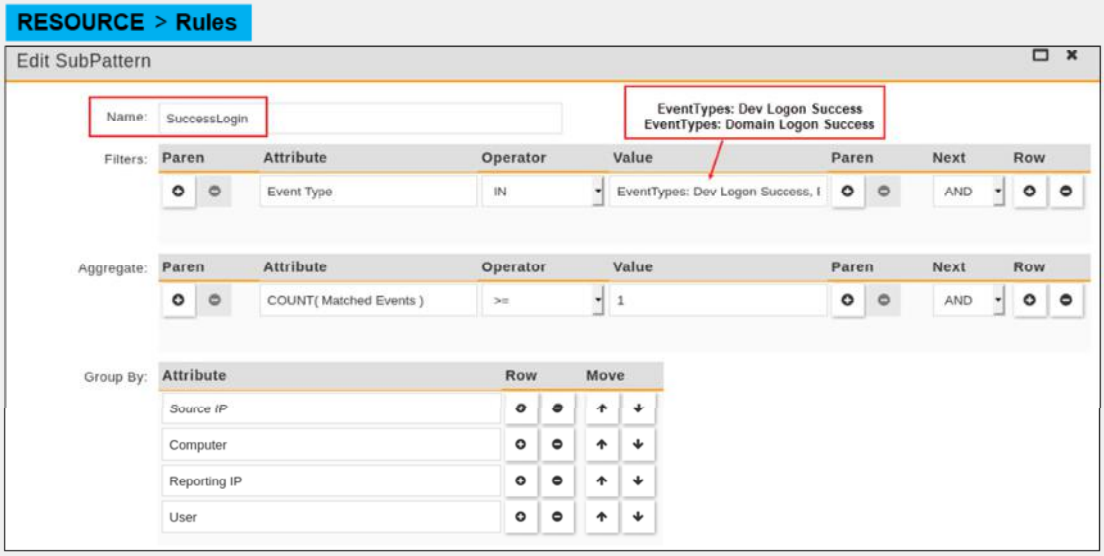
The **AcctDisabled** subpattern tracks Windows security events that have the IDs `629` and `4725`, which relate to user accounts being disabled.

After the events are grouped by **Reporting Device**, **Reporting IP**, and **Target User**, if one or more events match the aggregate condition then the subpattern is considered a match.

DO NOT REPRINT
© FORTINET

Successful Account Activity on a Disabled Account (Contd)

- Tracks device and domain login success events

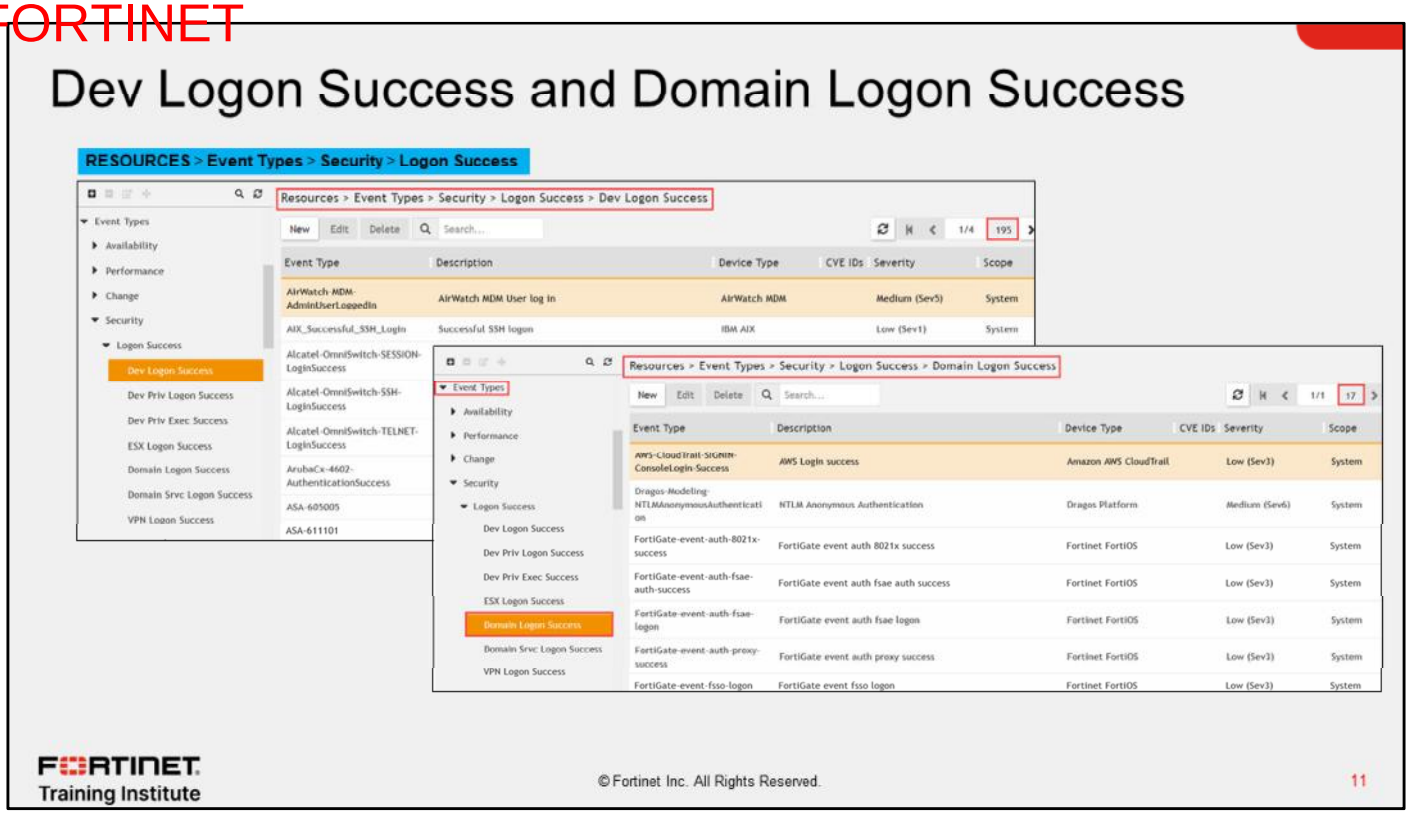


FORTINET Training Institute
© Fortinet Inc. All Rights Reserved.

The second subpattern in the rule is **SuccessLogin**, which tracks successful login event types from the **Dev Logon Success** and **Domain Logon Success** CMDB groups.

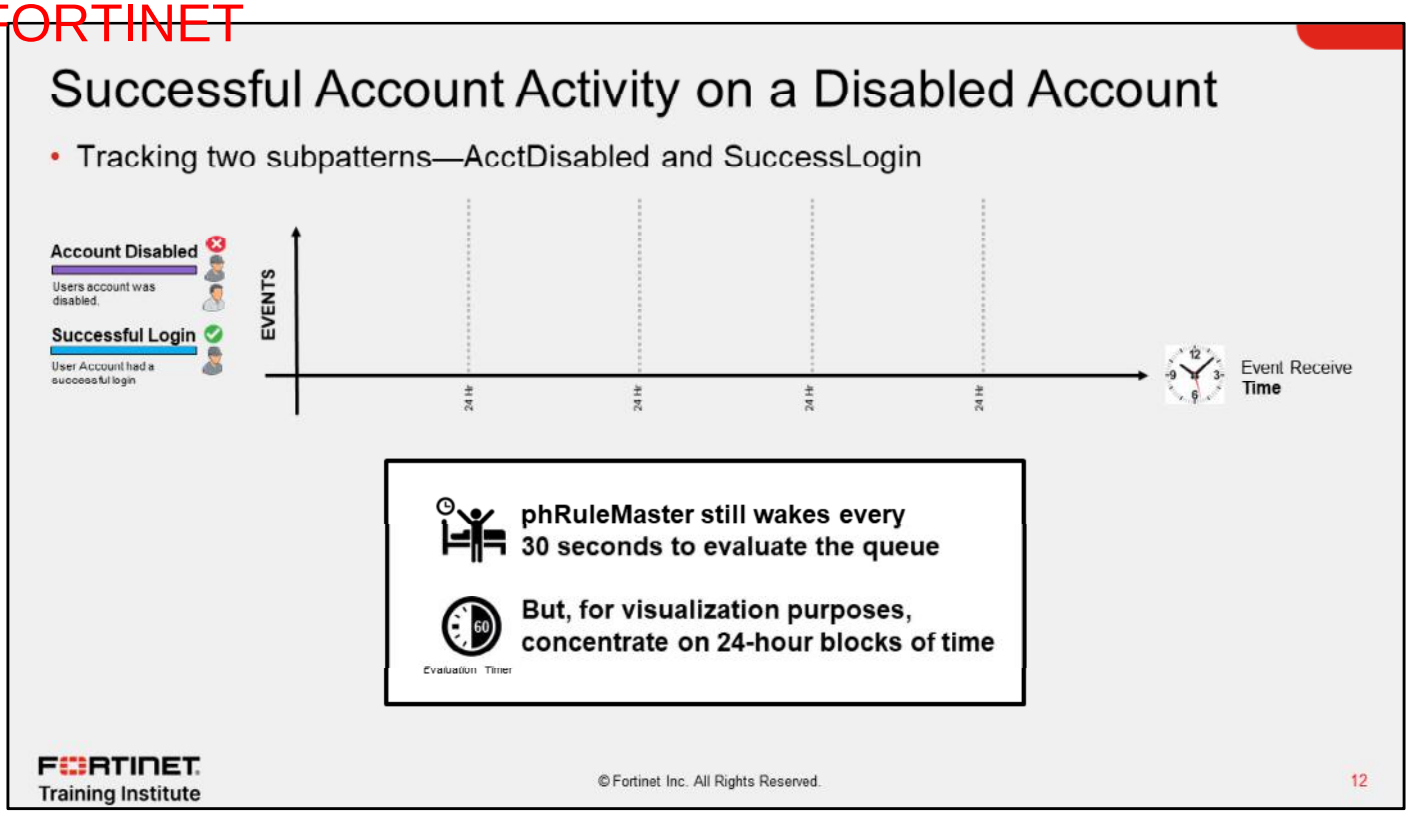
After the events are grouped by **Source IP**, **Computer**, **Reporting IP**, and **User**, if one or more events match the aggregate condition then the subpattern is considered a match.

DO NOT REPRINT
© FORTINET



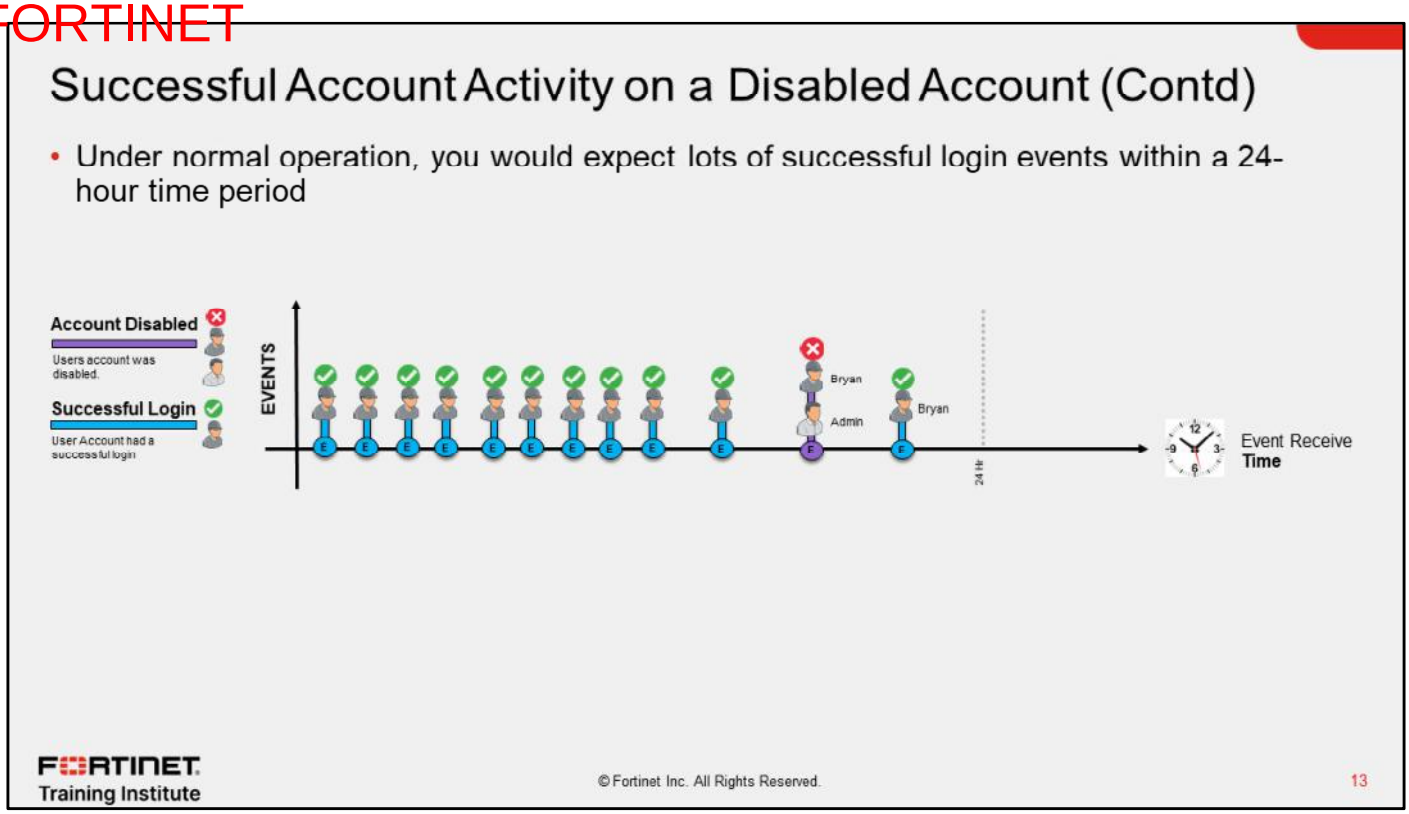
There are 195 different types of out-of-the-box device login success rules, and 10 types of domain login success rules. You can create your own event types under any category and reference them in a rule subpattern.

DO NOT REPRINT
© FORTINET



The phRuleMaster still wakes every 30 seconds to evaluate the queue. For visualization purposes, this lesson concentrates on a 24-hour block of time.

DO NOT REPRINT
© FORTINET



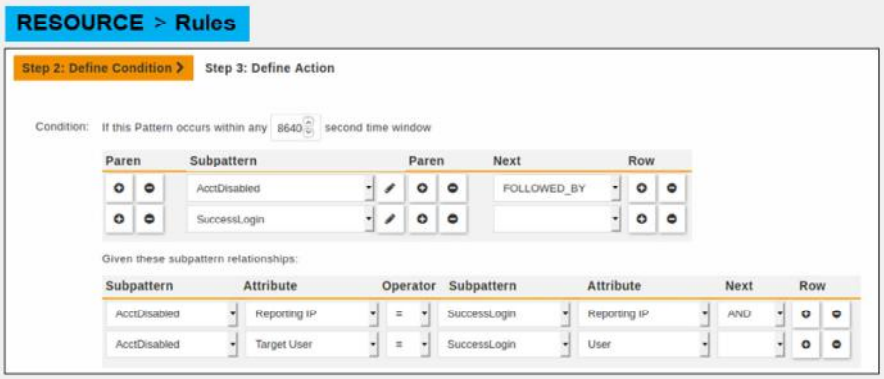
Under normal operation, you would expect lots of successful login events within a 24-hour time period, along with some odd account disabled events.

What would be unusual is an account being disabled then, a short time later, the same account being used for a successful login.

DO NOT REPRINT
© FORTINET

Rule—Subpattern Relationships

- **SubPattern** relationships are defined between the multiple patterns
- The **Target User** in the **AcctDisabled** subpattern has to be equal to the **User** in the **SuccessLogin** subpattern
- On the back end, you can also use the event receive time to identify the order in which the events were received



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

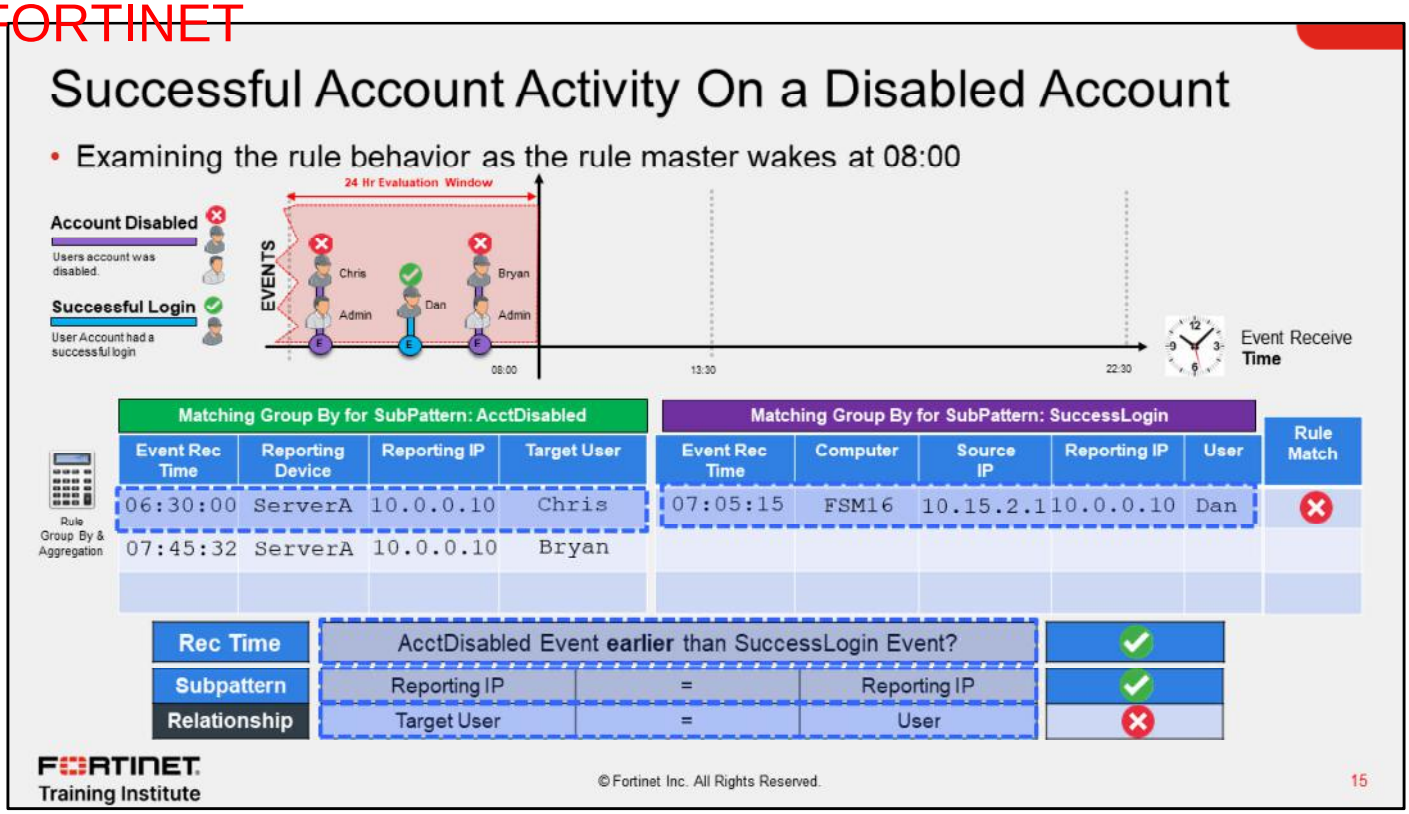
14

When there are multiple subpatterns, you may need to define a relationship between those subpatterns. In the example shown on this slide, the **Target User** in the **AcctDisabled** subpattern must be equal to the user in the **SuccessLogin** subpattern. Also, the **Reporting IP** in the **AcctDisabled** subpattern must be equal to the **Reporting IP** in the **SuccessLogin** subpattern.

If you do not define this relationship, FortiSIEM can generate false positive incidents. By defining this relationship, you are making a condition where the user who was actually disabled by the administrator is the same user who logged in successfully at a later time. Also, the same device that reported a disabled account event is also reporting the successful login event.

On the back end, you can also use the event receive time to identify the order in which the events were received. This is how FortiSIEM knows if the disabled account event occurred before the successful event.

DO NOT REPRINT
© FORTINET



This slide shows two tables—one for the AcctDisabled subpattern and the other for the SuccessLogin subpattern. The columns in the table are named according to the group by definition for each subpattern. The group by definition for the AcctDisabled subpattern is reporting device, reporting IP, and target user. The group by definition for the SuccessLogin subpattern is source IP, computer, reporting IP, and user.

The phRuleMaster evaluates the events for a 24-hour period. The AcctDisabled subpattern tracked two events in which the administrator had disabled the user accounts for Chris and Bryan. The SuccessLogin subpattern tracked one event for Dan, who had logged in successfully. The rule compares the disabled account for Chris and the successful login event for Dan. FortiSIEM goes through the analysis of the events and decides if it should trigger an incident:

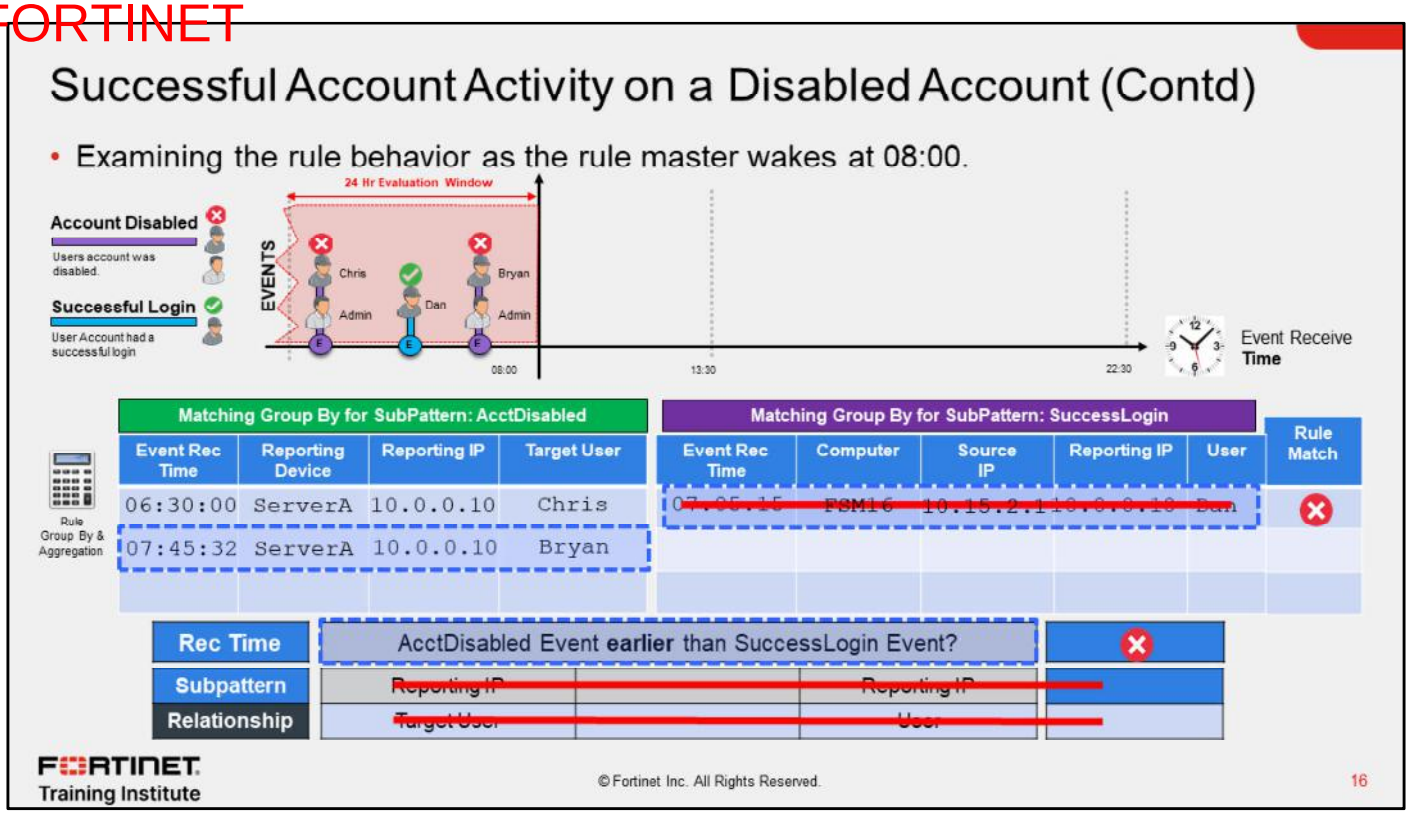
- Was the AcctDisabled event earlier than the SuccessLogin event? The answer is yes.
- Is the reporting IP of the AcctDisabled event the same as the reporting IP of the SuccessLogin event? The answer is yes.
- Is the target user of the AcctDisabled event the same as the user of the SuccessLogin event? The answer is no.

Therefore, no incident is triggered.

Now, the rule compares Bryan's disabled user account with the successful login event for Dan:
Was the **AcctDisabled** event earlier than the **SuccessLogin** event? The answer is no.

At this point, the rule does not need to process the rule relationship condition because the disabled login event was received after the successful login event. During this evaluation window, no incident is triggered because the only user who logged in successfully is Dan, and his account was not disabled by the administrator.

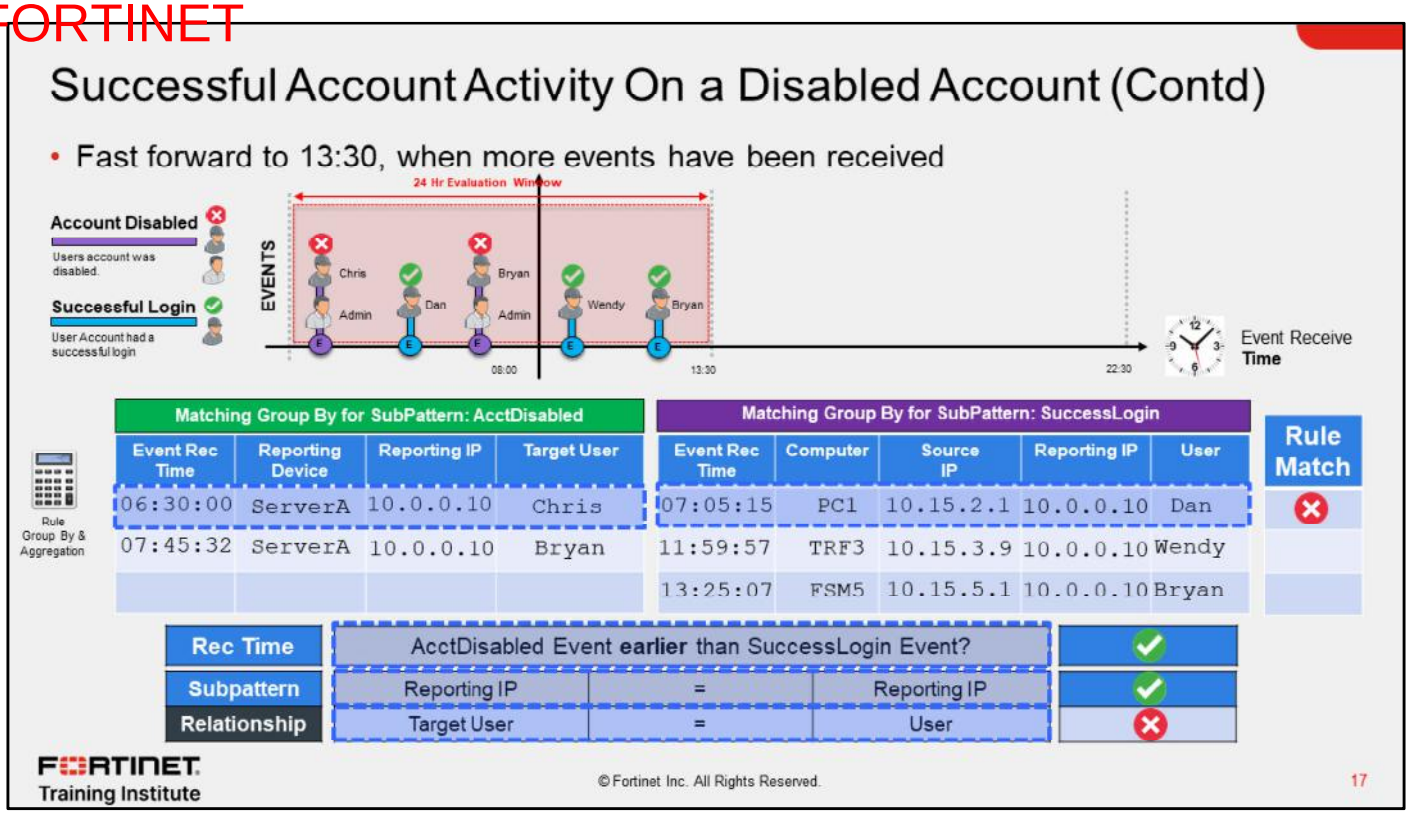
DO NOT REPRINT
© FORTINET



Now, the rule compares the disabled user account for Bryan and the successful login event for Dan: Was the **AcctDisabled** event earlier than the **SuccessLogin** event? The answer is no.

At this point, the rule does not need to process the rule relationship condition because the disabled login event was received after the successful login event. During this evaluation window, no incident is triggered because the only user who logged in successfully is Dan, and his account was not disabled by the administrator.

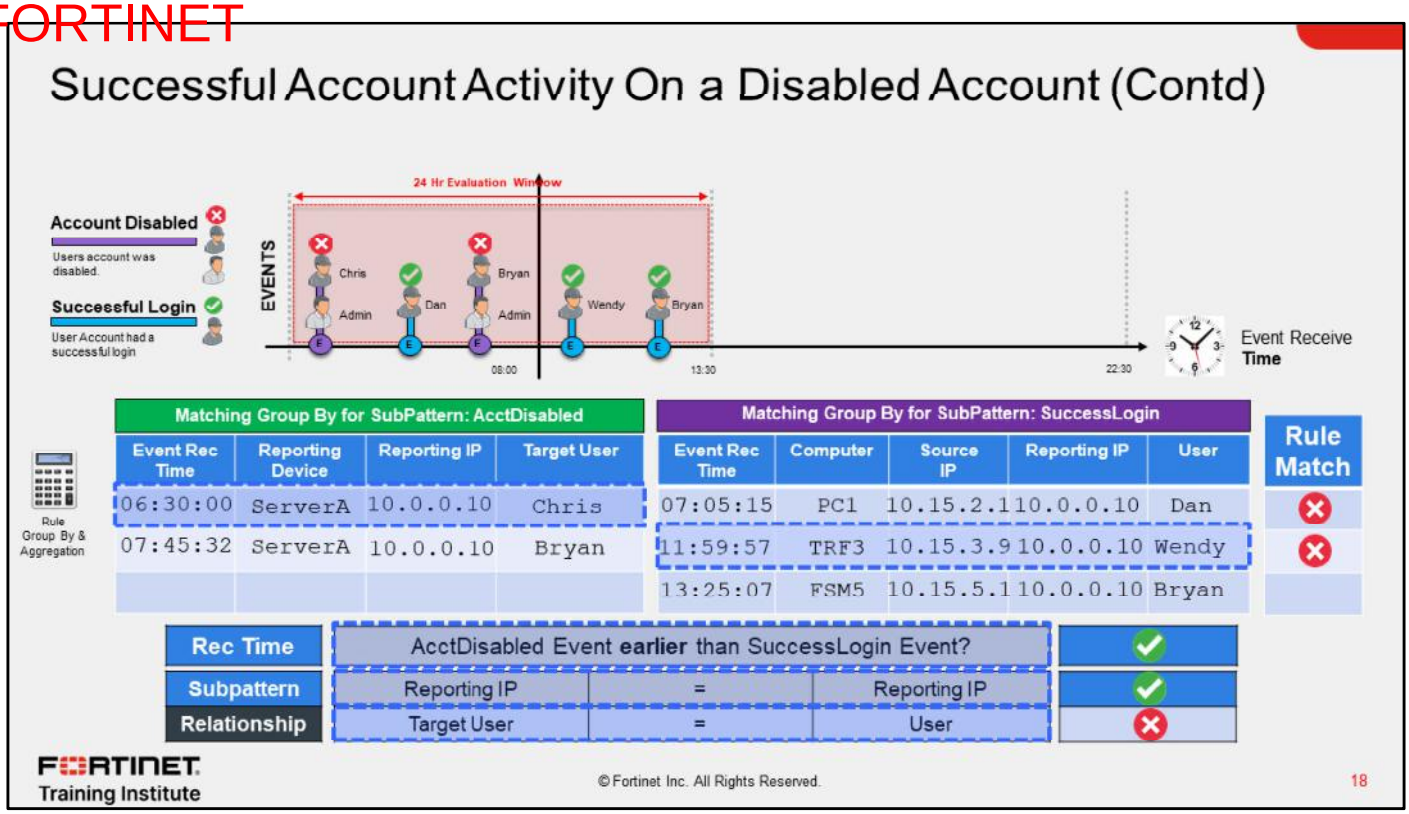
DO NOT REPRINT
© FORTINET



Now, fast forward to 13:30, when more events have been received. Users Wendy and Bryan had logged in successfully since the last evaluation.

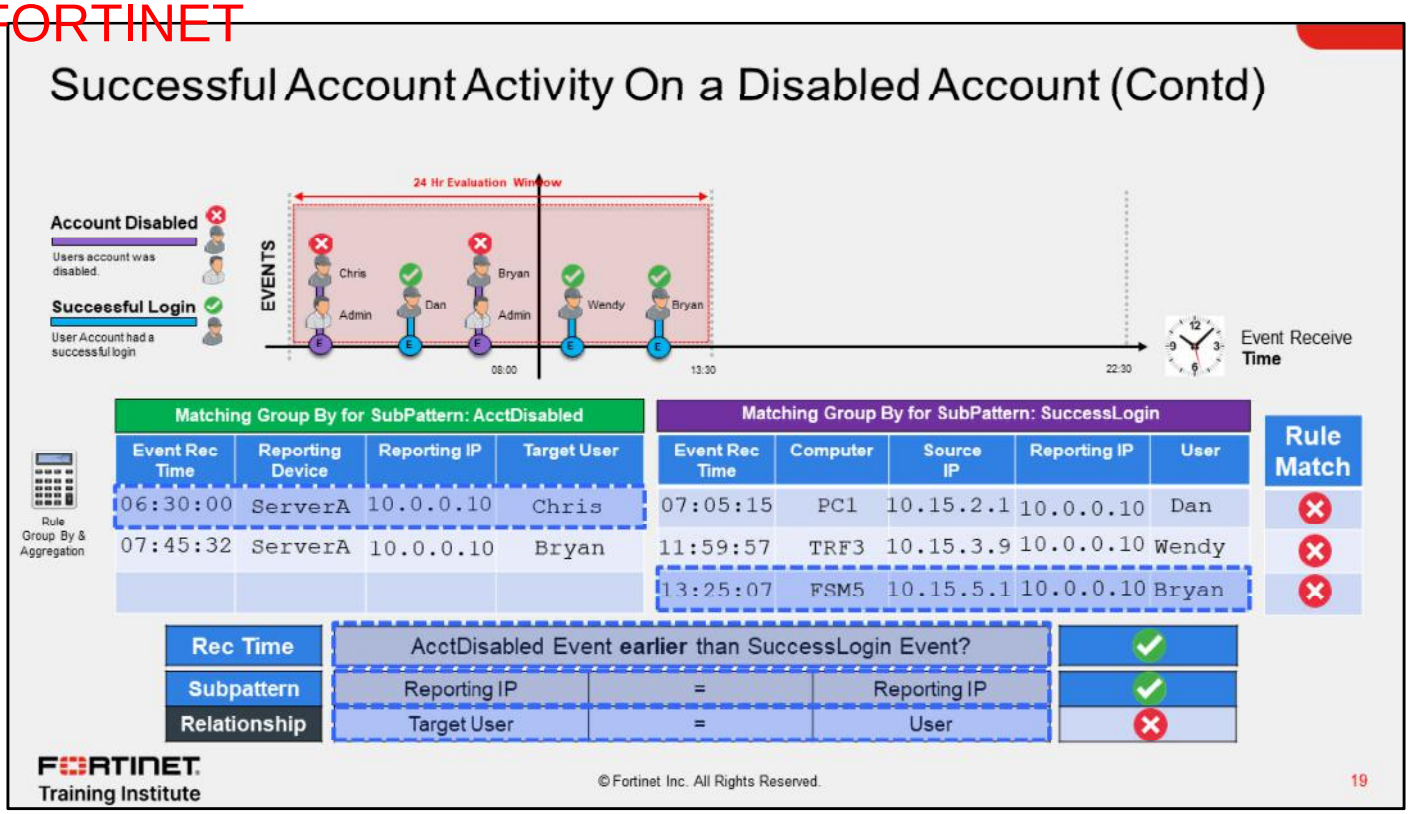
The rule compares Chris' disabled user account and the successful login event for Dan. No incident is triggered because the target user from the AcctDisabled event is not the same as the user from the SuccessLogin event.

DO NOT REPRINT
© FORTINET



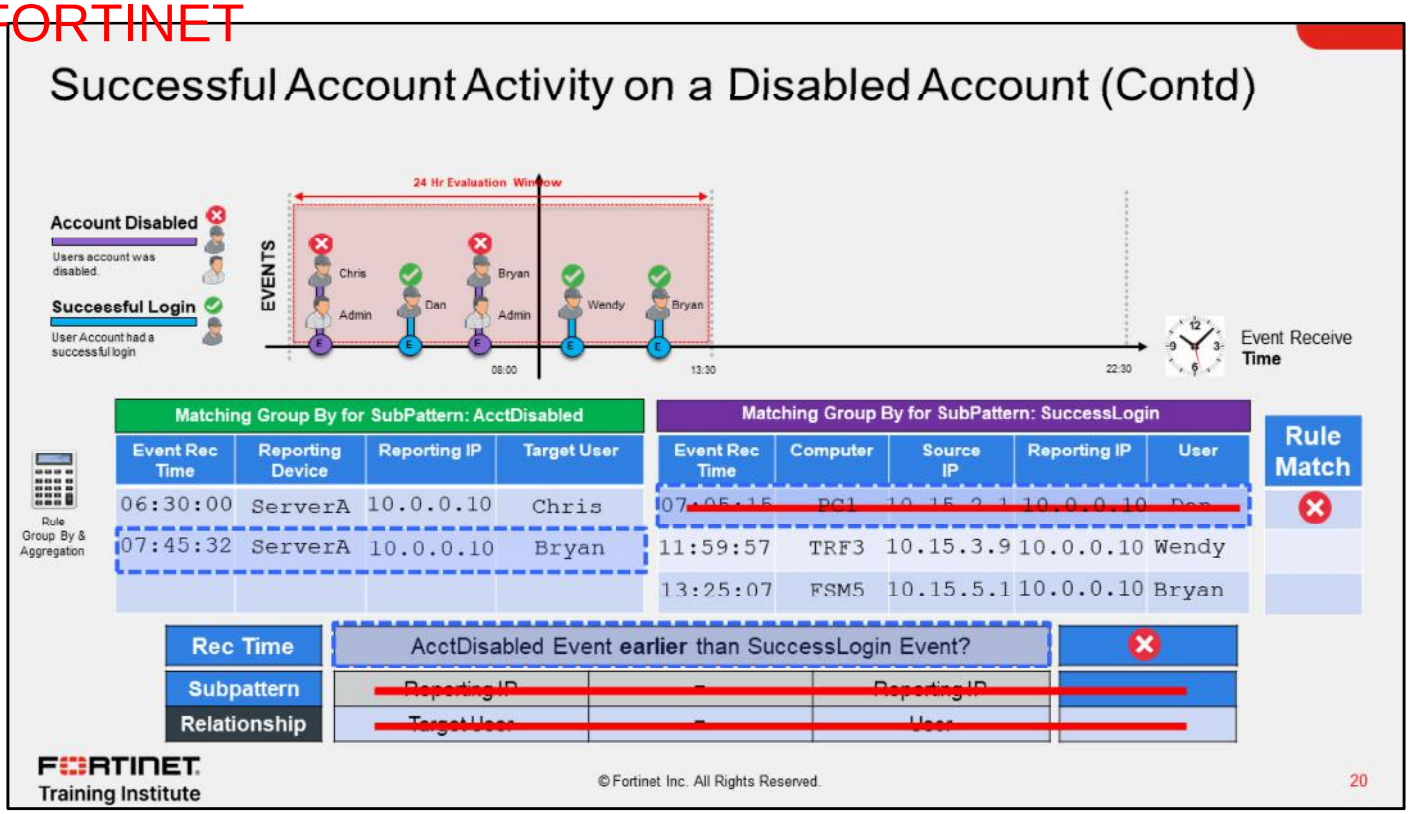
The rule compares the disabled user account for Chris and the successful login event for Wendy. No incident is triggered because the target user from the AcctDisabled event is not the same as the user from the SuccessLogin event.

DO NOT REPRINT
© FORTINET



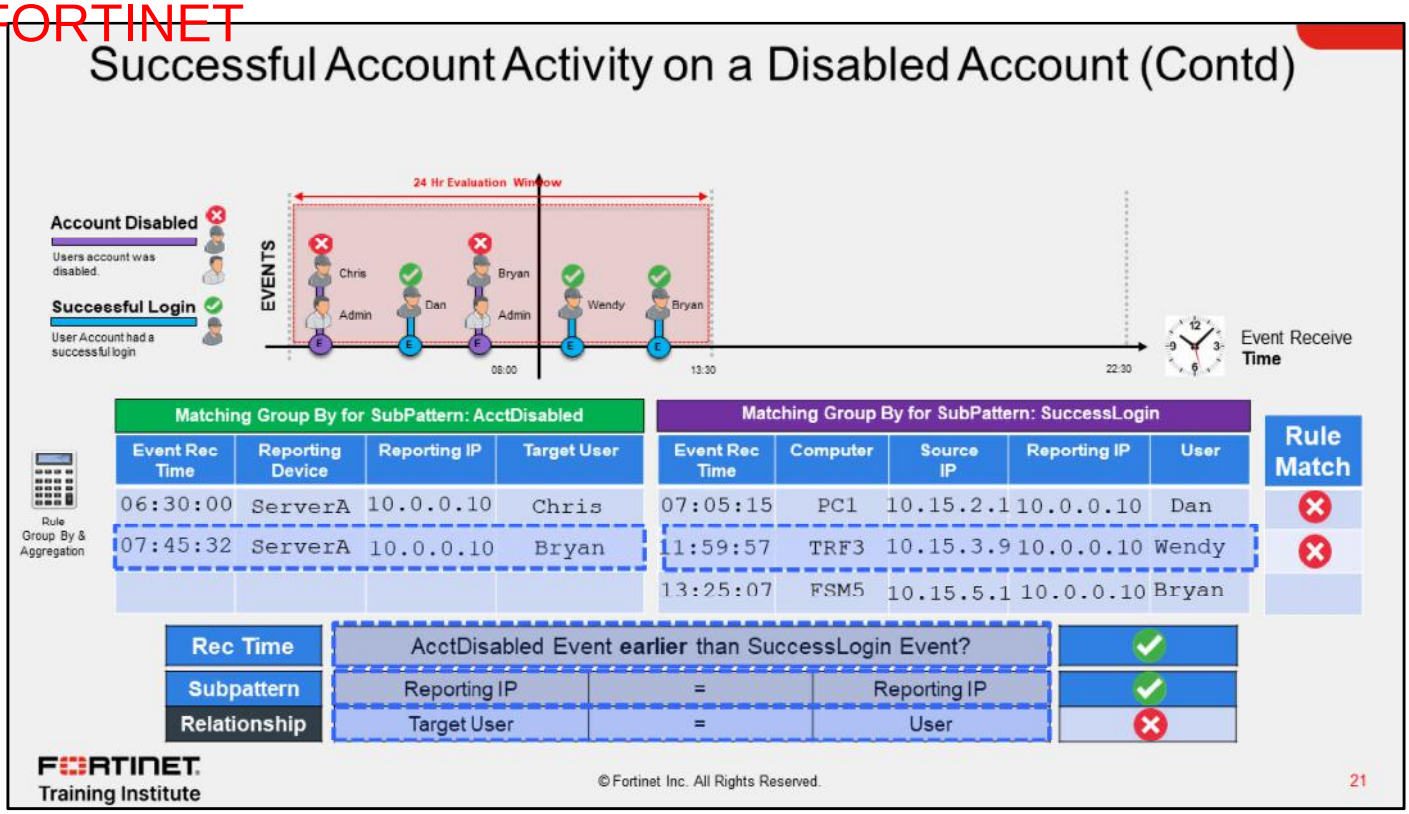
The rule compares the disabled user account for Chris and the successful login event for Bryan. No incident is triggered because the target user from the AcctDisabled event is not same as the user from the SuccessLogin event.

DO NOT REPRINT
© FORTINET



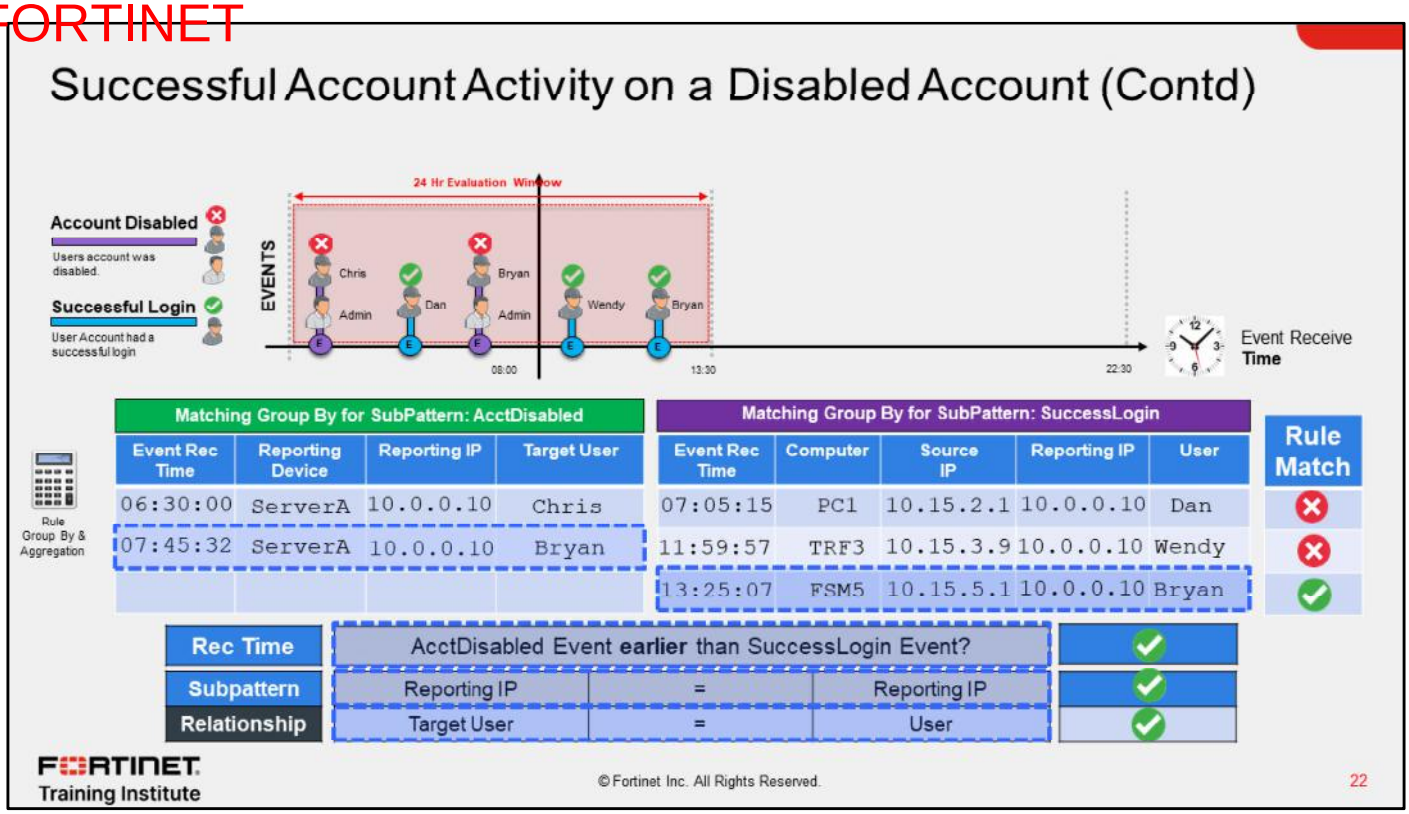
The rule compares the disabled user account for Bryan and the successful login event for Dan. No incident is triggered because the AcctDisabled event occurred earlier than the SuccessLogin event.

DO NOT REPRINT
© FORTINET



The rule compares the disabled account for Bryan and the successful login event for Wendy. No incident is triggered because the target user from the AcctDisabled event is not the same as the user from the SuccessLogin event.

DO NOT REPRINT
© FORTINET



Finally, the rule compares the disabled user account for Bryan and the successful login event for Bryan. An incident is triggered because Bryan logged in after his account was disabled. It is the same device that reported Bryan's disabled account event and his successful login event.

DO NOT REPRINT
© FORTINET

Generate Incident

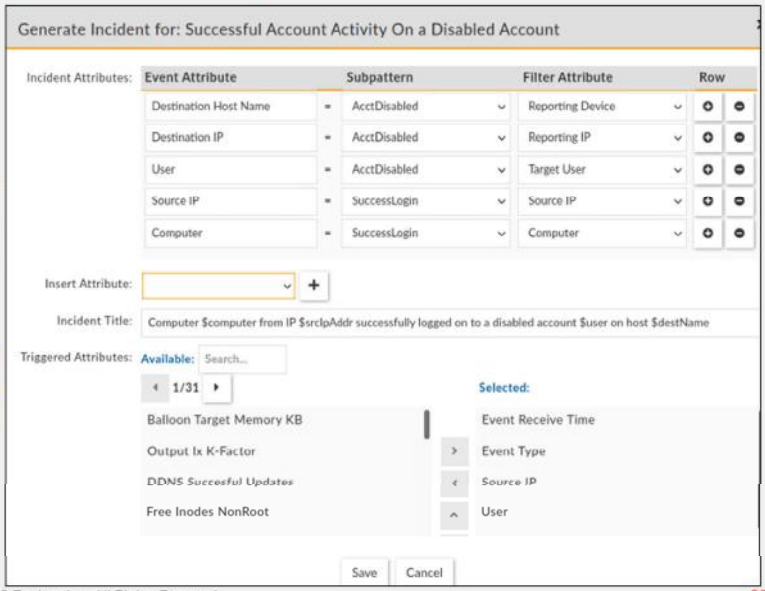
- Now you can set the incident attributes
- You can select them from either subpattern or both

Subpattern	Reporting IP	Target User
AcctDisabled	10.0.0.10	Bryan

Subpattern	Reporting IP	User
SuccessLogin	10.0.0.10	Bryan

FORTINET Training Institute

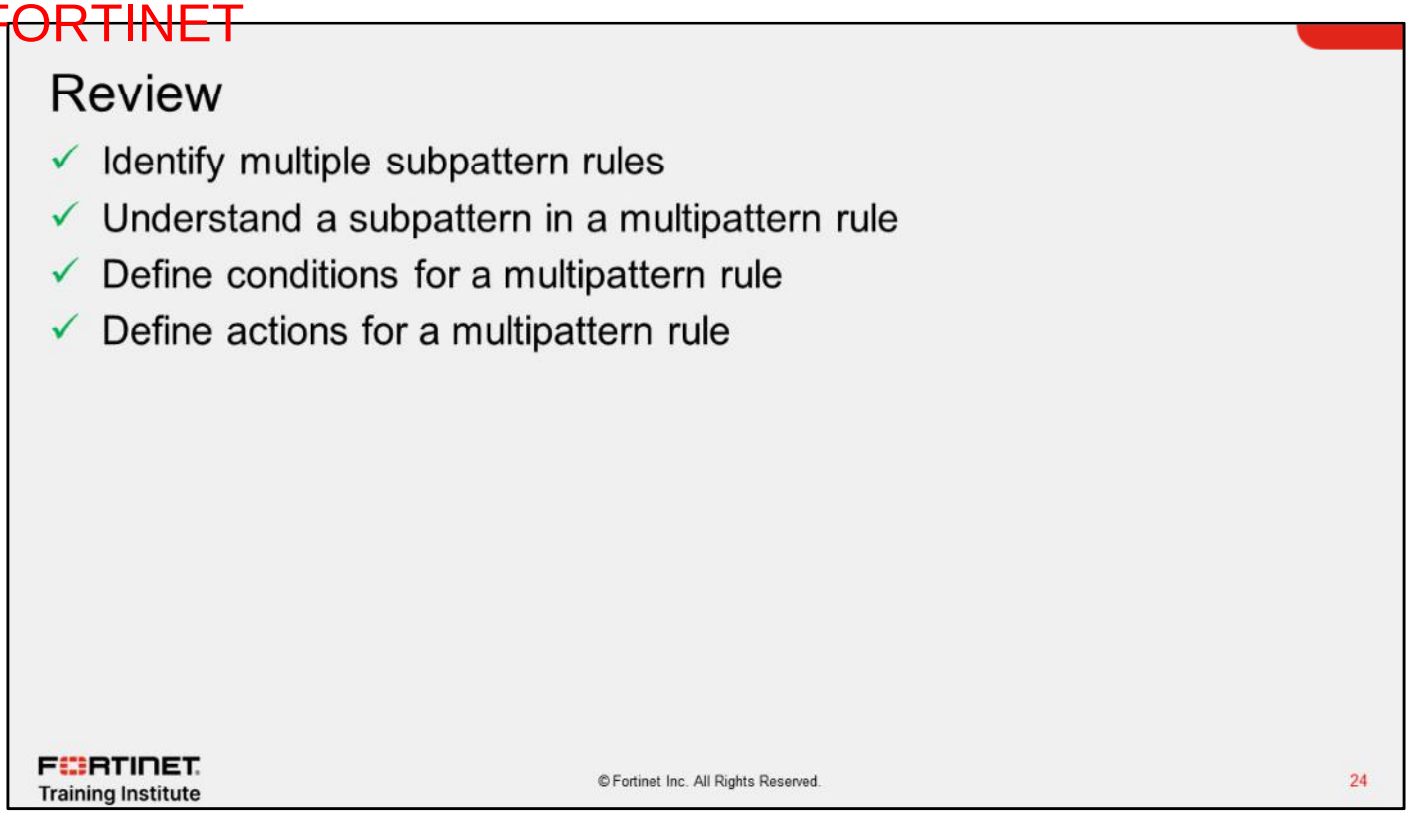
RESOURCE > Rules



© Fortinet Inc. All Rights Reserved. 23

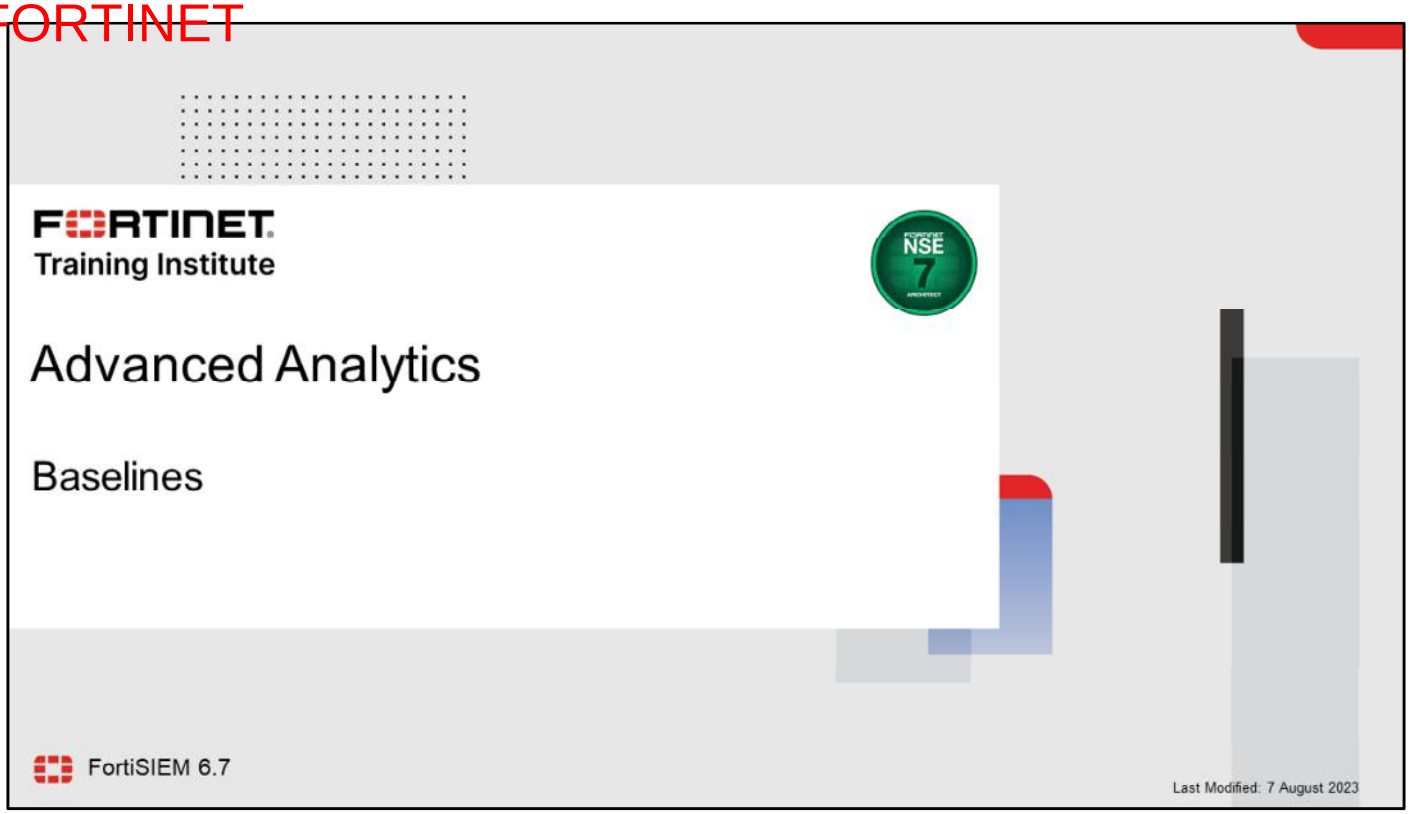
You can select the incident attributes from either subpattern or both. Each matching event returns the same attributes on the **Incident** tab for each subpattern. Therefore, some fields may be empty. You can customize the incident title by inserting attributes from the **Insert Attributes** drop-down list.

DO NOT REPRINT
© FORTINET



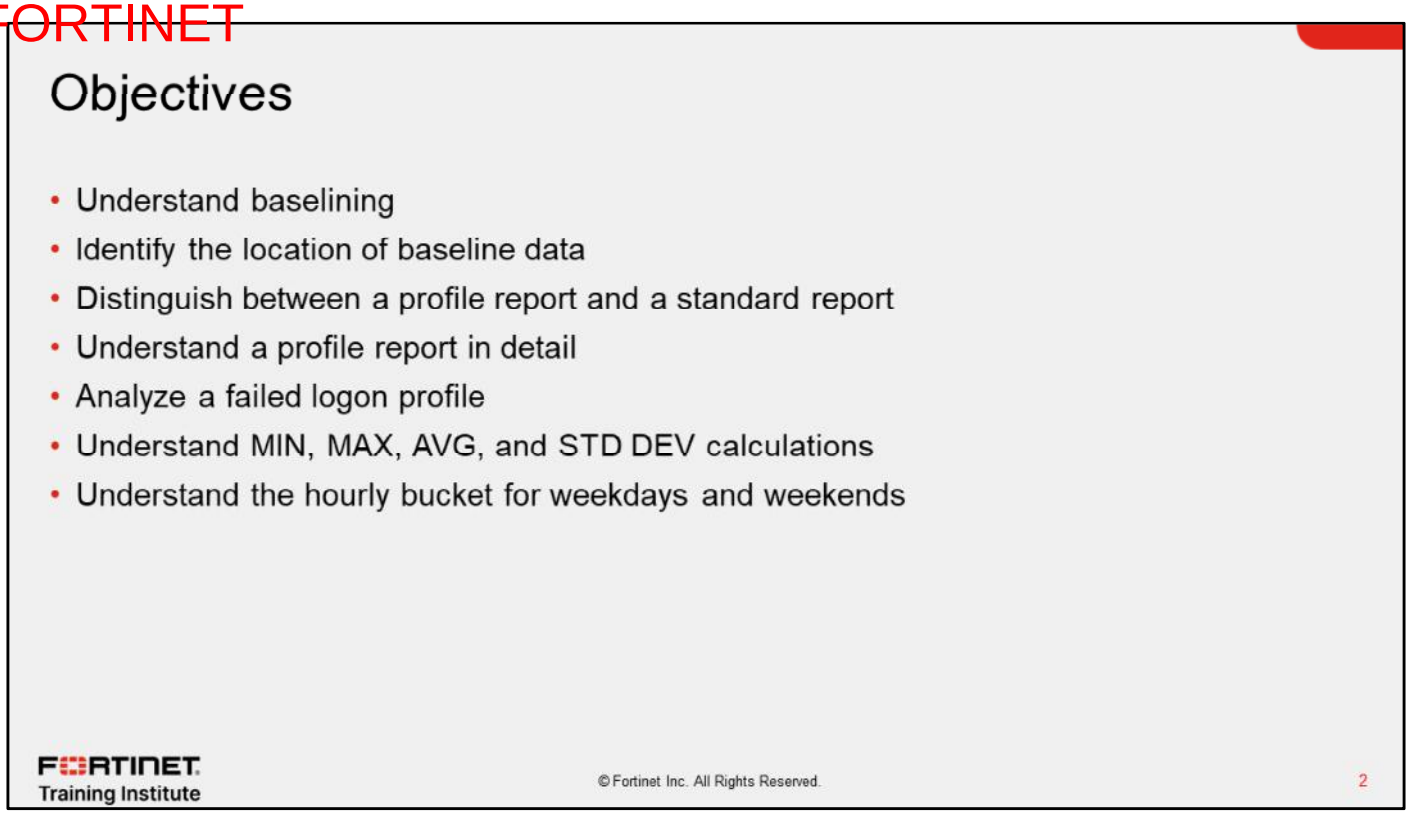
By mastering the objectives covered in this lesson, you learned about multiple subpattern rules, and about the procedure to construct a multiple subpattern rule.

DO NOT REPRINT
© FORTINET



In this lesson, you will learn about baselines on FortiSIEM and analyze the back-end profile report that is responsible for tracking baseline data.

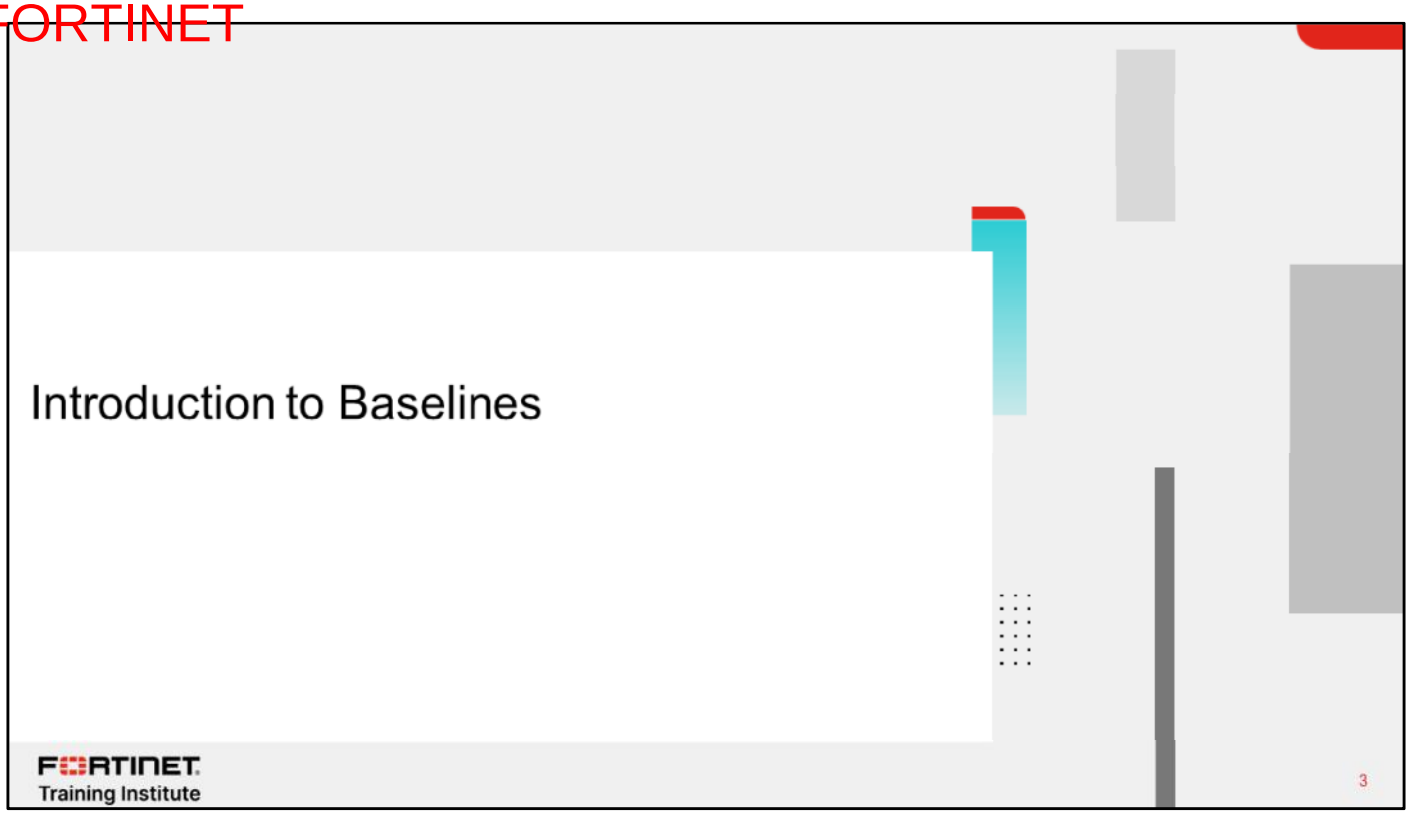
DO NOT REPRINT
© FORTINET



After completing this lesson, you should be able to achieve the objectives shown on this slide.

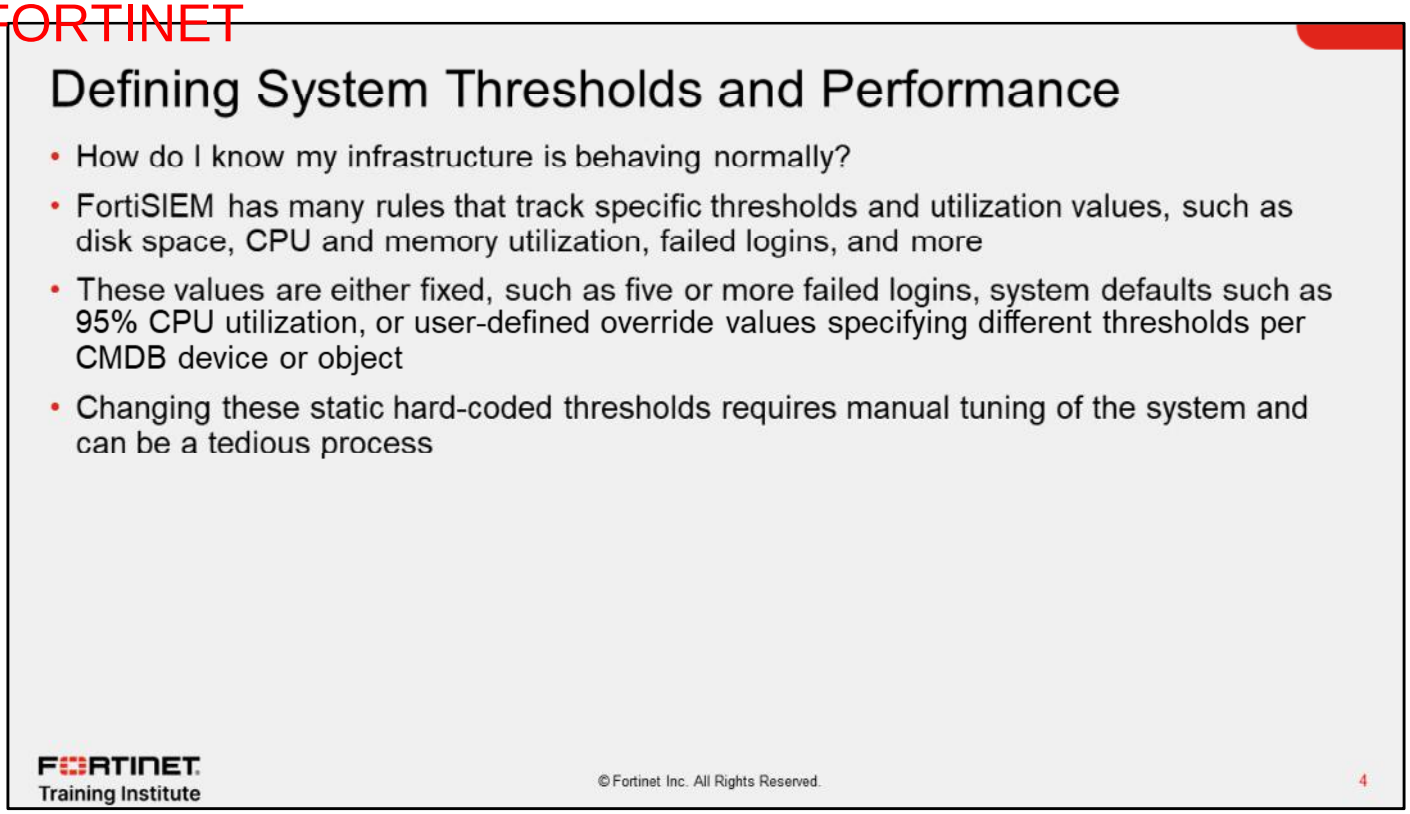
By demonstrating competence in baselines, you will be able to understand the mathematics involved in calculating statistical values for various data from various devices and event types.

DO NOT REPRINT
© FORTINET



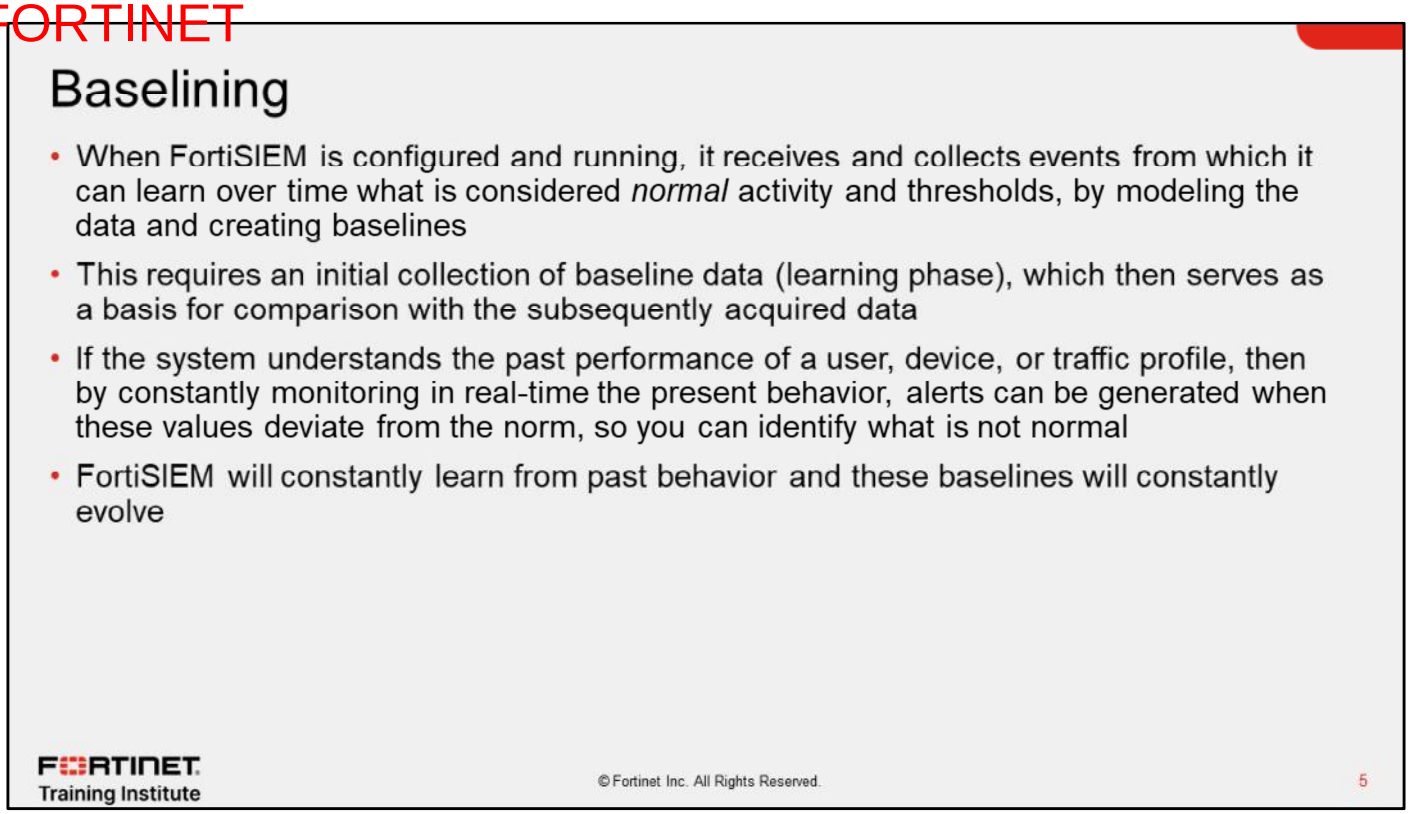
In this section, you will learn about the baseline process and examine a few use case examples.

DO NOT REPRINT
© FORTINET



FortiSIEM has many rules that track specific thresholds and utilization values, such as disk space, CPU and memory utilization, failed logins, and more. These values are either fixed, such as five or more failed logins, system defaults, such as 95% CPU utilization, or user-defined override values specifying different thresholds per CMDB device or object. Changing these static hard-coded thresholds requires manual tuning of the system and can be a tedious process.

DO NOT REPRINT
© FORTINET

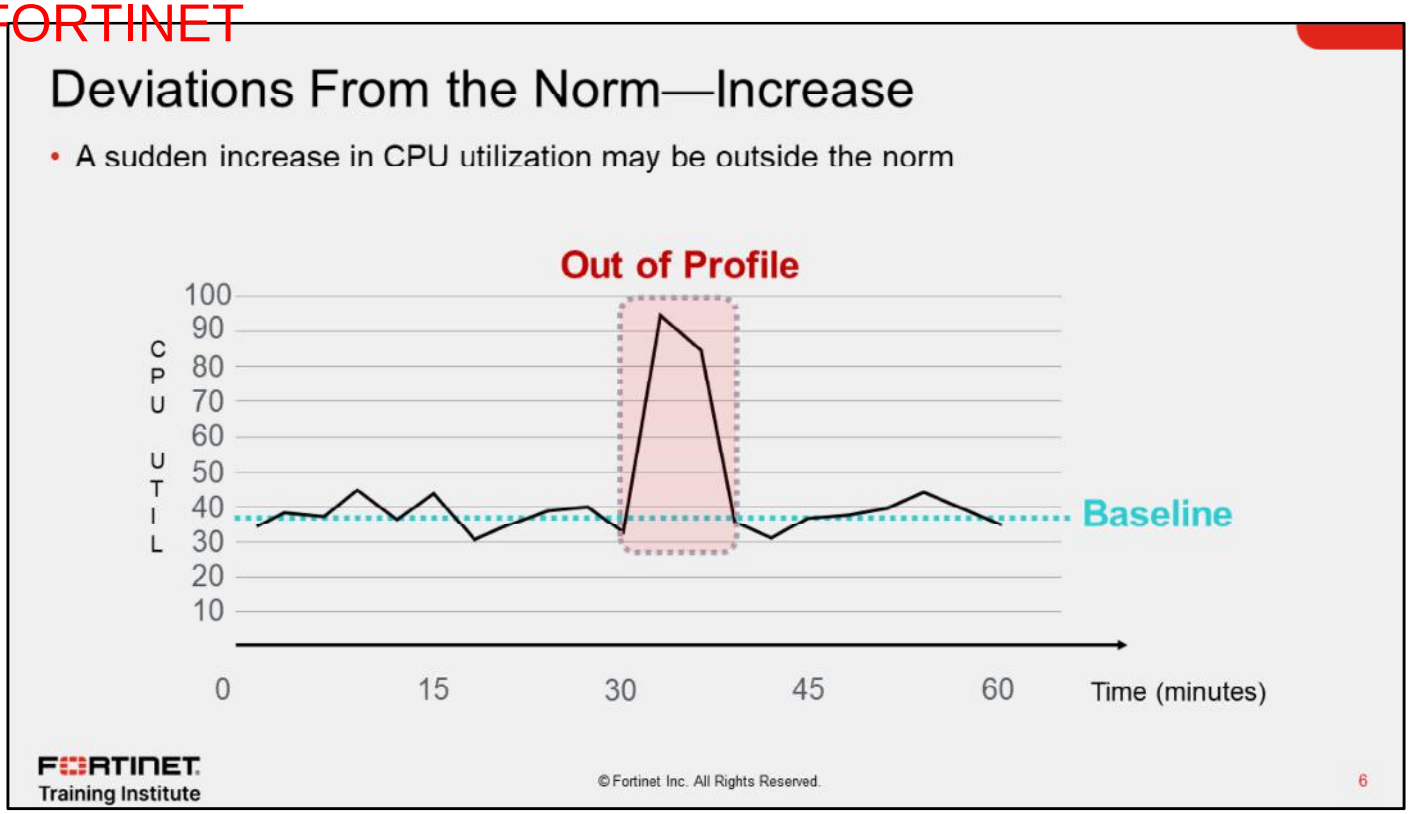


When FortiSIEM is configured and running, it receives and collects events from which it can learn over time what is considered *normal* activity and thresholds, by modeling the data and creating baselines.

This requires an initial collection of baseline data, which is the learning phase. This baseline data serves as a basis for comparison with the subsequently acquired data. If the system understands the past performance of a user, device, or traffic profile, then by constantly monitoring in real-time the present behavior, alerts can be generated when these values deviate from the norm, so you can identify what is not normal.

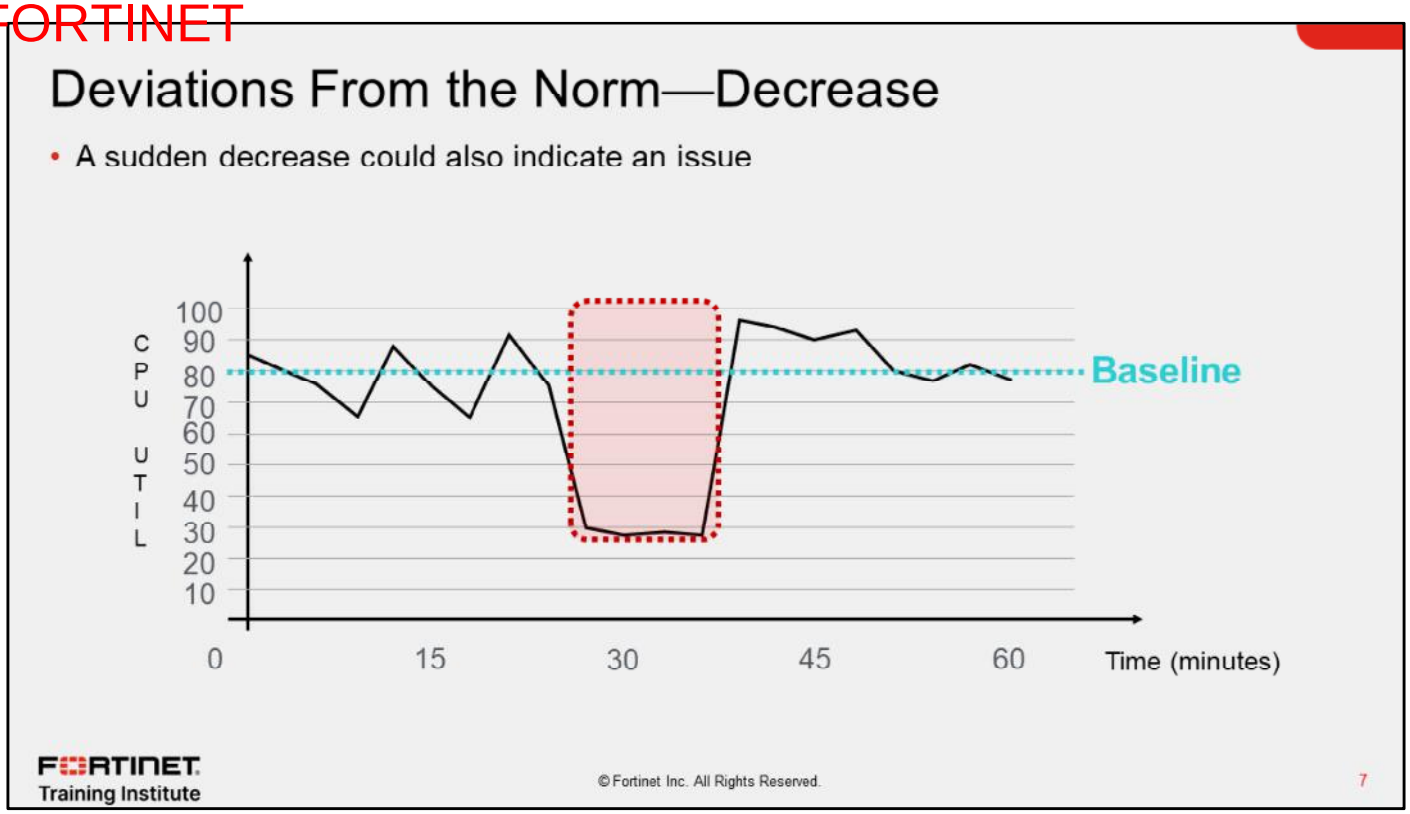
FortiSIEM will constantly learn from past behavior and these baselines will constantly evolve.

DO NOT REPRINT
© FORTINET



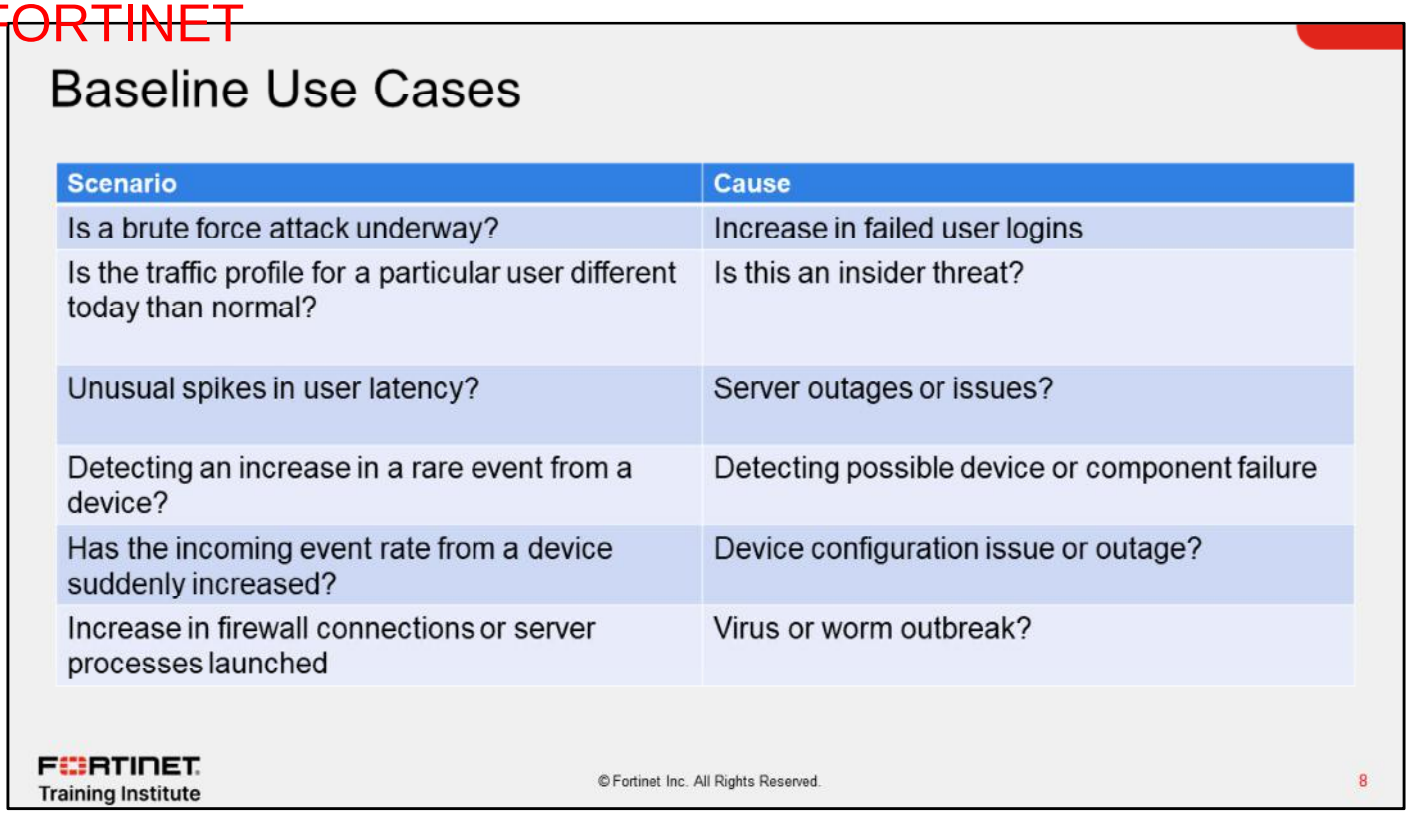
This is an example of a baseline profile for sudden increases in CPU utilization that may be outside the normal behavior. In this case, the CPU utilization spikes above the baseline value, which is considered an abnormal behavior by the baseline profile.

DO NOT REPRINT
© FORTINET



A sudden decrease could also indicate an issue. The baseline profile also tracks for data that is below the baseline average value.

DO NOT REPRINT
© FORTINET



These are few of the scenarios that would be applicable for a baseline profile.

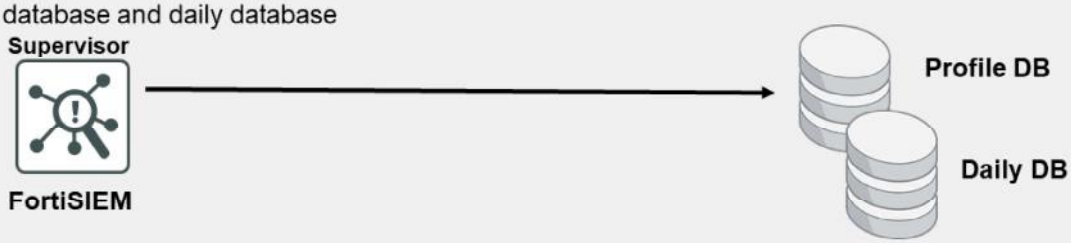
If there is a brute force attack, then most likely there is increased failed user logins. If the traffic profile for a particular user is different on a specific day compared to what it usually is, then it is considered an insider threat. If there is an unusual spike in user latency then most likely it could be a server outage.

If FortiSIEM detects an increase in a rare event from a device, then it is mostly likely a device or a component failure. If the incoming event rate from a device suddenly increases, then it could indicate a configuration issue on the device. If there is an increase in firewall connections or server processes are launched, then it could indicate a virus or worm outbreak.

DO NOT REPRINT
© FORTINET

FortiSIEM Profile Database

- Two small dedicated SQLite databases are used for the baseline data
  - Profile database and daily database



- Located on the disk 1 (OS disk) of the supervisor node
  - `/opt/phoenix/cache/profile.db`
  - `/opt/phoenix/cache/daily.db`
- Stores running trends of learned baseline data calculated for many parameters, such as traffic and device resource usage profiles, running averages, and standard deviation values

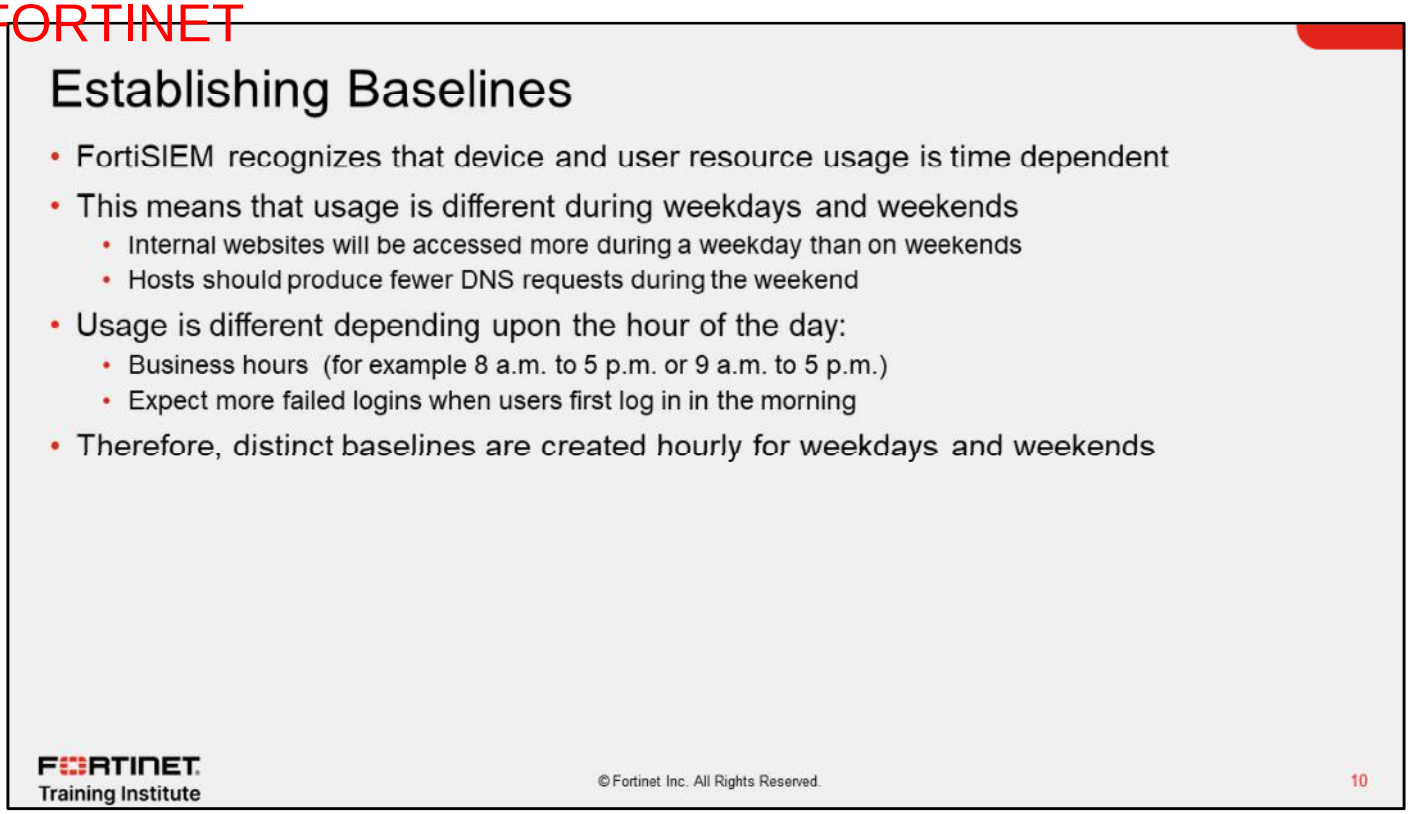
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

9

FortiSIEM uses two small dedicated SQLite databases used for the baseline data—profile database and daily database. These databases are located on disk 1 of the supervisor node. The databases will store running trends of learned baseline data calculated for many parameters, such as traffic and device resource usage profiles, running averages, and standard deviation values.

DO NOT REPRINT
© FORTINET



FortiSIEM recognizes that device and user resource usage is time dependent. This means that usage is different during weekdays and weekends. For example, internal websites will be accessed more frequently during a weekday than during weekends, and hosts should produce fewer DNS requests during the weekend. Also, the usage is different depending upon the hour of the day. During business hours there will naturally be more traffic. You can expect more failed logins when the users first log in in the morning on business days. To take various scenarios into consideration, FortiSIEM builds distinct baselines for hours of the weekdays and weekends.

DO NOT REPRINT
© FORTINET

Buckets

- Hourly buckets are used to store data for the specific baselines
- Each baseline report has a set of keys that represents the baselined metrics and a collection of resultant values
- The supervisor and any worker nodes operate in parallel to analyze the previous hour's data for each baseline, with the supervisor summarizing the results and storing the data in the SQLite database

At 2 p.m. analyze the last hour of data => (1pm to 1:59:59 p.m.)



Hour of day for 1 p.m.

- This process is done for each hour of the day as new baselines values are learned

FORTINET Training Institute

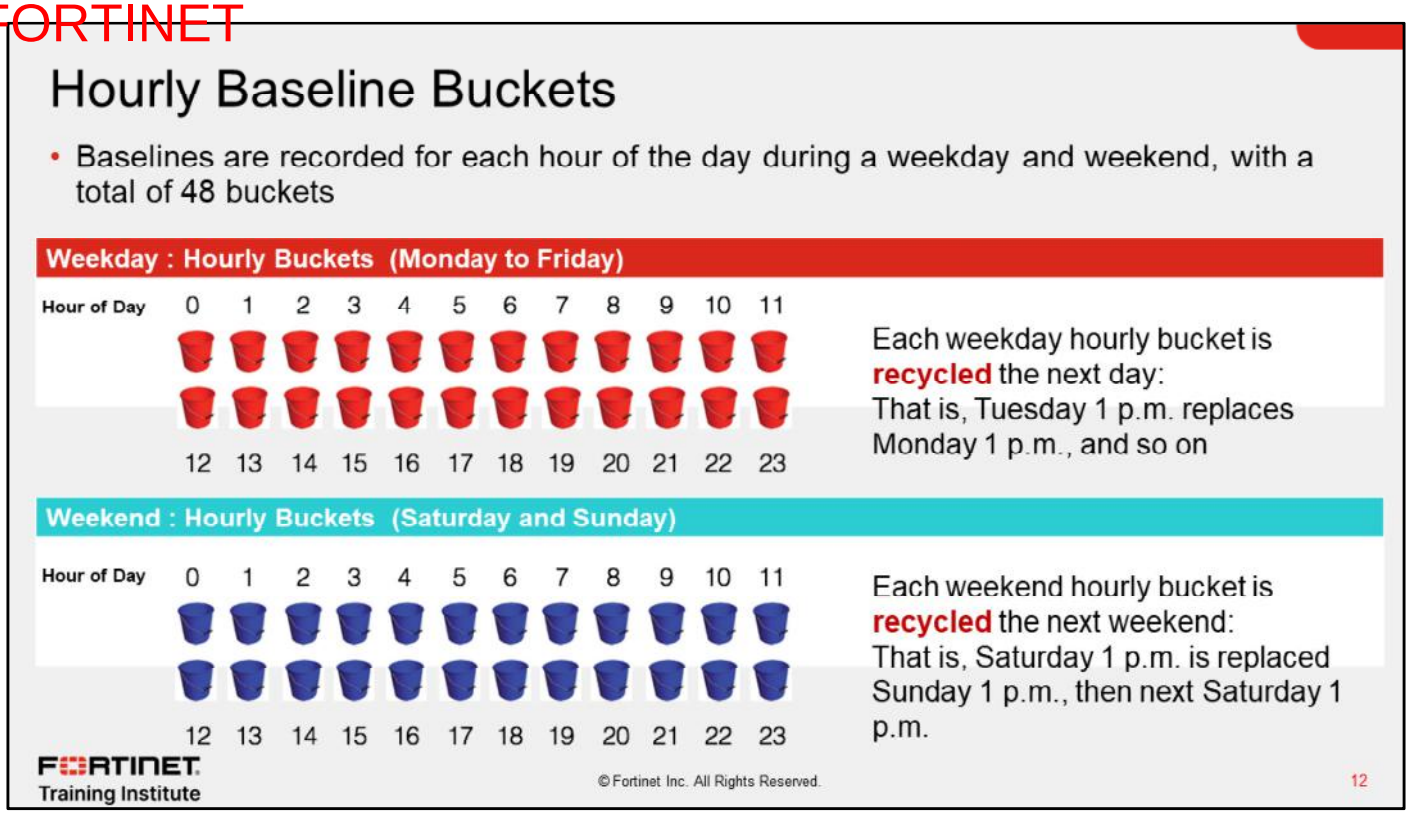
© Fortinet Inc. All Rights Reserved.

11

FortiSIEM uses hourly buckets to store data for specific baseline profiles. Each baseline report has a set of keys that represents the baselined metrics and a collection of resultant values. The supervisor and any worker nodes operate in parallel to analyze the previous hour's data for each baseline, with the supervisor summarizing the results and storing the data in the SQLite database.

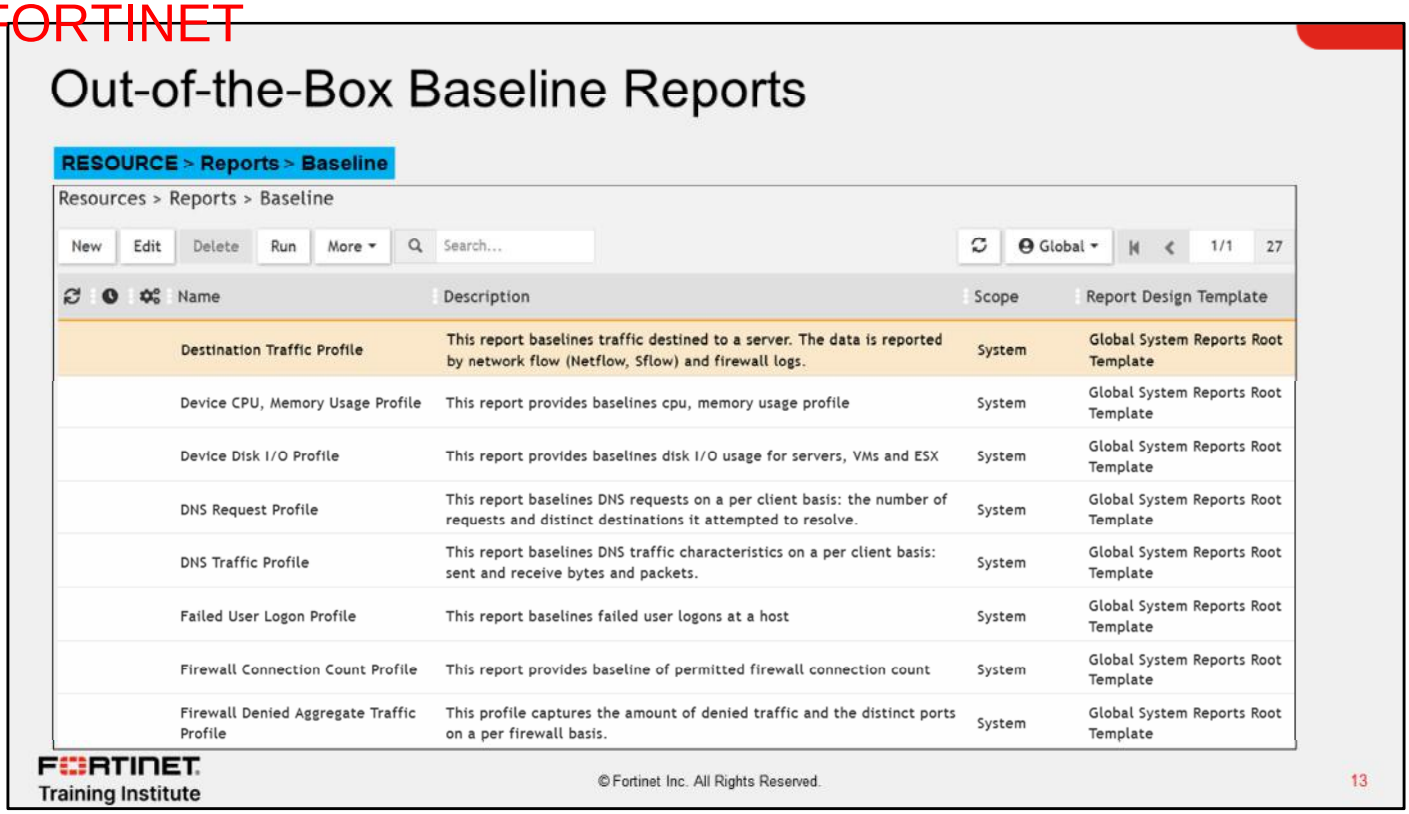
The term used is *hour of day* for the hour that is being analyzed. When FortiSIEM collects data between 1 p.m. to 1:59:59 p.m., that data will be analyzed at 2 p.m. and that will be for the 1 p.m. hour of day. This process is done for each hour of the day as new baselines values are learned.

DO NOT REPRINT
© FORTINET



Baselines are recorded for each hour of the day during a weekday and weekend, with a total of 48 buckets. Each weekday hourly bucket is recycled the next day, which means Tuesday 1 p.m. replaces the 1 p.m. value from Monday. Similarly, each weekend hourly bucket is recycled the next weekend, which means Saturday 1 p.m. is replaced by Sunday 1 p.m. and so on.

DO NOT REPRINT
© FORTINET



When you navigate to baseline reports, you can view the out-of-the-box baseline reports. You can create your own baseline report from here.

Do not confuse this with the baseline profile, which can be created only in the back-end.

DO NOT REPRINT
© FORTINET

Baseline Report Sample—Outbound Firewall Permitted TCP/UDP Port Usage

- Current baseline data can be returned and filtered for each hour of the day, by running a baseline report

RESOURCE > Reports > Baseline

Action | [1] DNS Request Profile | [2] Inbound Firewall Permitted TCP/... | [3] Outbound Firewall Permitted TCP... | Event Type = PH_PROF_ET_105_OUT_PORT_USAGE | Run

Baseline Report - Outbound Firewall Permitted TCP/UDP Port Usage Profile

1/19 265

Profile Date Type	Hour Of Day	IP Protocol	Destination TCP/UDP Port	Avg Distinct Source IP	Avg Distinct Dest IP	Avg Total Flows	Avg Total Bytes64	Std Dev Total Flows	Std Dev Total Bytes64
1 (Work Day)	0	17 (UDP)	53 (DOMAIN)	2	1.16	486.65	129.58 KB	18.18	5.26 KB
1 (Work Day)	0	6 (TCP)	443 (HTTPS)	1.43	2.43	86.81	6.51 MB	1.38	281.1 KB
1 (Work Day)	0	17 (UDP)	443	1	10.95	51.72	406.86 KB	10.2	193.6 KB
1 (Work Day)	0	6 (TCP)	5228	1	1.02	26.98	8.57 MB	0.16	8.41 MB
1 (Work Day)	0	17 (UDP)	123 (NTP)	1	4	21.7	113.97 KB	13.3	540.71 KB
1 (Work Day)	0	6 (TCP)	80 (HTTP)	1	1	1	457.62 KB	0	0B
1 (Work Day)	1	17 (UDP)	53 (DOMAIN)	2	1.03	478.19	127.06 KB	30.2	8.56 KB
1 (Work Day)	1	6 (TCP)	443 (HTTPS)	1.36	2.5	86.8	6.31 MB	1.13	163.12 KB
1 (Work Day)	1	17 (UDP)	443	1	10.01	49.66	411.2 KB	11.49	190.68 KB
1 (Work Day)	1	6 (TCP)	5228	1	1	26.68	8.1 MB	0.47	8.24 MB
1 (Work Day)	1	17 (UDP)	123 (NTP)	1	4	21.18	37.74 KB	11.56	358.19 KB
1 (Work Day)	1	6 (TCP)	80 (HTTP)	1	2	3	151.06 MB	0	0B
1 (Work Day)	2	17 (UDP)	53 (DOMAIN)	2	1.16	492.53	131.53 KB	36.83	10.81 KB
1 (Work Day)	2	6 (TCP)	443 (HTTPS)	1.36	2.36	87.15	6.43 MB	1.49	197.16 KB

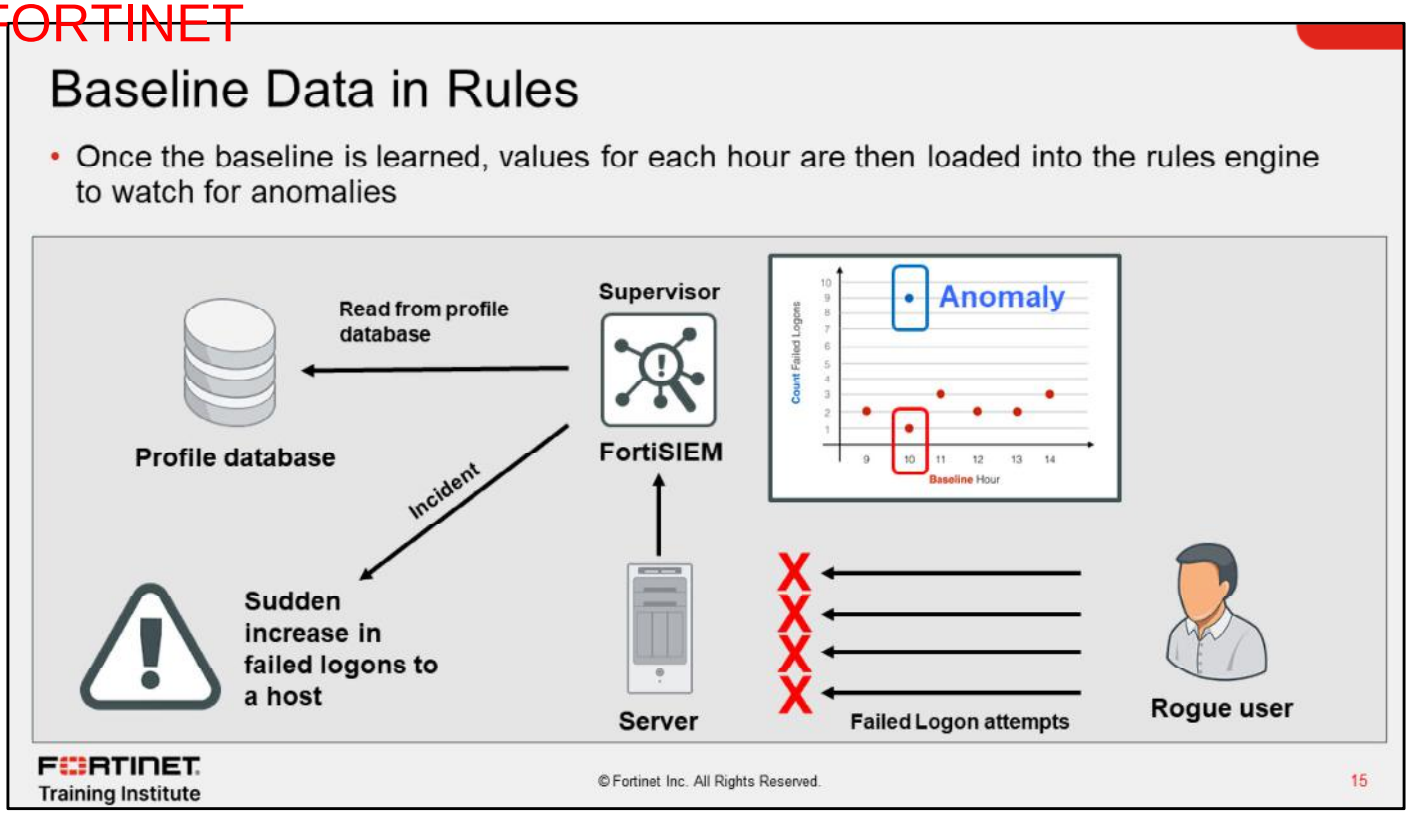
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

14

Current baseline data can be returned and filtered for each hour of the day, by running a baseline report. This data is queried from the profile database for each hour.

DO NOT REPRINT
© FORTINET

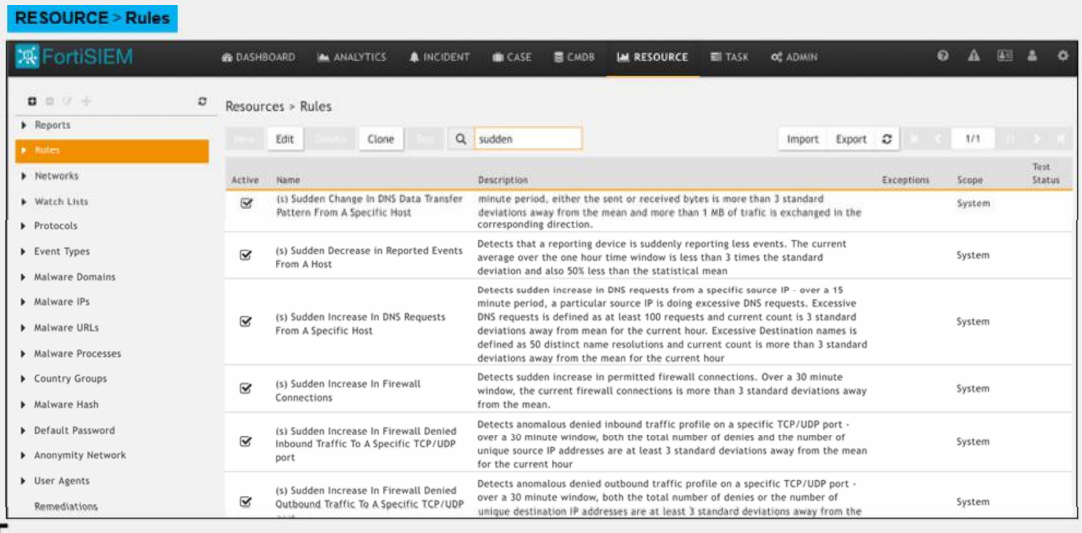


The baseline data can be used by the rule engine to evaluate aggregate conditions in a rule. The rule engine loads the baseline values for each hour from the profile database into the memory and then computes the aggregate condition against the current data. If the result points to an anomaly from the baseline values, then the rule can trigger an incident.

DO NOT REPRINT
© FORTINET

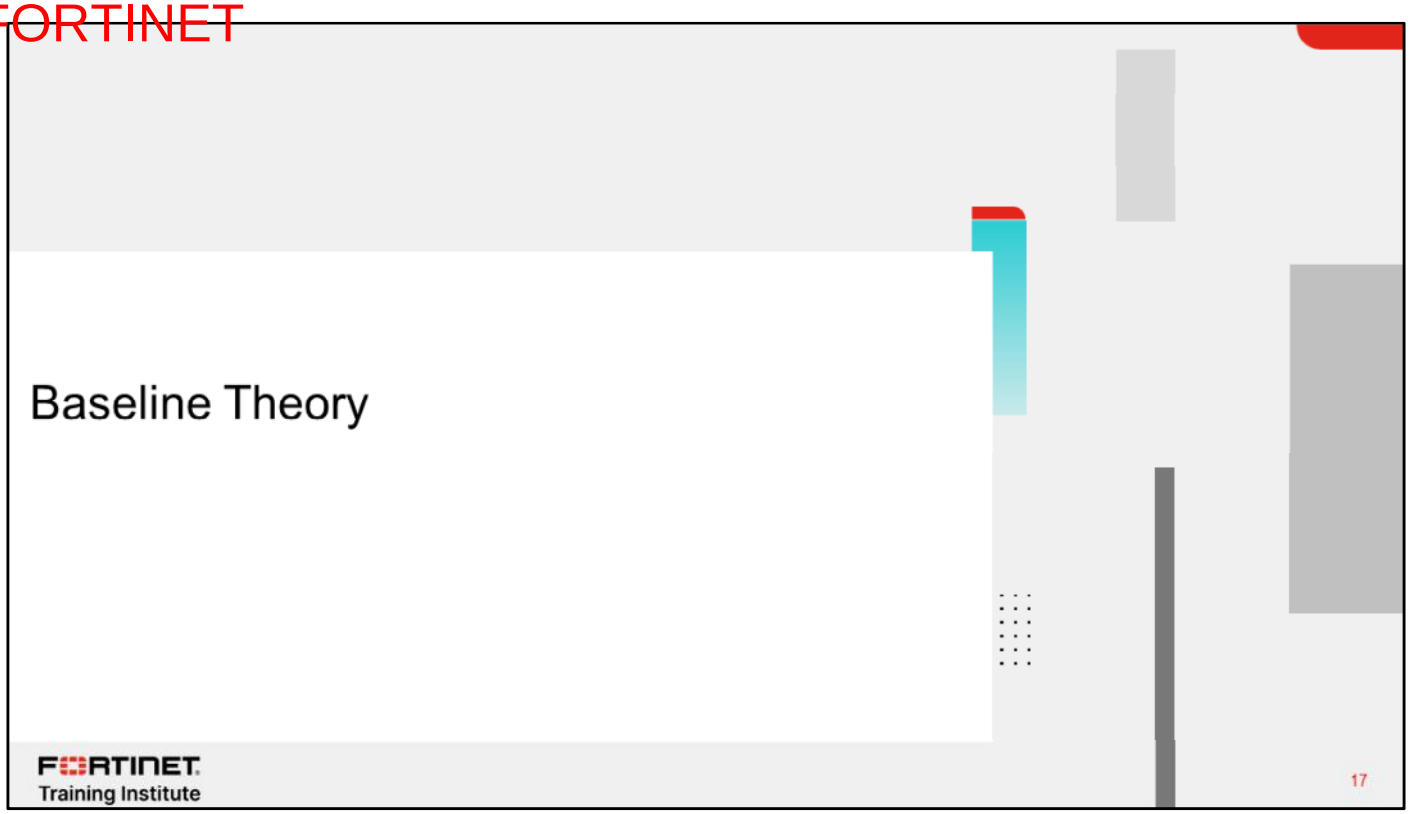
Baseline Anomaly Rules

- Corresponding baseline-related rules are also provided out-of-the-box
- The rule names start with the term *Sudden*



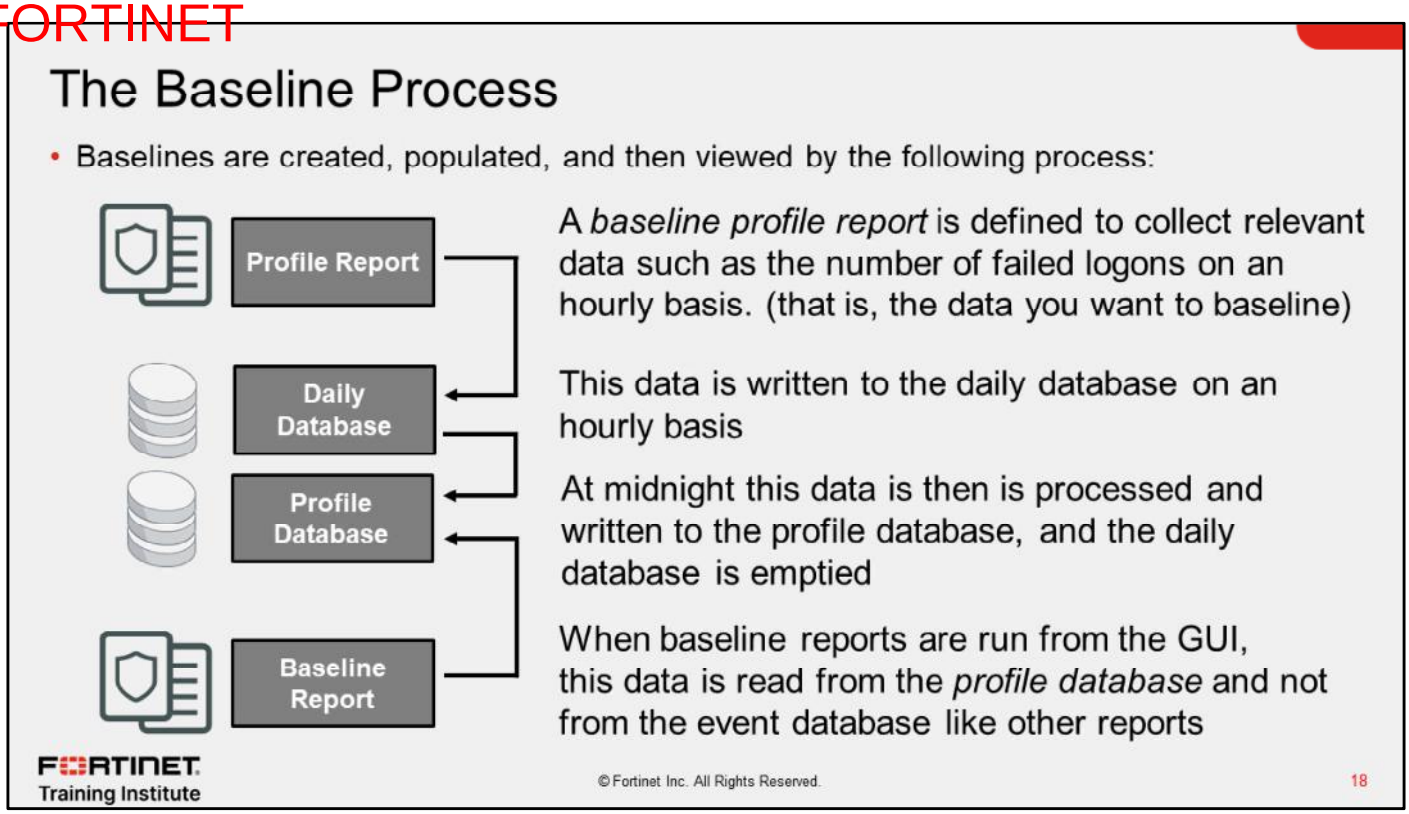
There are several out-of-the-box rules that refer to baseline data to compute aggregate conditions and generate incidents. The rule names start with the term *Sudden*.

DO NOT REPRINT
© FORTINET



In this section, you will learn the theory behind baselines on FortiSIEM.

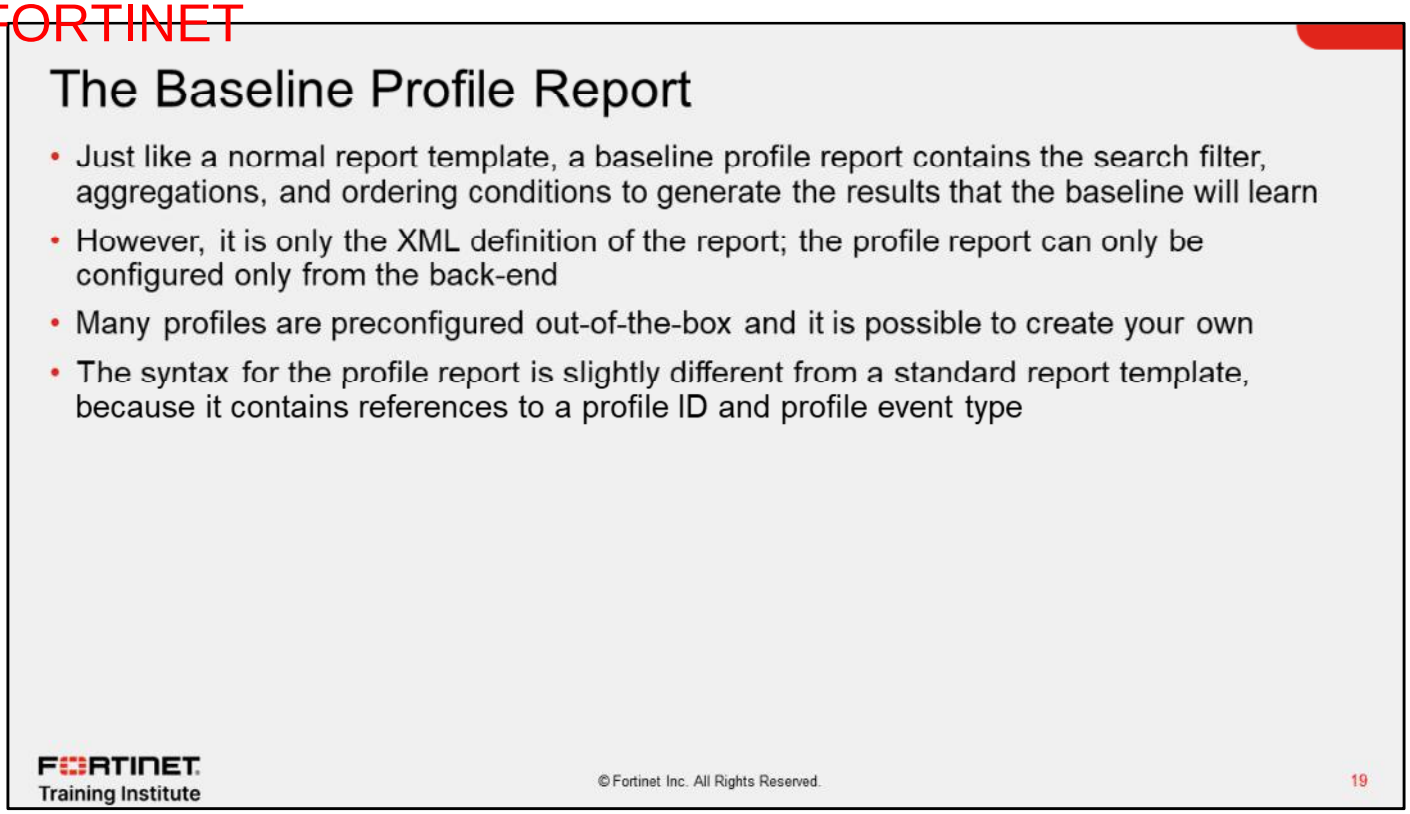
DO NOT REPRINT
© FORTINET



A profile report is defined to collect relevant data, such as the number of failed logons on an hourly basis. This is the data you want to baseline. This data is written to the daily database on an hourly basis.

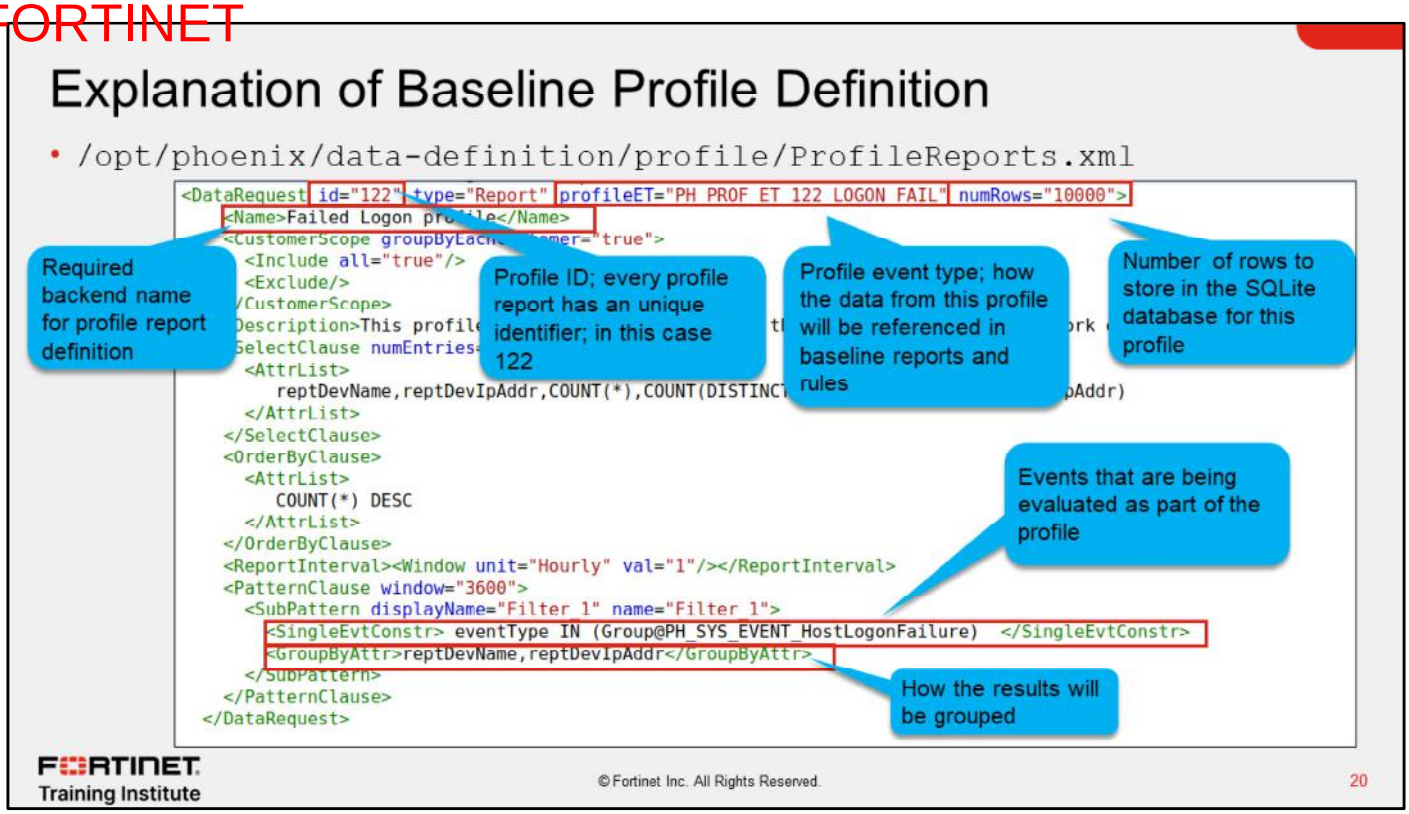
At midnight, this data is processed and written to the profile database, and the daily database is emptied. When baseline reports are run from the GUI, FortiSIEM uses the data from the profile database and not from the event database like other reports.

DO NOT REPRINT
© FORTINET



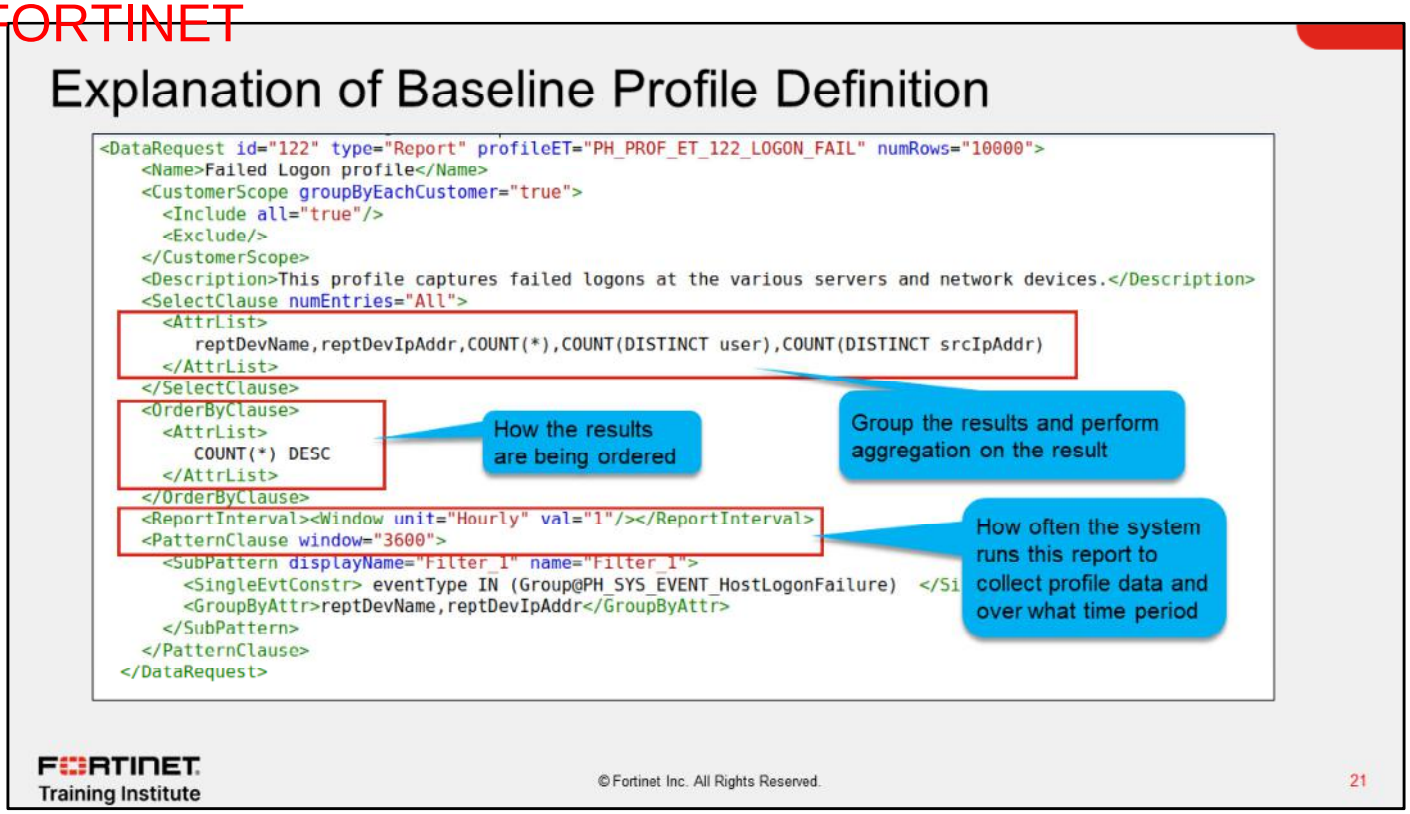
Just like a normal report template, a baseline profile report contains the search filter, aggregations, and ordering conditions to generate the results that the baseline will learn. The profile report is the XML definition of the report, because the profile report can be configured only from the back-end. Many profiles are preconfigured out-of-the-box and it is possible to create your own. The syntax for the profile report is slightly different from a standard report template, because it contains the reference to a profile ID and profile event type.

DO NOT REPRINT
© FORTINET



The profile reports are defined in the `ProfileReports.xml` file in the FortiSIEM back-end. This slide shows an example of the failed logon profile, which has an unique identifier—122. The profile event type defines the way the data from this profile will be referenced in the baseline reports and rules. The profile also defines the number of rows to store in the SQLite database for this profile. You must assign a name to the profile, which does not appear in the GUI. It is defined for the purpose of identifying the profile. You must also define the events that will be evaluated by this profile, and define how the results will be grouped.

DO NOT REPRINT
© FORTINET

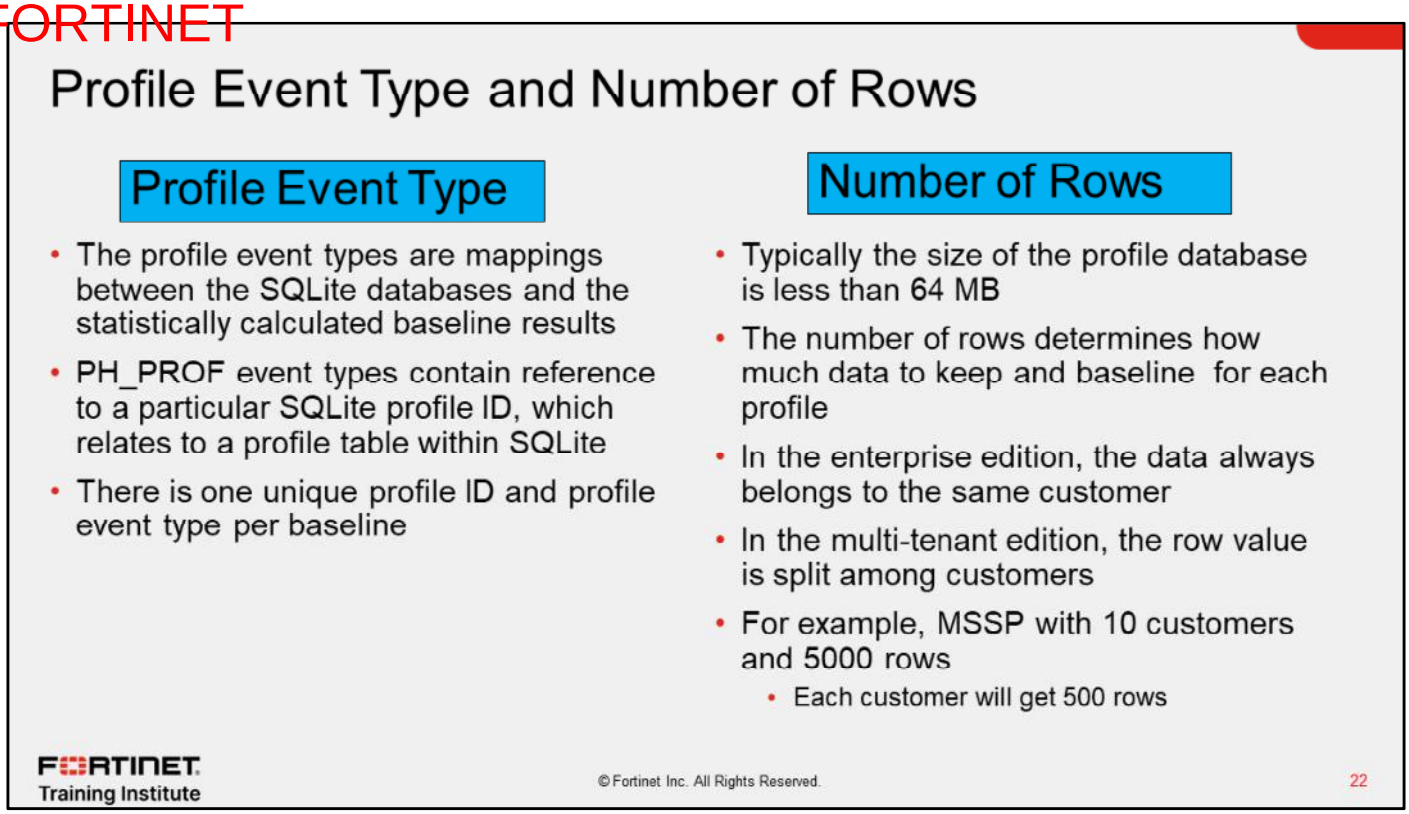


For the baseline profile to perform aggregation, you must first group the results, which are defined along with the aggregate function.

In the example shown on this slide, the results will be grouped by reporting device name and reporting device IP address. After the results are grouped into those two columns, three more columns will be added for the count functions, which are the aggregate function. After that, the results will be arranged in descending order according to the count column value.

The window specifies how often the baseline profile will collect data points and over what time period. In this example, the data will be collected every hour for 24 hours. This means that the baseline profile will consider all failed logon events for an hour and perform the necessary calculations and then store that result in the daily database bucket.

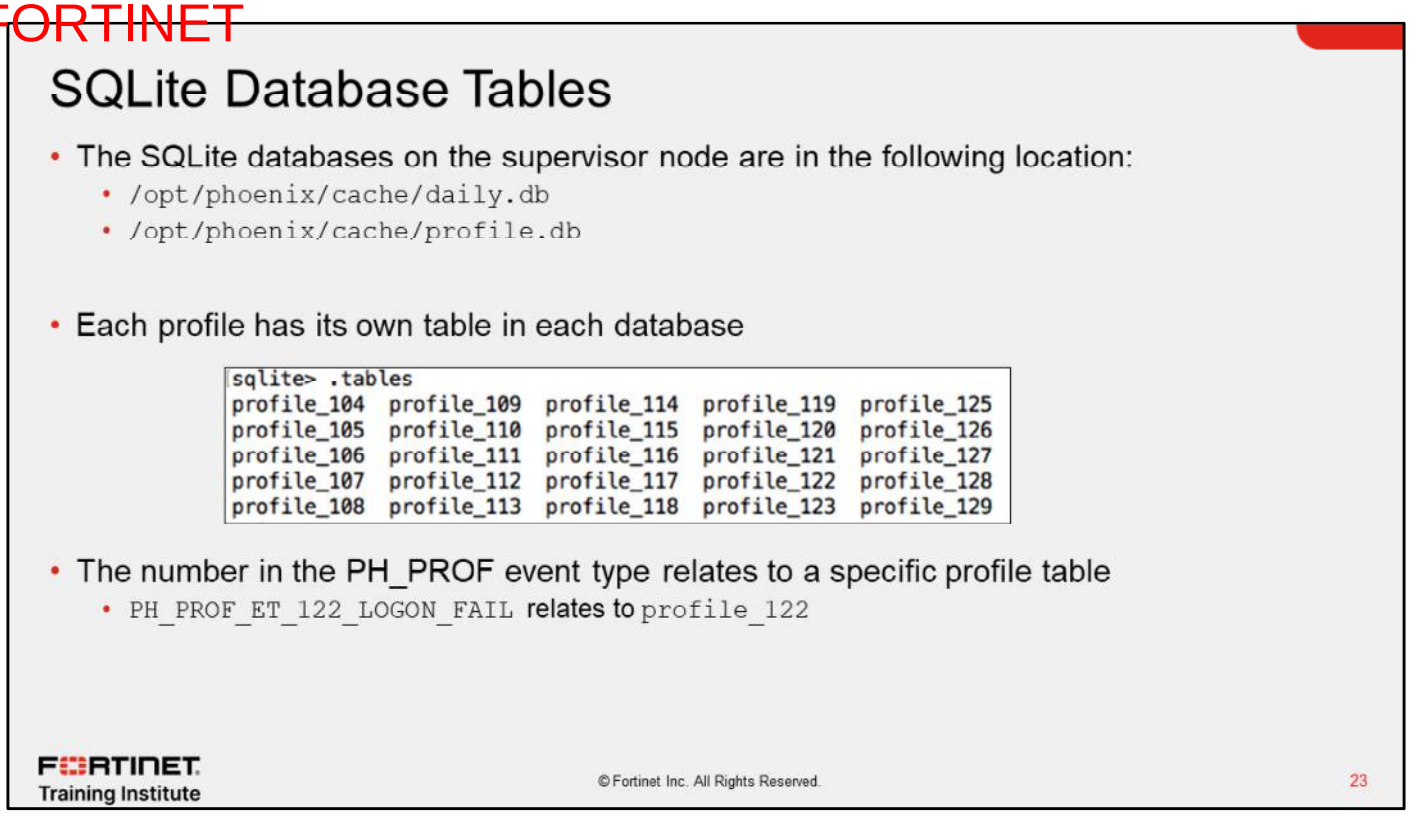
DO NOT REPRINT
© FORTINET



The profile event types are mappings between the SQLite databases and the statistically calculated baseline results. PH_PROF event types contain a reference to a particular SQLite profile ID, which relates to a profile table within SQLite. There is one unique profile ID and profile event type per baseline.

The number of rows determines how much data to keep and baseline for each profile. Typically, the size of the profile database is less than 64 MB, but could be as high as 500 MB in the multi-tenant edition. In the enterprise edition, the data always belongs to the same customer. In the multi-tenant edition, this value is split among customers. For example, if an MSSP has 10 customers and 5000 rows, then each customer will get 500 rows.

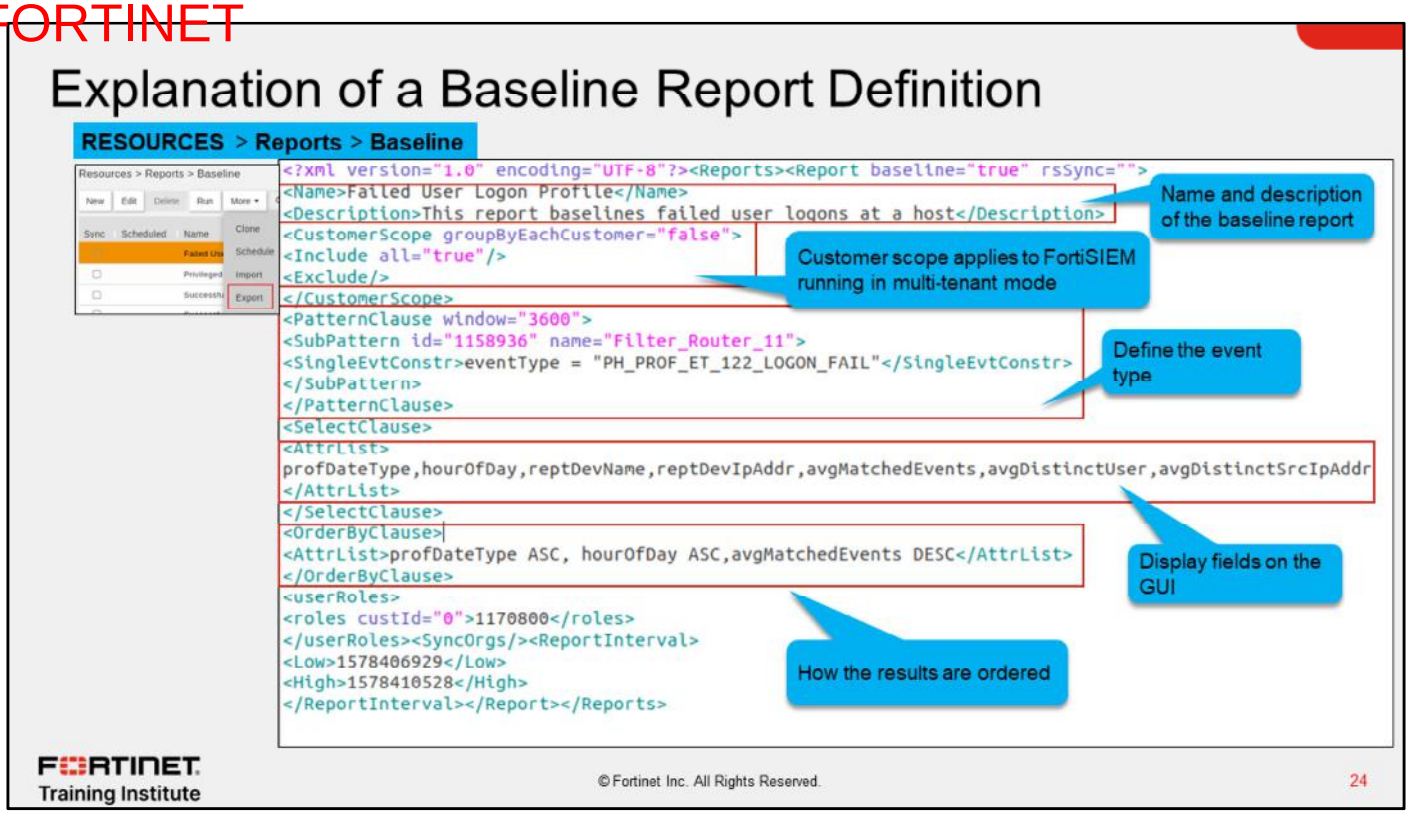
DO NOT REPRINT
© FORTINET



The SQLite databases on the supervisor node are located in the cache directory. The specific directory paths are shown on this slide.

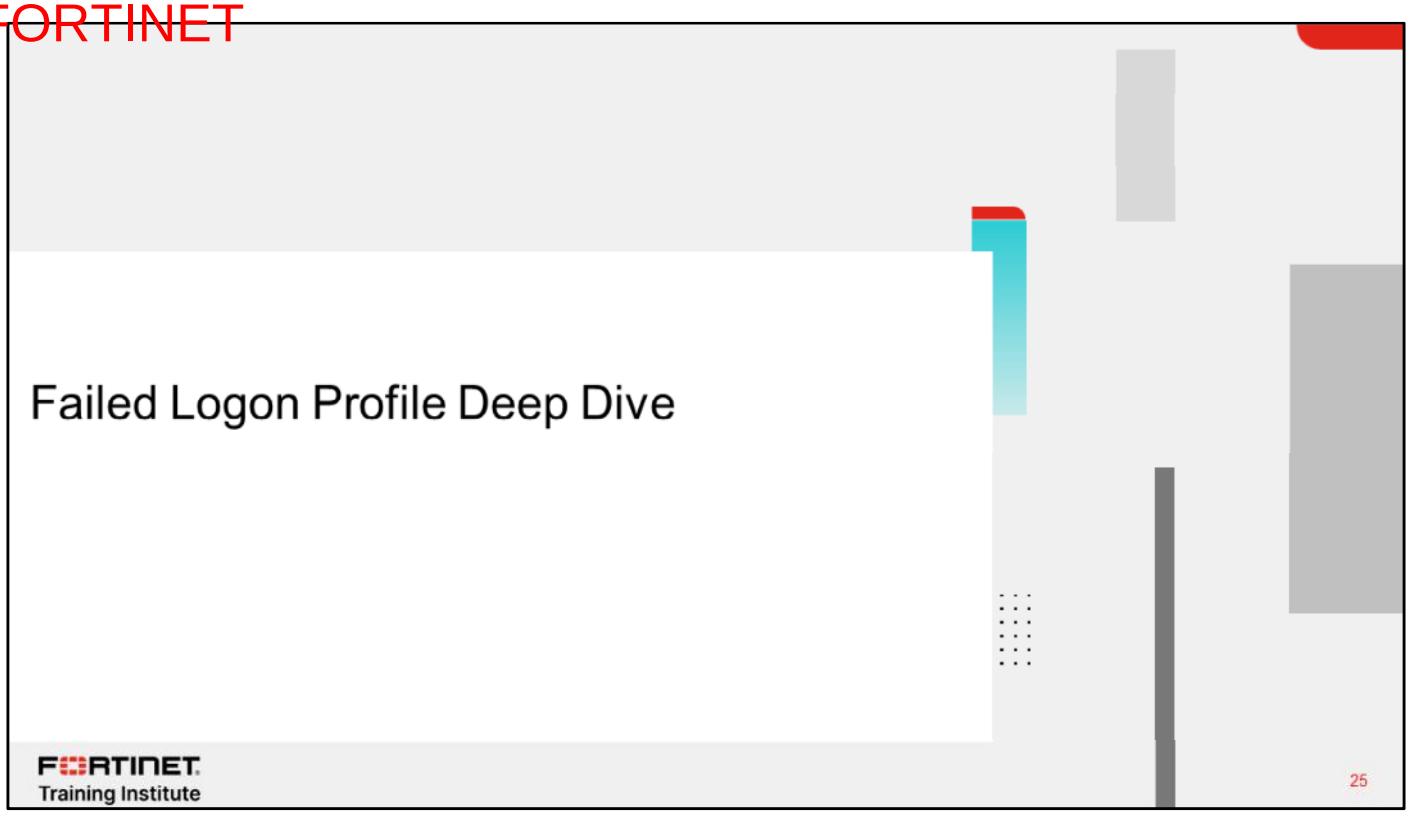
Each profile has its own table in each database. The number in the PH_PROF event type relates to a specific profile table. In the case of the failed logon profile it is `profile_122`.

DO NOT REPRINT
© FORTINET



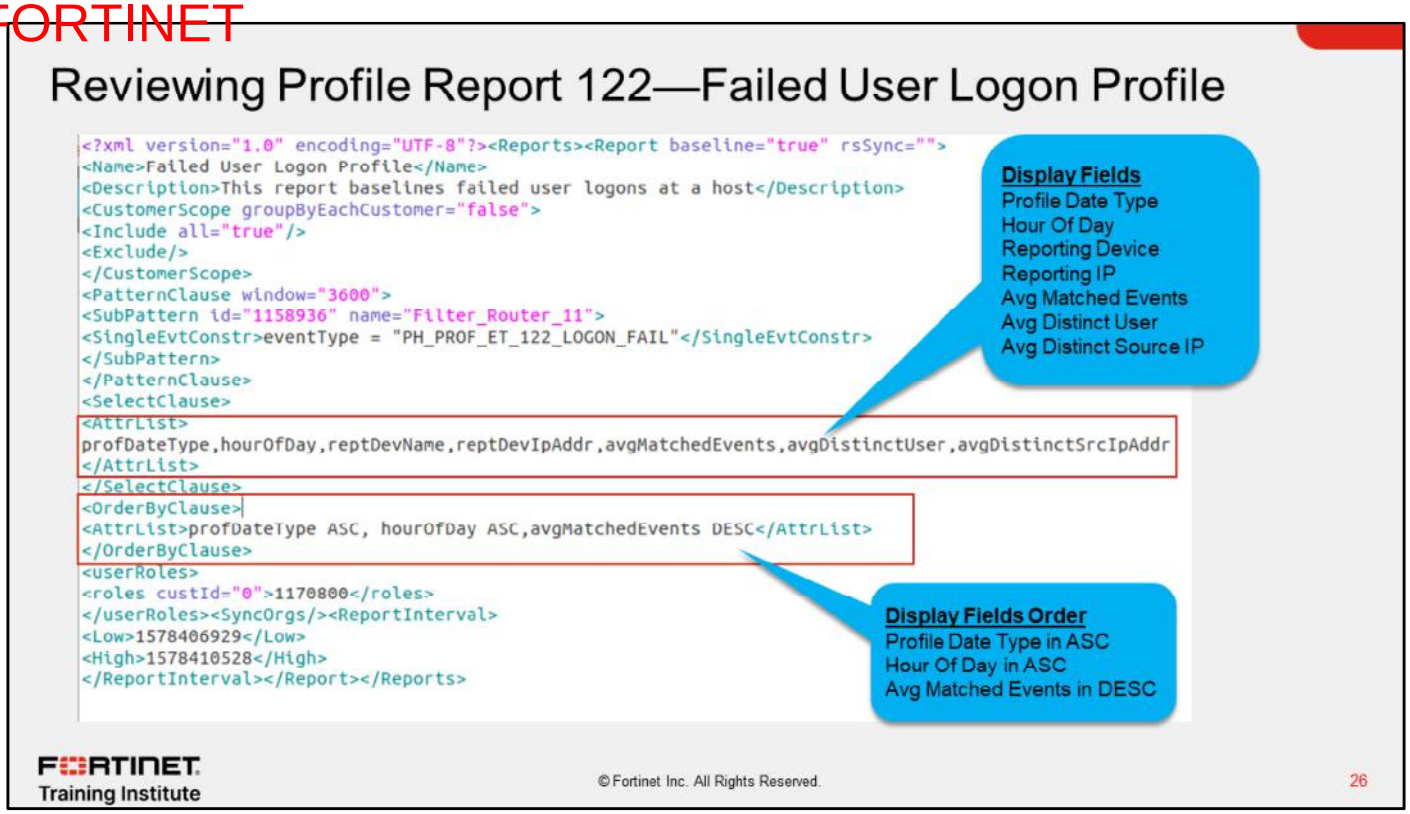
This slide shows the XML definition of a baseline report. This should not be confused with a baseline profile. You can select and export a baseline report in XML format. All the attributes that you see in the XML file can be defined through the GUI.

DO NOT REPRINT
© FORTINET



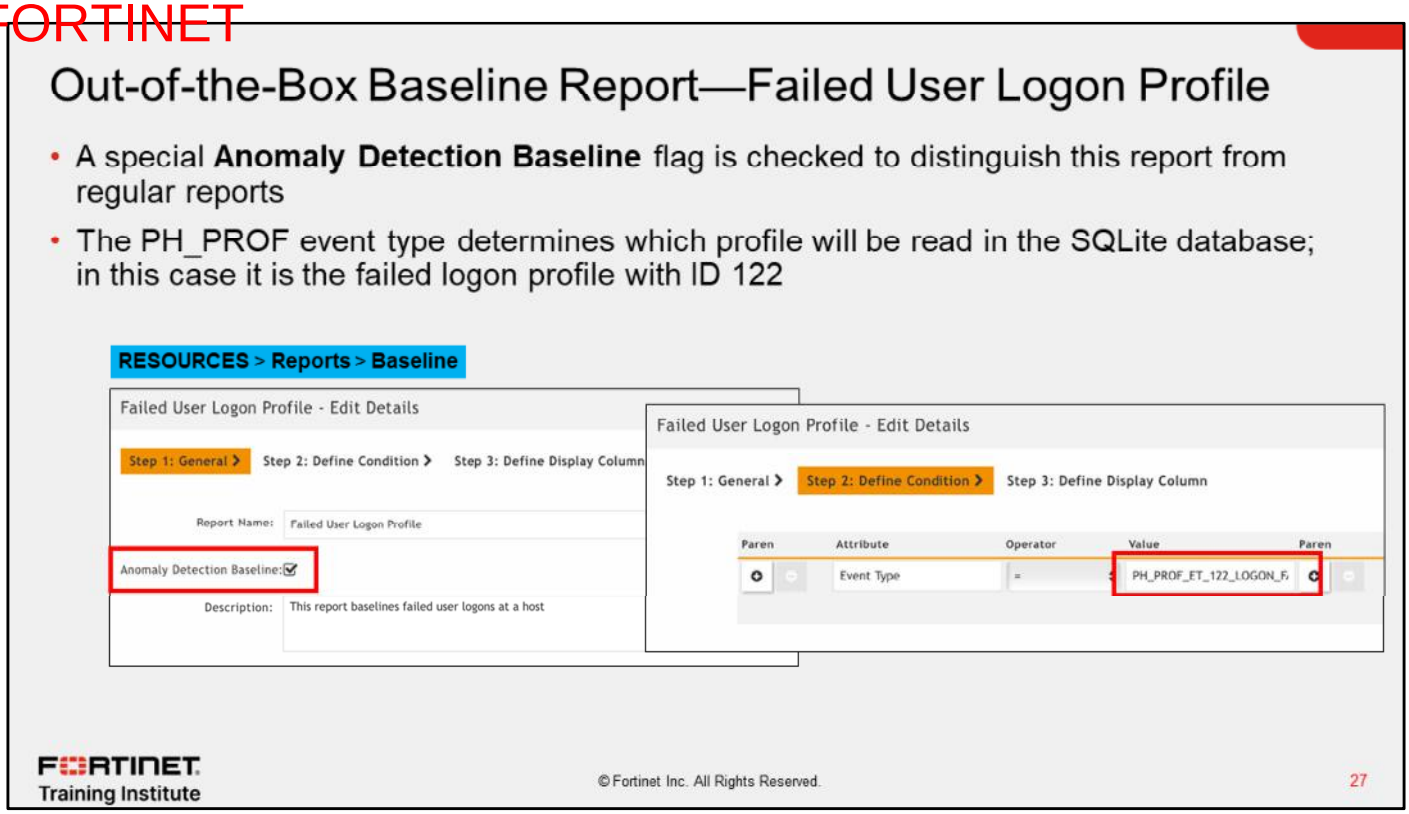
In this section, you will learn about the failed logon profile in depth.

DO NOT REPRINT
© FORTINET



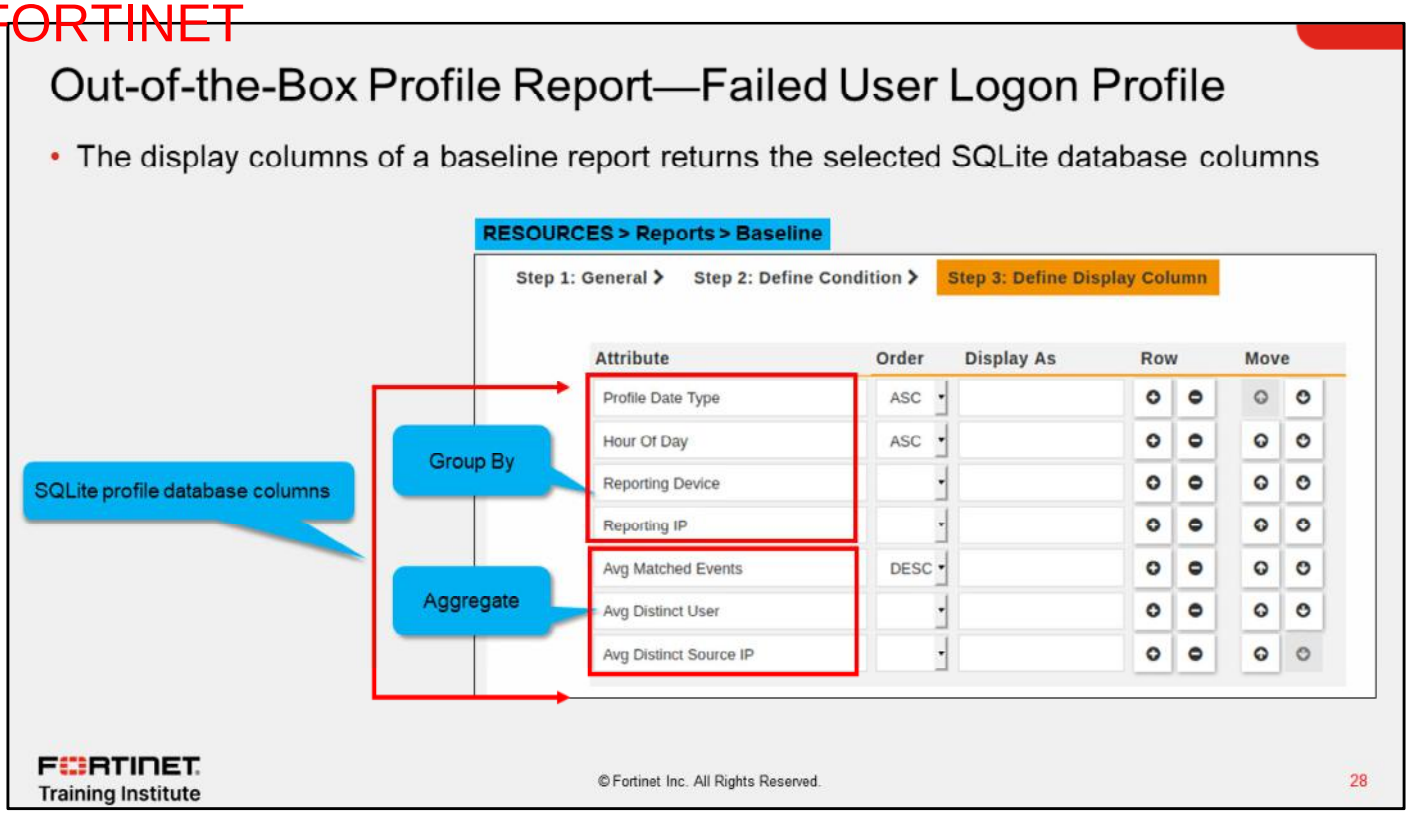
This is the XML definition of the failed user logon profile report that you have seen on previous slides. You will analyze the report output and correlate that data with SQLite database attributes. The report will display seven columns and order them in ascending order based on profile date type and hour of day. It will also order them in descending order based on average matched events.

DO NOT REPRINT
© FORTINET



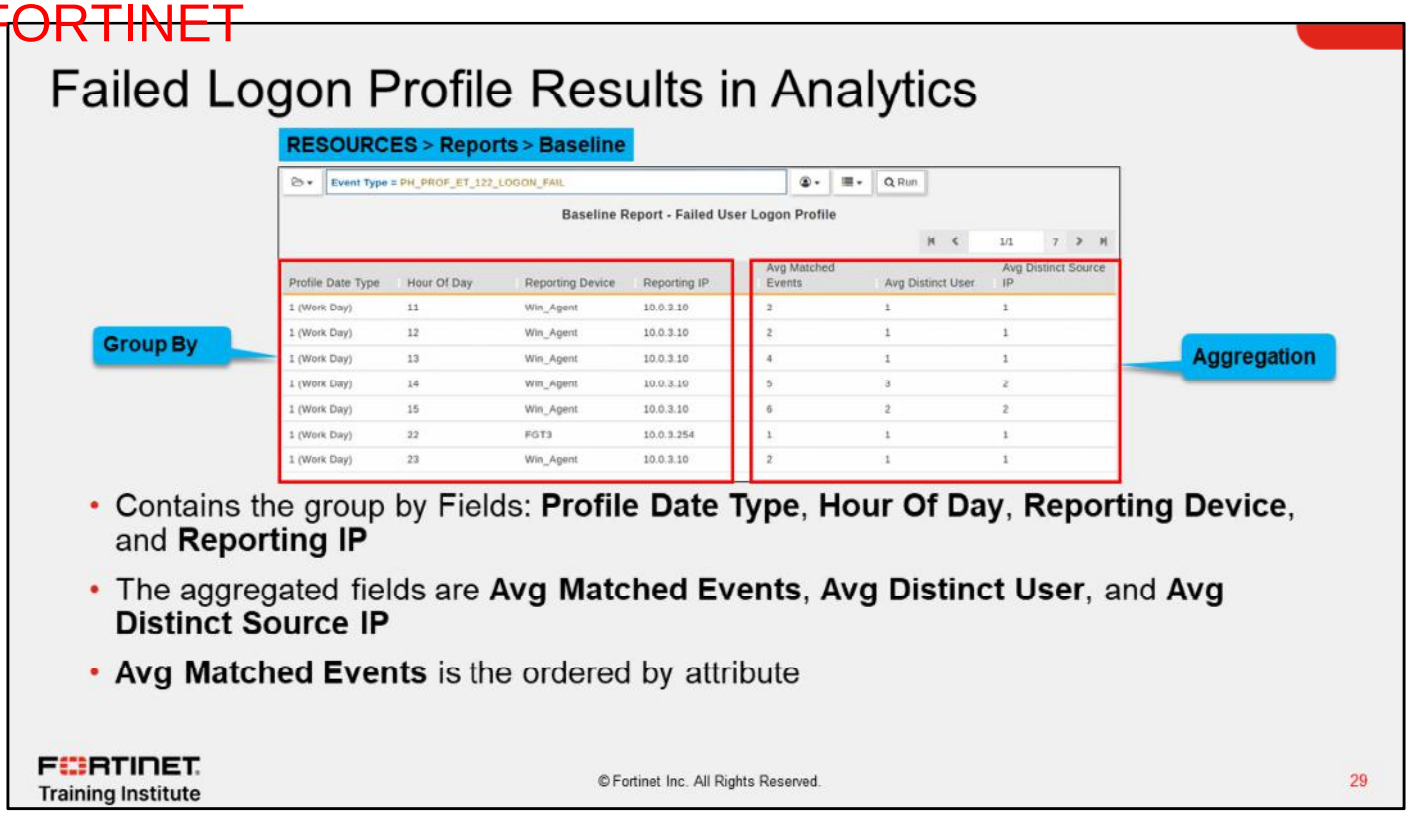
When you edit the report, you will notice that a special flag called **Anomaly Detection Baseline** is set to distinguish this report from regular reports. The PH_PROF event type determines which profile will be read in the SQLite database—in this case it is the failed logon profile with ID 122.

DO NOT REPRINT
© FORTINET



The display column of the baseline report returns the selected SQLite database columns from profile 122. Profile date type, hour of day, reporting device, and reporting IP are the group by attributes. Average matched events, average distinct user, and average distinct source IP are the aggregate functions.

DO NOT REPRINT
© FORTINET



When you run the baseline report, FortiSIEM will query the profile database and display the results, which contains the group by fields column and the aggregate fields column.

DO NOT REPRINT
© FORTINET

SQLite Aggregate Fields

Aggregate Condition	SQLite Profile Table Field
MIN(COUNT(Matched Events))	minMatchedEvents
MAX(COUNT(Matched Events))	maxMatchedEvents
AVG(COUNT(Matched Events))	avgMatchedEvents
StdDev(COUNT(Matched Events))	sdevMatchedEvents

Aggregate Condition	SQLite Profile Table Field
MIN(AVG(CPU Util))	minCpuUtil
MAX(AVG(CPU Util))	maxCpuUtil
AVG(AVG(CPU Util))	avgCpuUtil
StdDev(AVG(CPU Util))	sdevCpuUtil

Aggregate Condition	SQLite Profile Table Field
MIN(SUM(SentBytes)	minSentBytes
MAX(SUM(SentBytes)	maxSentBytes
AVG(SUM(SentBytes)	avgSentBytes
StdDev(SUM(SentBytes)	sdevSentBytes

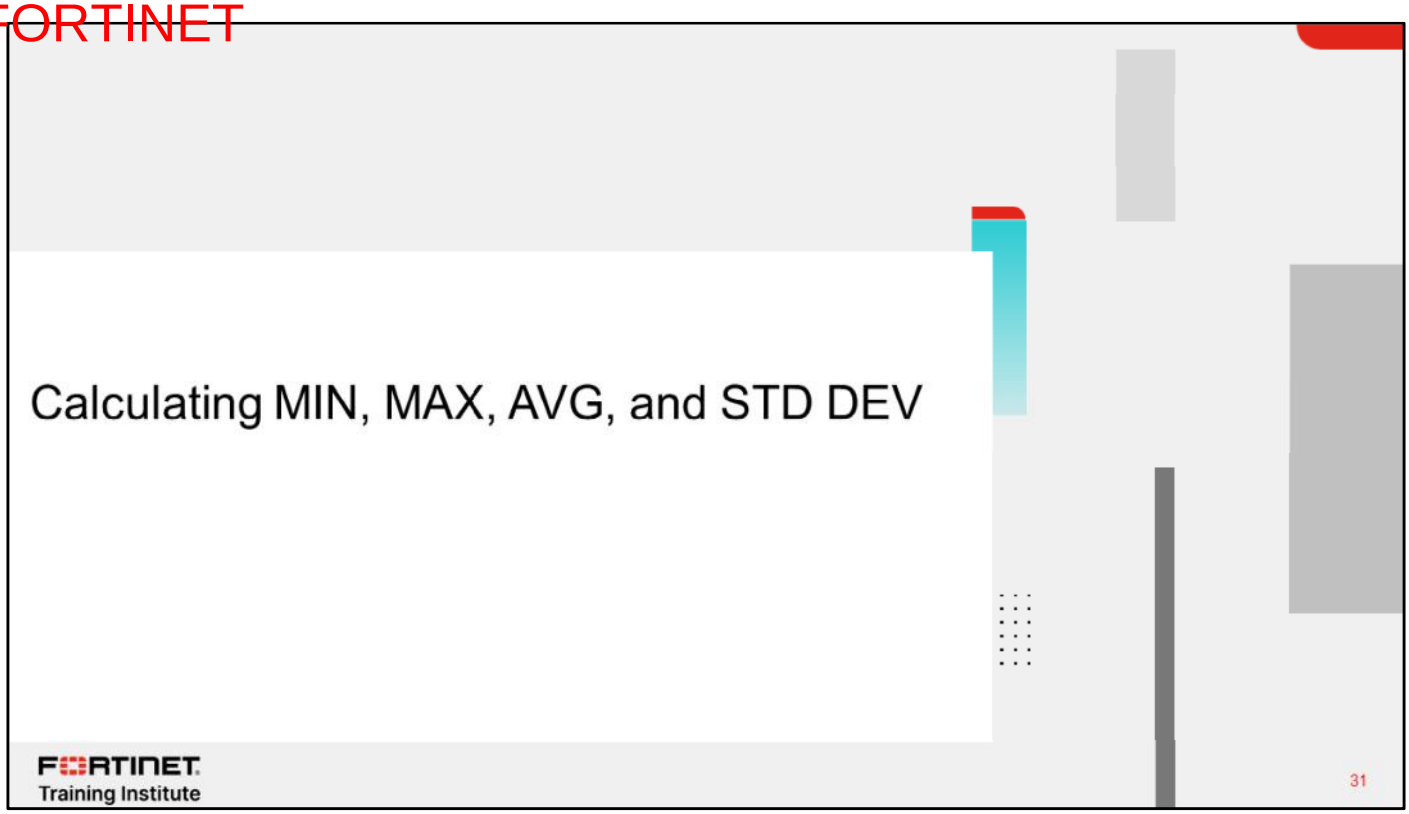
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

30

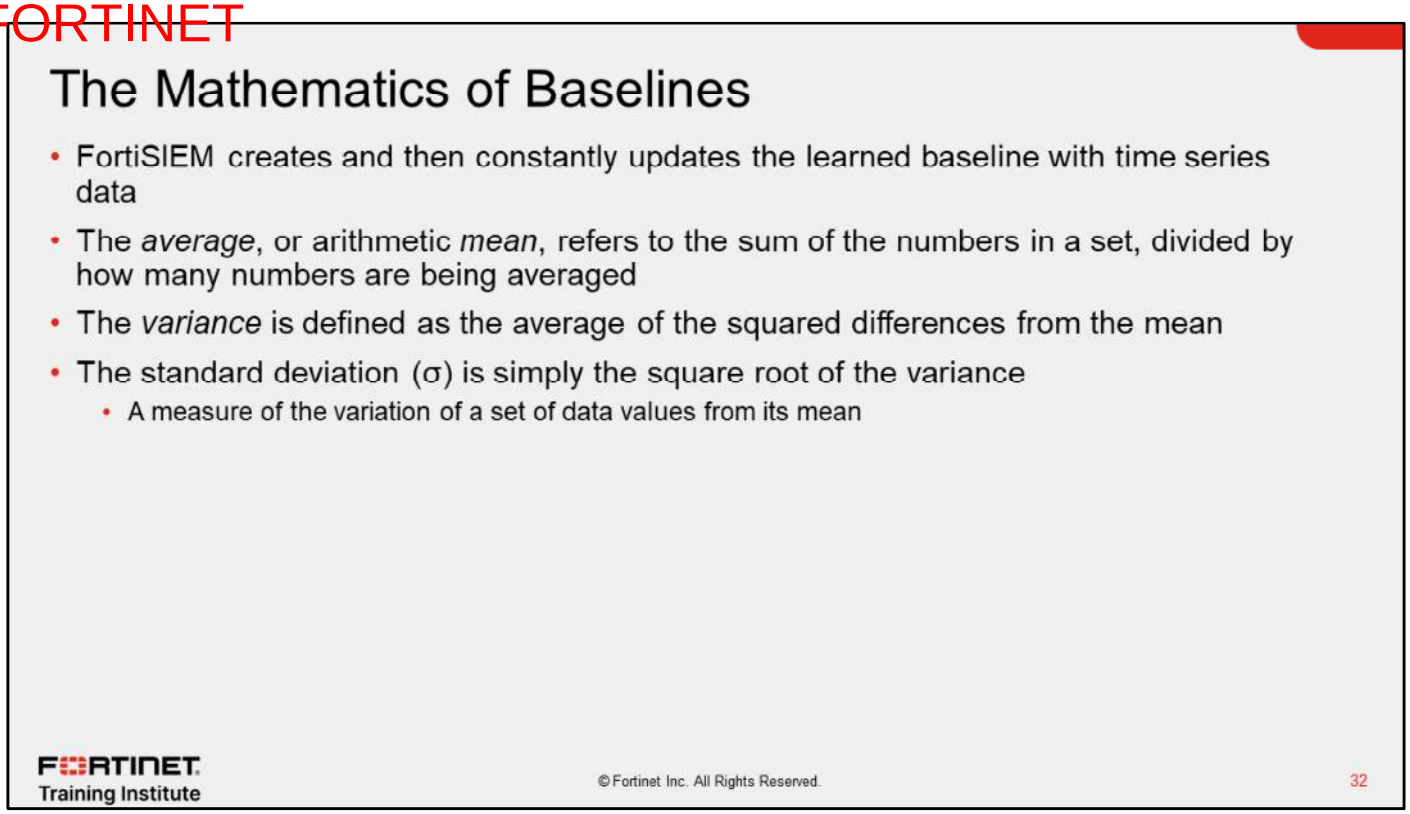
This slide compares various aggregate conditions and their equivalent terms in the SQLite profile table.

DO NOT REPRINT
© FORTINET



In this section, you will learn about calculating MIN, MAX, AVG, and STD DEV.

DO NOT REPRINT
© FORTINET

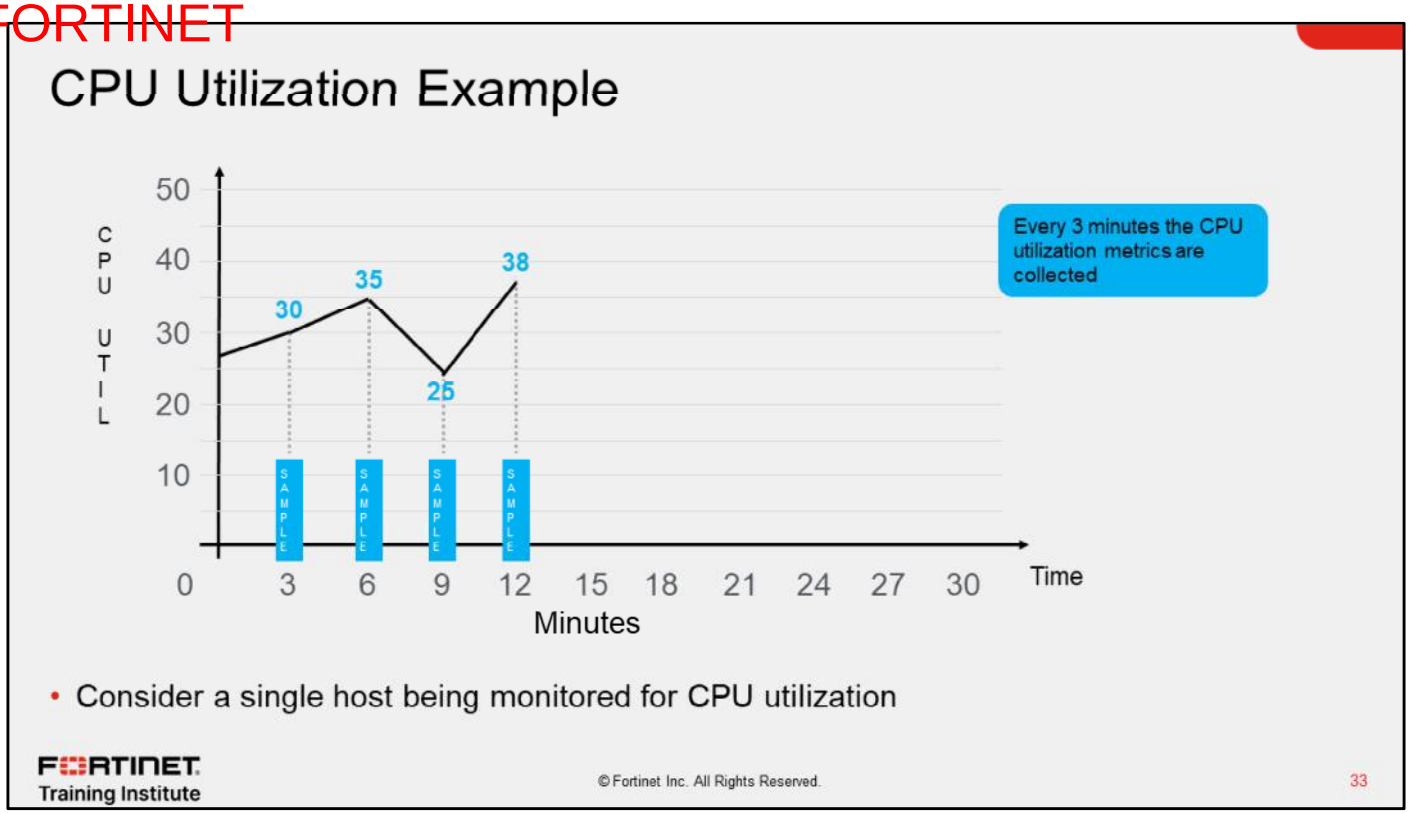


FortiSIEM creates and then constantly updates the learned baseline with time series data.

The *average*, or arithmetic *mean*, refers to the sum of the numbers in a set divided by how many numbers are being averaged. The *variance* is defined as the average of the squared differences from the mean.

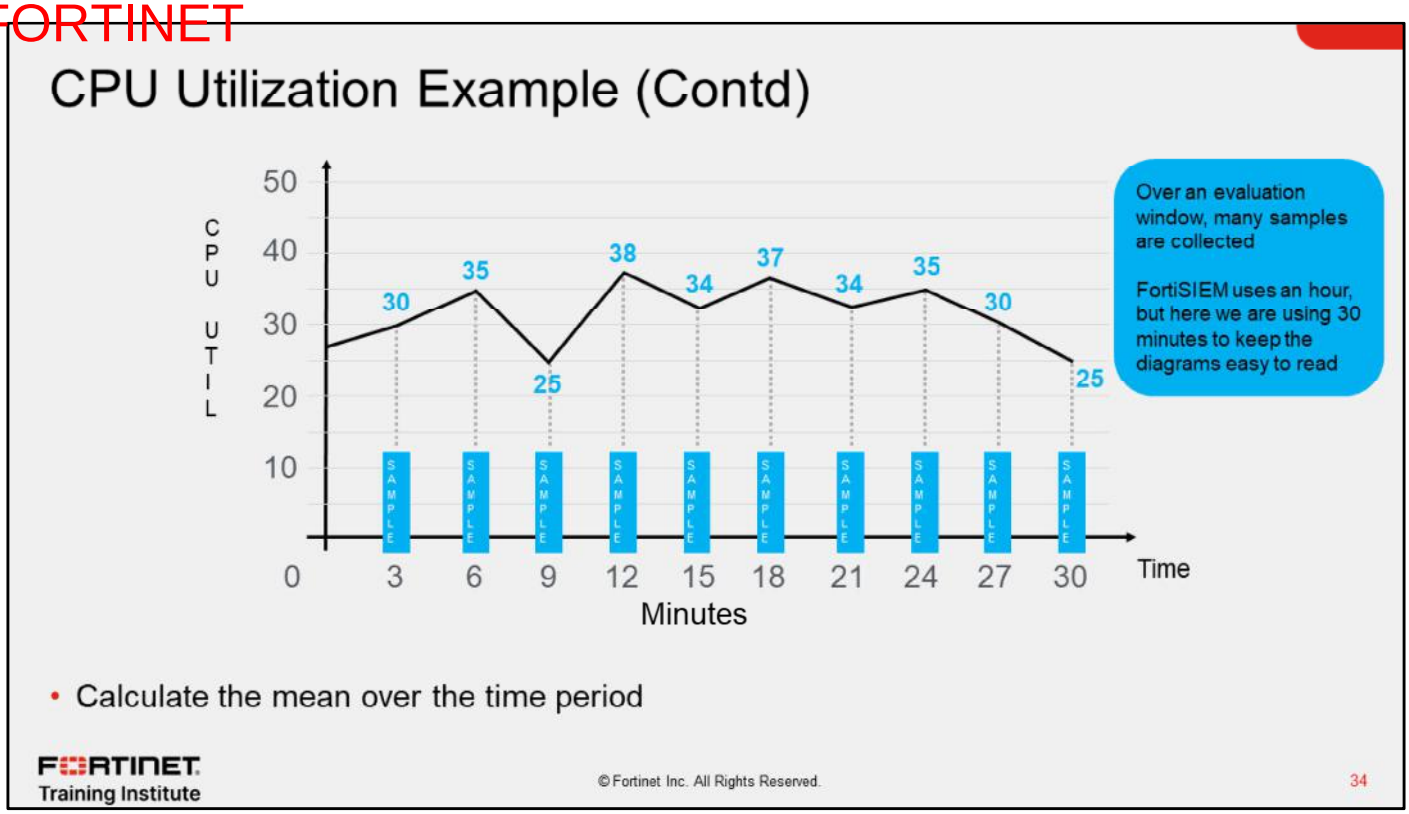
The standard deviation (σ) is the square root of the *variance*, which is basically a measure of the variation of a set of data values from its mean.

DO NOT REPRINT
© FORTINET



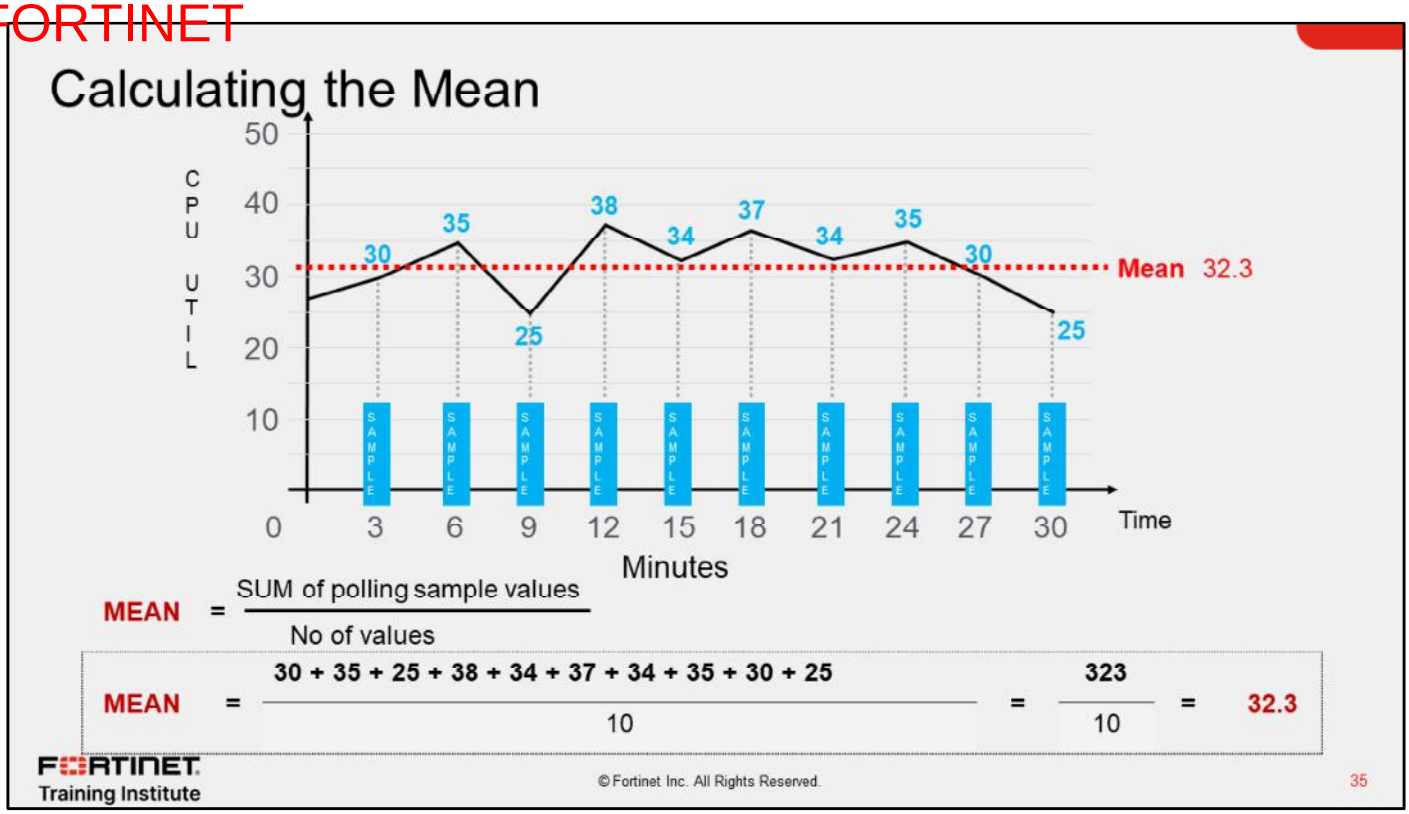
Consider the example shown on this slide, where a single host is being monitored for CPU utilization. Every 3 minutes the CPU utilization metrics are collected. To simplify the demonstration, assume that only a single sample is collected every 3 minutes.

DO NOT REPRINT
© FORTINET



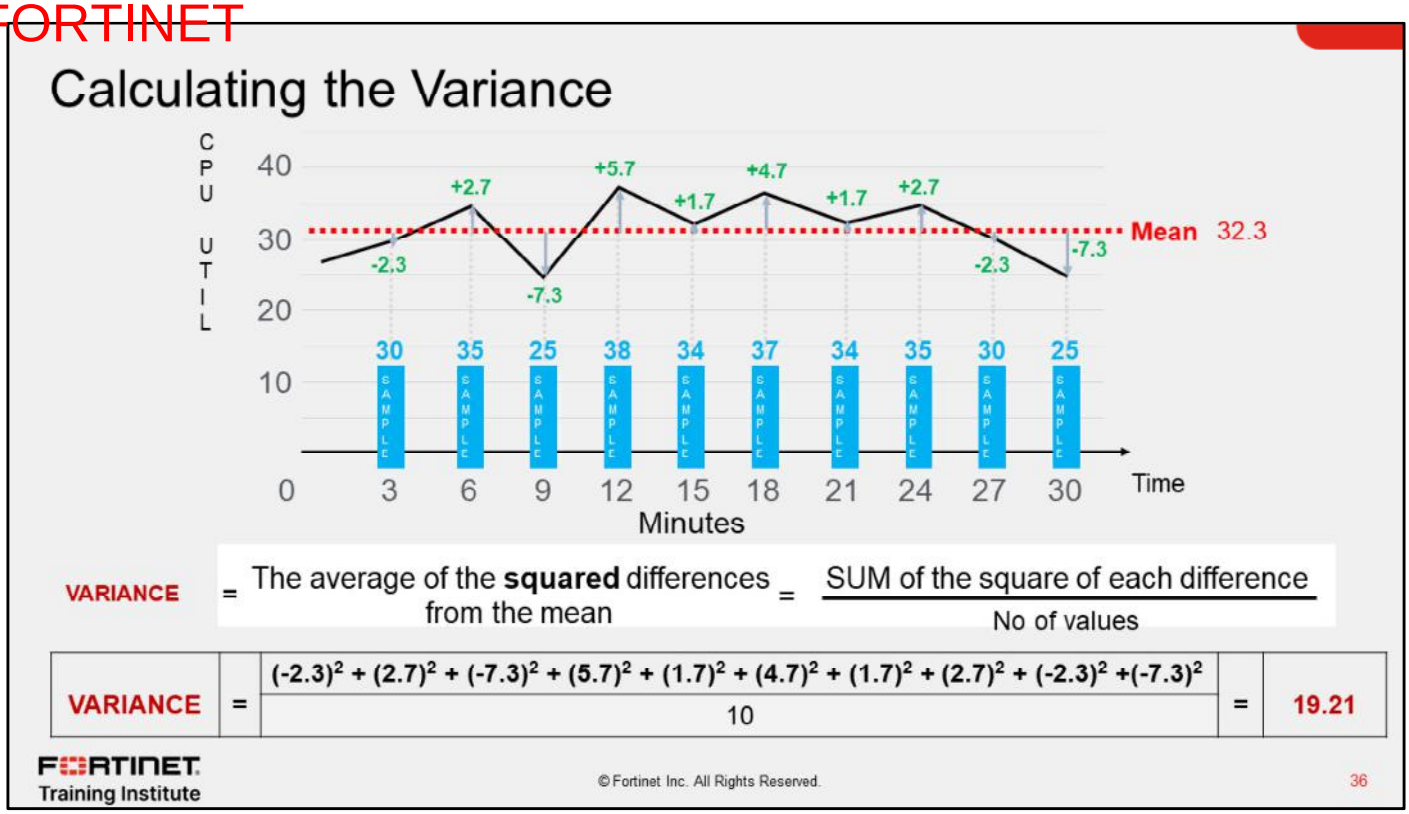
First, the mean value will be calculated over an evaluation window. Over an evaluation window, many samples are collected. FortiSIEM uses an hour for evaluating samples, but a 30-minute window will be used for demonstrating the calculations.

DO NOT REPRINT
© FORTINET



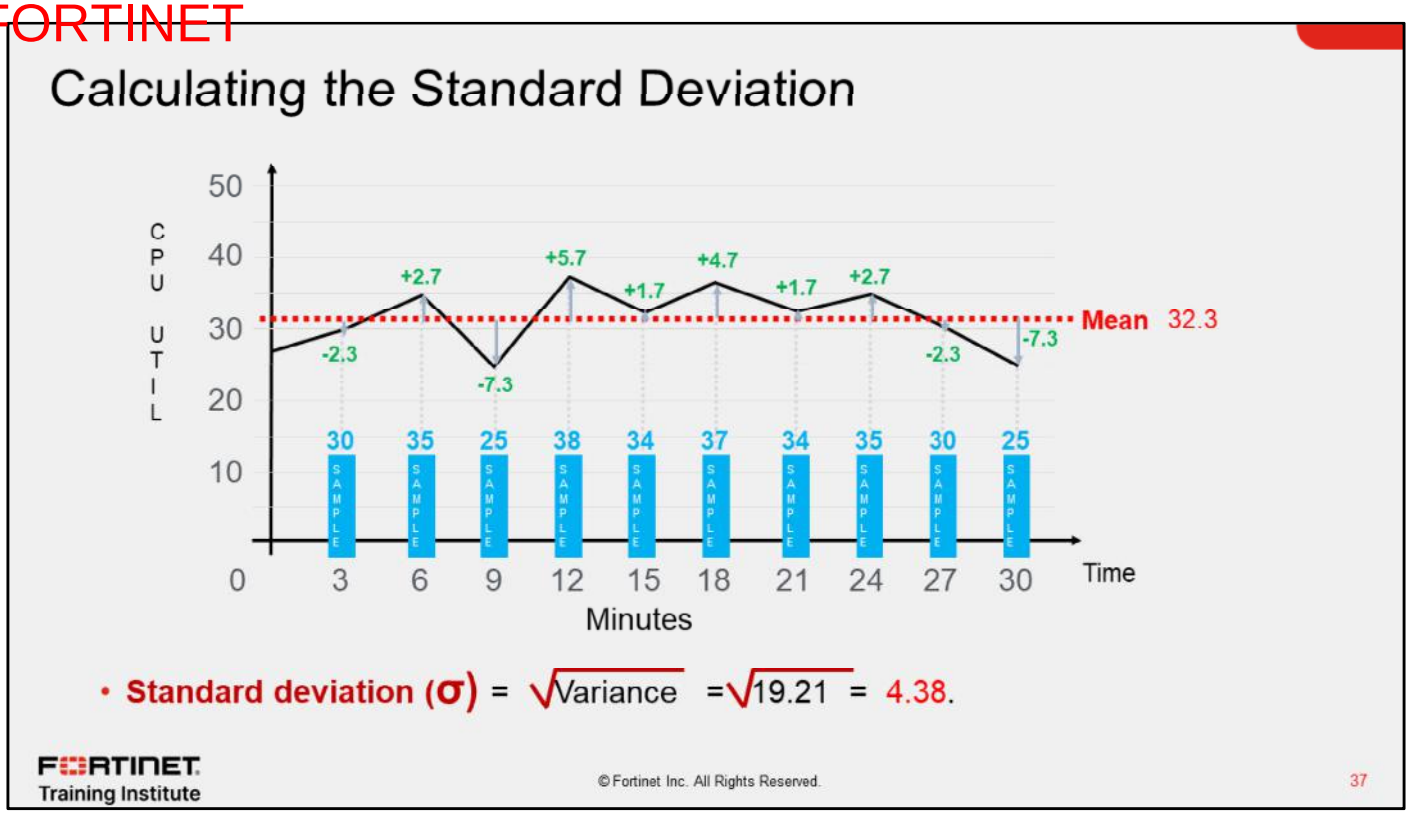
The mean is the sum of polling sample values divided by the number of values. In the example shown on this slide, there are ten samples with distinct CPU utilization values. When you add the values from each of these samples and divide by the number of samples, then you will get the mean value of 32.3.

DO NOT REPRINT
© FORTINET



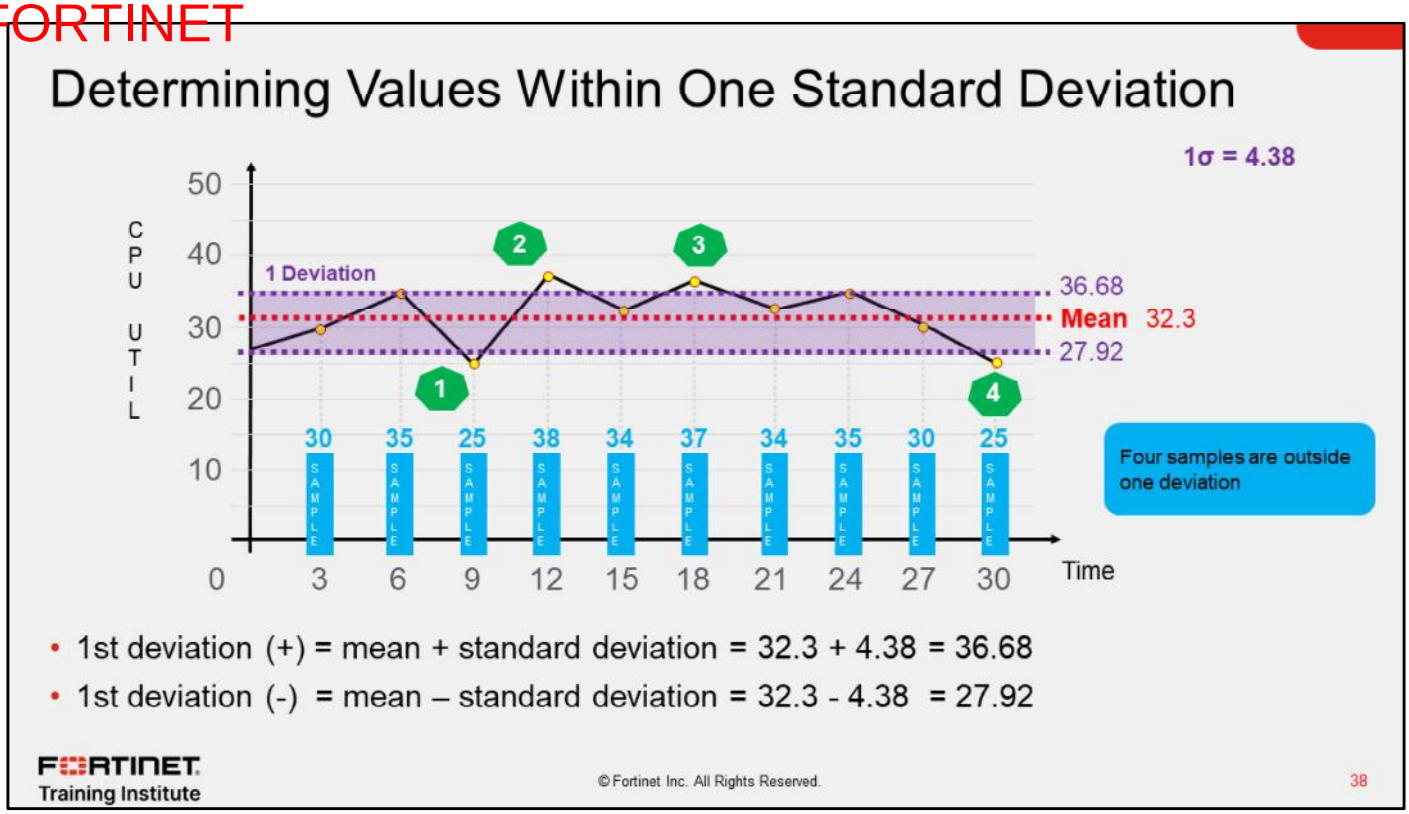
Variance is the average of the the squared differences from the mean. In order to calculate variance, you need to first calculate the variation of each sample from the mean value. If the CPU utilization is higher than the mean value, then it will be a positive difference. If the CPU utilization is below the mean value, then it will be a negative difference. You square all the difference values and add them. The sum is divided by the number of samples, which in this case is 10. The final result is the variance value.

DO NOT REPRINT
© FORTINET



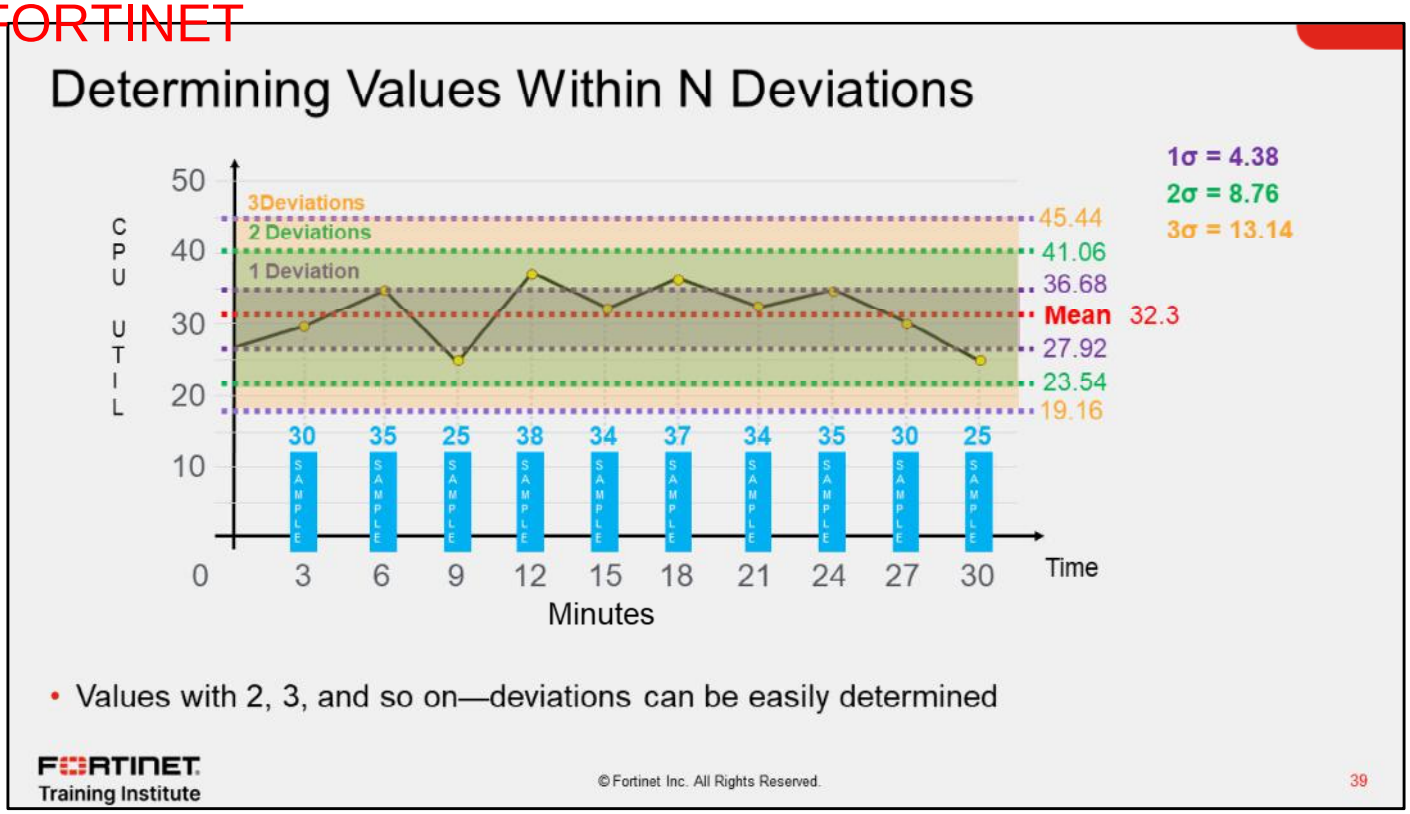
Standard deviation is the square root of the variance value.

DO NOT REPRINT
© FORTINET



Once you determine the standard deviation, you can determine the positive and negative deviation from the mean value. To determine the first positive deviation, you need to add the deviation value to the mean value. To determine the first negative deviation, you need to subtract the deviation from the mean value. You will notice that four samples are outside the first deviation.

DO NOT REPRINT
© FORTINET



You can determine the second and third deviation by multiplying the standard deviation by two and three respectively, and then calculating the positive and negative deviation from the mean by adding or subtracting that value from the mean.

DO NOT REPRINT
© FORTINET

Applying the Logic Back to FortiSIEM

- Since FortiSIEM uses hourly buckets, values are recorded for each profile every hour (AVG, MIN, MAX, Standard Deviation)
- Hence, the 24 buckets for a weekday, and 24 for weekend; one per hour



- Since there is only one data point, the results are real values

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

40

Since FortiSIEM uses hourly buckets, values are recorded for each profile every hour. In the backend, FortiSIEM calculates the minimum, maximum, average, and standard deviation for attributes that are related to each profile and stores those in buckets every hour. Hence, there are 24 buckets for a weekday, and 24 buckets for weekends; one per hour. Since there is only one data point, the results are real values.

DO NOT REPRINT
© FORTINET

Day One—Real Values Are Stored Hourly in the Daily Database



Hour Of Day	Host IP	Host Name	Min CPU Util	AVG CPU Util	Max CPU Util	Std Dev CPU Util	numPoints
9	1.1.1.1	ServerA	32.31	32.31	32.31	0	1
10	1.1.1.1	ServerA	32.52	32.52	32.52	0	1
11	1.1.1.1	ServerA	45.10	45.10	45.10	0	1
12	1.1.1.1	ServerA	43.61	43.61	43.61	0	1
13	1.1.1.1	ServerA	33.71	33.71	33.71	0	1
14	1.1.1.1	ServerA	28.10	28.10	28.10	0	1
15	1.1.1.1	ServerA	42.01	42.01	42.01	0	1

FORTINET Training Institute

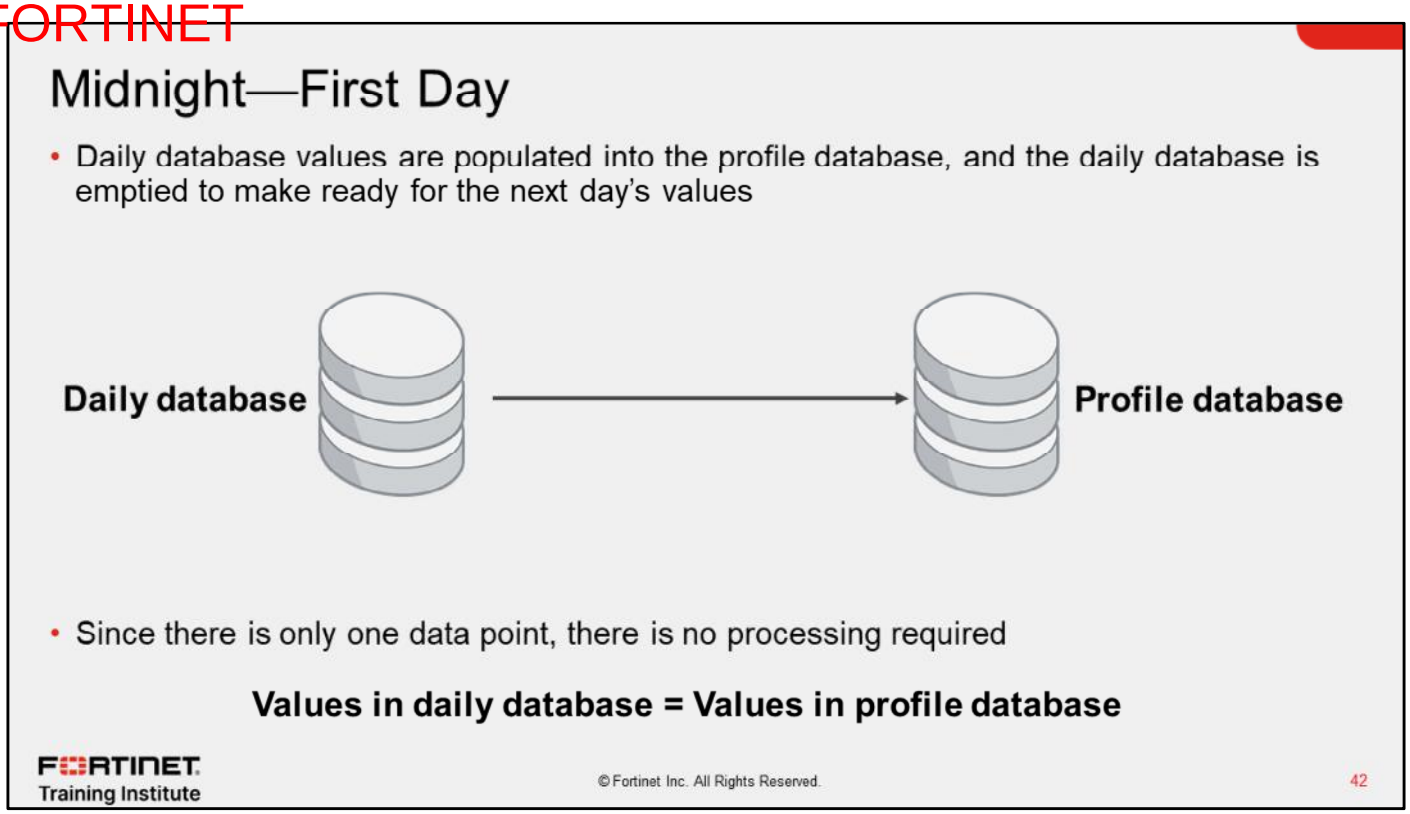
© Fortinet Inc. All Rights Reserved.

41

This is an example of the CPU utilization that is reported by a server on the first day. Every value that is reported for every hour of the day is considered a real value. Hence the minimum, maximum, and average will be the same value since there is no comparison with any other values against the daily database.

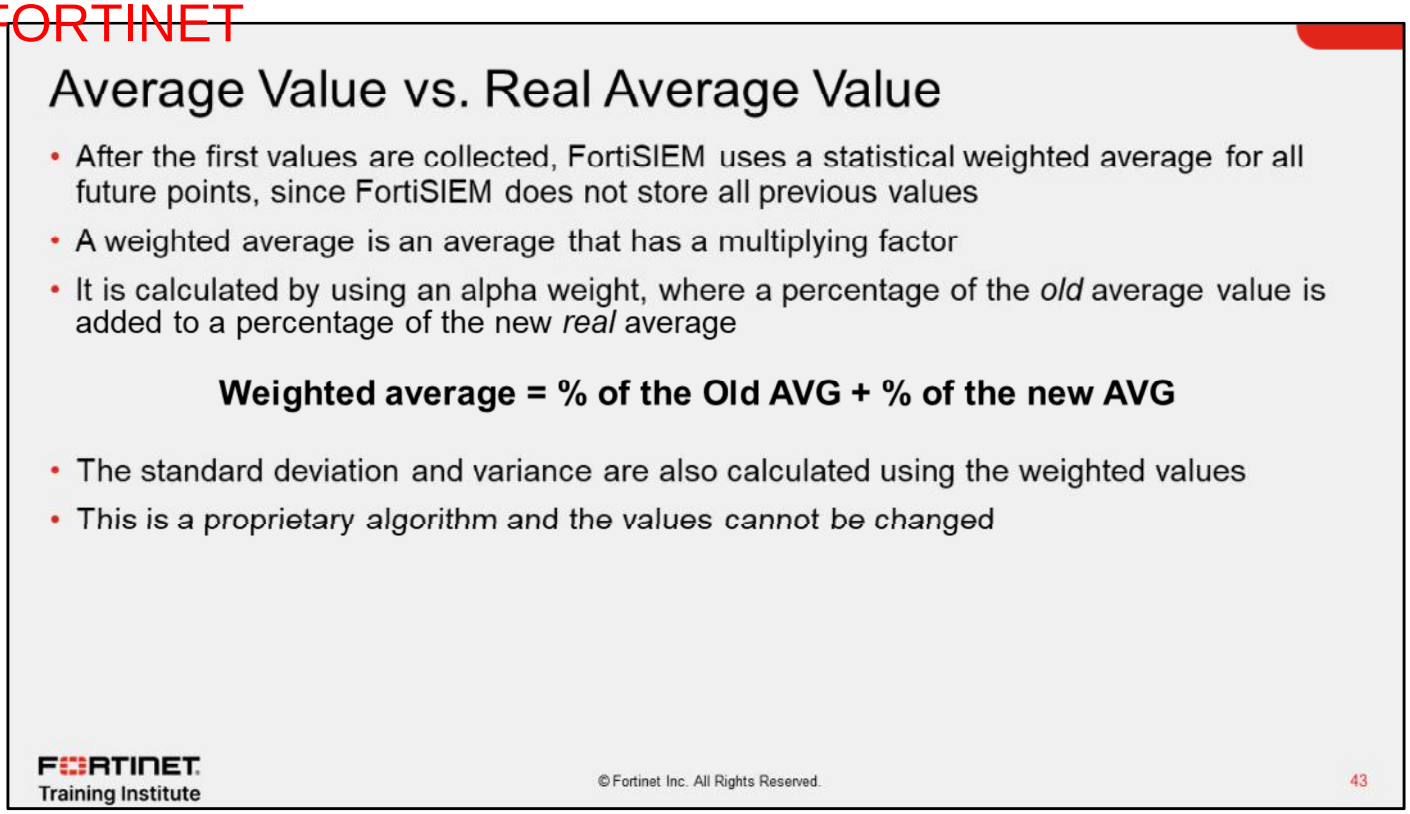
The standard deviation will be zero since there is no secondary data to compare the deviation for every hour. The numPoints column will update each hour of day with 1 since there is only one value in the database for that hour.

DO NOT REPRINT
© FORTINET



At midnight on the first day, the daily database values are populated into the profile database, and the daily database is purged to make ready for the next day's values. After the first day, the values in the profile database will be the same as the values in the daily database. There is no processing of data since there is only one data point for each hour.

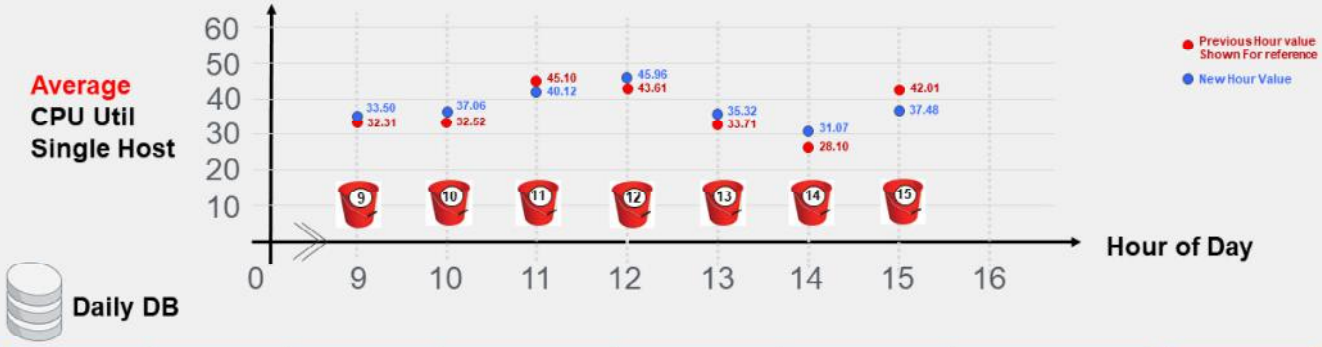
DO NOT REPRINT
© FORTINET



You might think that the average is just the real average value, but that's not true. FortiSIEM uses a statistical weighted average for all future points, since FortiSIEM does not store all previous values. A weighted average is an average that has a multiplying factor. It is calculated by using an alpha weight, where a percentage of the *old* average value is added to a percentage of the new *real* average value. The standard deviation and variance are also calculated using the weighted values. This is a proprietary algorithm and the values cannot be changed.

DO NOT REPRINT
© FORTINET

Day Two—Real Value Stored Hourly in Daily Database



Hour Of Day	Host IP	Host Name	Min CPU Util	AVG CPU Util	Max CPU Util	Std Dev CPU Util	numPoints
9	1.1.1.1	ServerA	33.50	33.50	33.50	0	1
10	1.1.1.1	ServerA	37.06	37.06	37.06	0	1
11	1.1.1.1	ServerA	40.12	40.12	40.12	0	1
12	1.1.1.1	ServerA	45.96	45.96	45.96	0	1
13	1.1.1.1	ServerA	35.32	35.32	35.32	0	1
14	1.1.1.1	ServerA	31.07	31.07	31.07	0	1
15	1.1.1.1	ServerA	37.48	37.48	37.48	0	1

FORTINET Training Institute

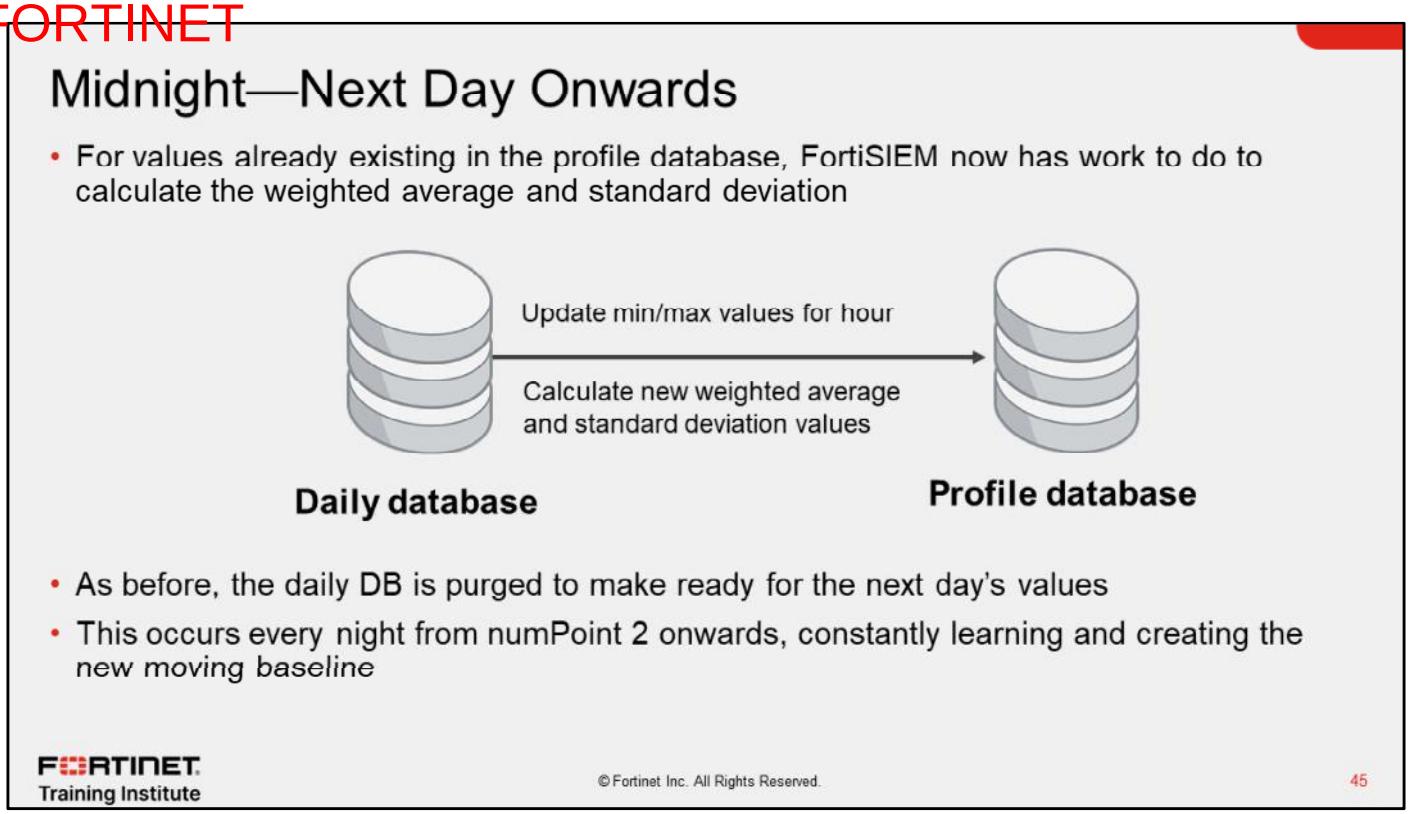
© Fortinet Inc. All Rights Reserved.

44

On the second day, the daily database is again populated with real data values. On this slide, the data values for the previous day are shown for reference only.

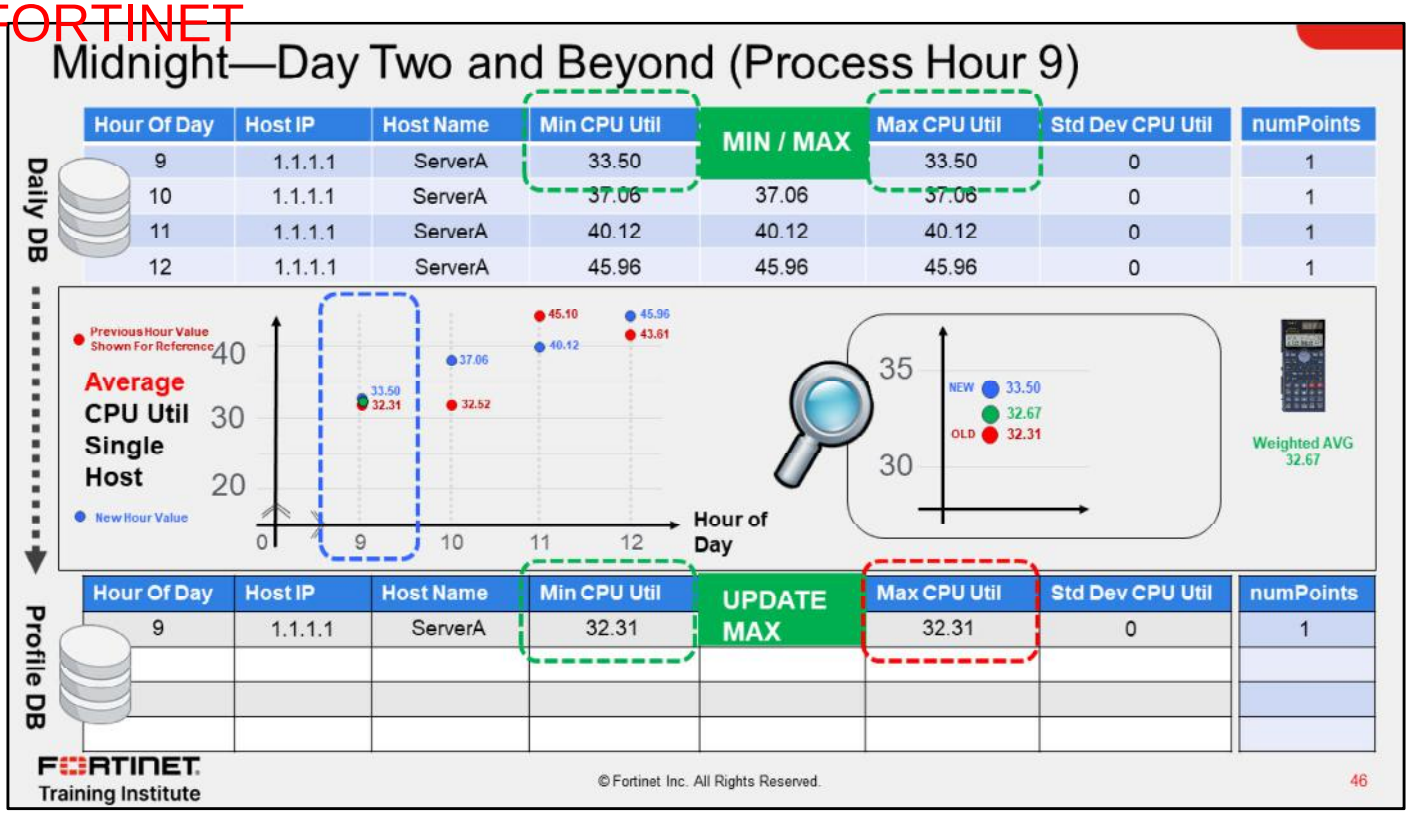
Since the daily database has been emptied at midnight, it does not have the data values for the previous day. It will populate the minimum, maximum, average, and standard deviation as if it is receiving this data for the first time. So the minimum, maximum, and average will be the same value with standard deviation as 0. The numPoints column will be updated again with 1.

DO NOT REPRINT
© FORTINET



At midnight on the second day, FortiSIEM will need to perform some calculations before storing those values in the profile database. It will update the minimum and maximum values for every hour of the day by comparing the values stored in the profile database for day one and those values that are currently in the daily database of the second day. It will also calculate the new weighted average and standard deviation values. After performing those calculations, it will store those values in the profile database and update the numPoints field in the profile database to 2. This means that the data stored is based on two data points for every hour. In this way FortiSIEM is constantly learning and creating new baselines in the profile database. The daily database will be emptied for the third day.

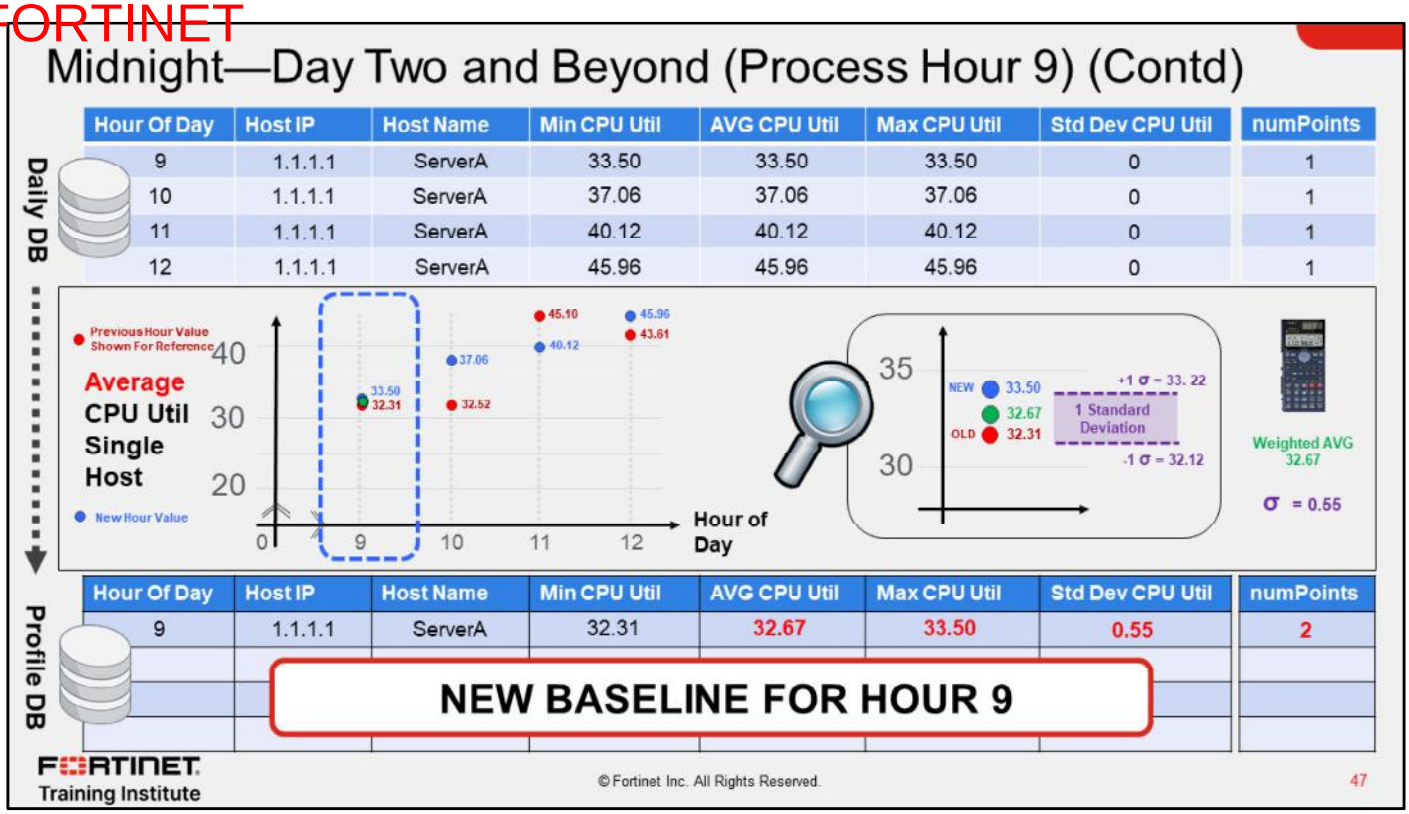
DO NOT REPRINT
© FORTINET



Now, you will analyze the calculation that is performed in the back-end by FortiSIEM for hour 9 at midnight.

The daily database contains the new CPU utilization data for day two, which is 33.50. The profile database has the data for the same hour from the previous day, which is 32.31. The weighted average between day one and day two is 32.67. FortiSIEM compares the minimum and maximum value between the daily database and the profile database. It concludes that the minimum value does not need to be changed and it will keep that value as 32.31. The profile database which contains the previous values for hour 9, will be updated with the new learned baseline.

DO NOT REPRINT
© FORTINET

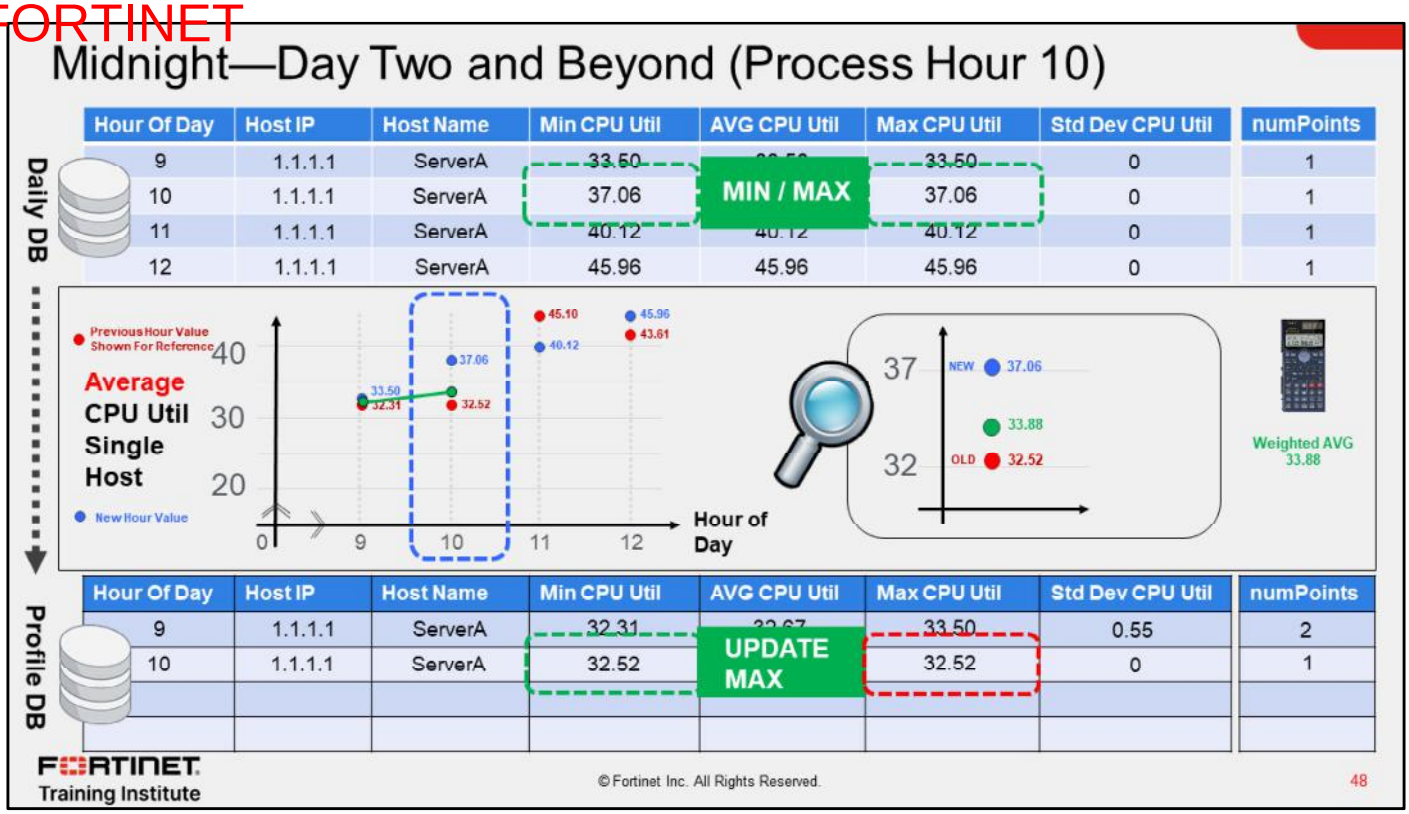


Hour Of Day	Host IP	Host Name	Min CPU Util	AVG CPU Util	Max CPU Util	Std Dev CPU Util	numPoints
9	1.1.1.1	ServerA	33.50	33.50	33.50	0	1
10	1.1.1.1	ServerA	37.06	37.06	37.06	0	1
11	1.1.1.1	ServerA	40.12	40.12	40.12	0	1
12	1.1.1.1	ServerA	45.96	45.96	45.96	0	1

Hour Of Day	Host IP	Host Name	Min CPU Util	AVG CPU Util	Max CPU Util	Std Dev CPU Util	numPoints
9	1.1.1.1	ServerA	32.31	32.67	33.50	0.55	2

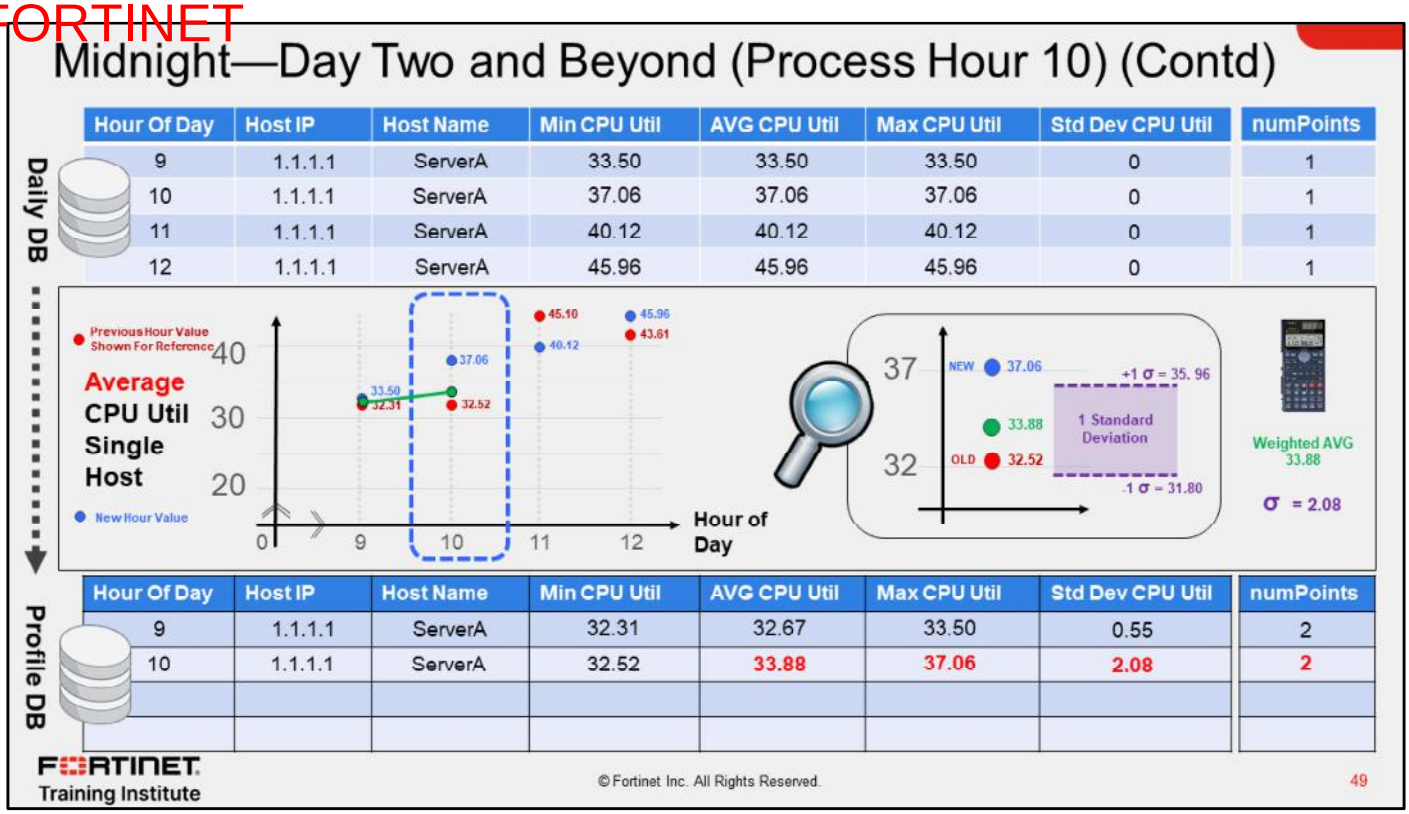
The maximum CPU utilization value for hour 9 needs to be updated in the profile database. It will update the maximum value in the profile database for hour 9 to 33.50. It will also update the average CPU utilization value for that hour to 32.67. The numPoints column will be updated to 2 because the calculation is based on two data points. It will then calculate the standard deviation between those two data points and update that to 0.55. That will be the new baseline for hour 9.

DO NOT REPRINT
© FORTINET



FortiSIEM will perform similar calculation for hour 10. The daily database contains the new CPU utilization data for day two, which is 37.06. The profile database has the data for the same hour from the previous day, which is 32.52. The weighted average between day one and day two is 33.88. FortiSIEM compares the minimum and maximum value between the daily database and the profile database. It concludes that the minimum value does not need to be changed and it will keep that value as 32.52.

DO NOT REPRINT
© FORTINET

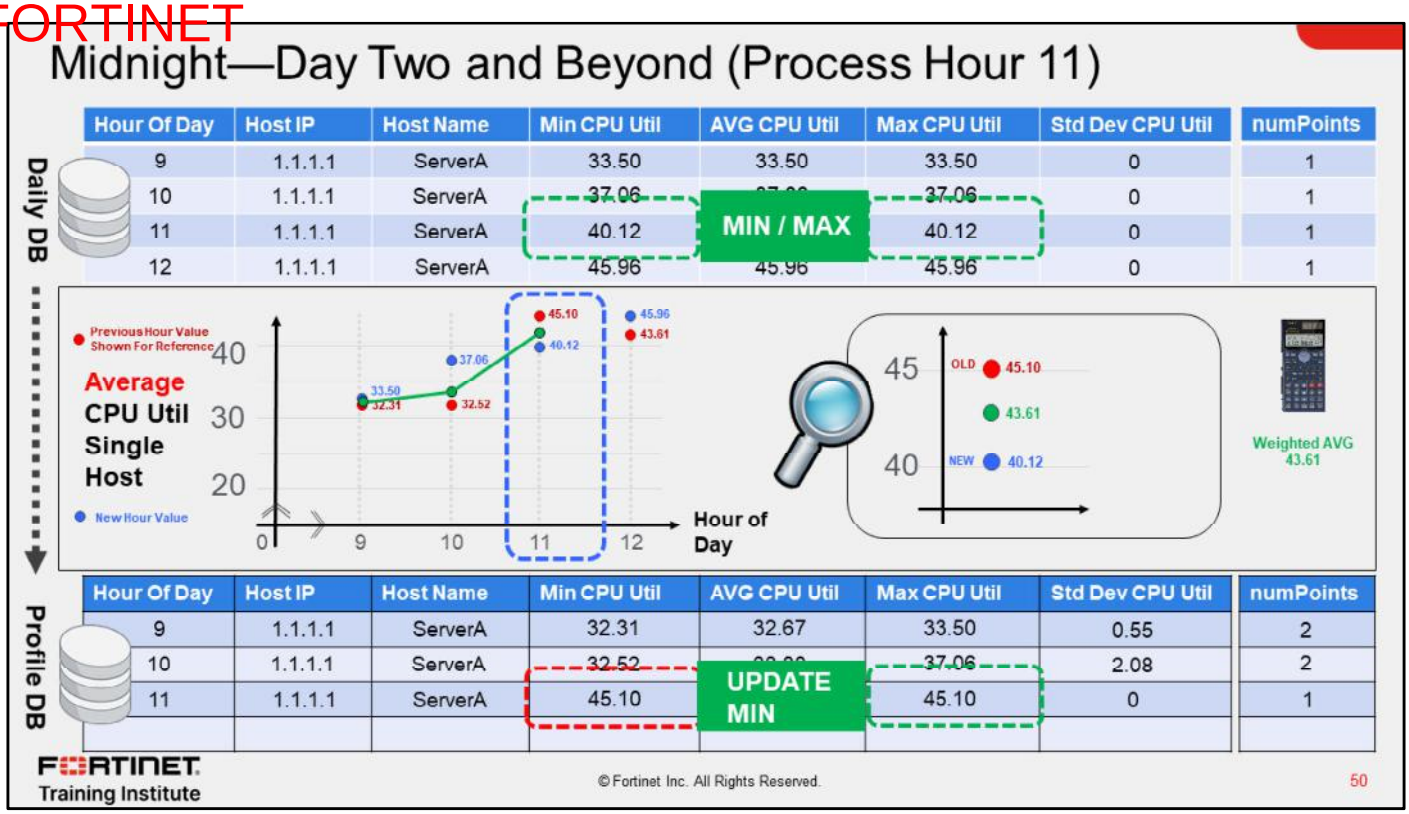


Hour Of Day	Host IP	Host Name	Min CPU Util	AVG CPU Util	Max CPU Util	Std Dev CPU Util	numPoints
9	1.1.1.1	ServerA	33.50	33.50	33.50	0	1
10	1.1.1.1	ServerA	37.06	37.06	37.06	0	1
11	1.1.1.1	ServerA	40.12	40.12	40.12	0	1
12	1.1.1.1	ServerA	45.96	45.96	45.96	0	1

Hour Of Day	Host IP	Host Name	Min CPU Util	AVG CPU Util	Max CPU Util	Std Dev CPU Util	numPoints
9	1.1.1.1	ServerA	32.31	32.67	33.50	0.55	2
10	1.1.1.1	ServerA	32.52	33.88	37.06	2.08	2

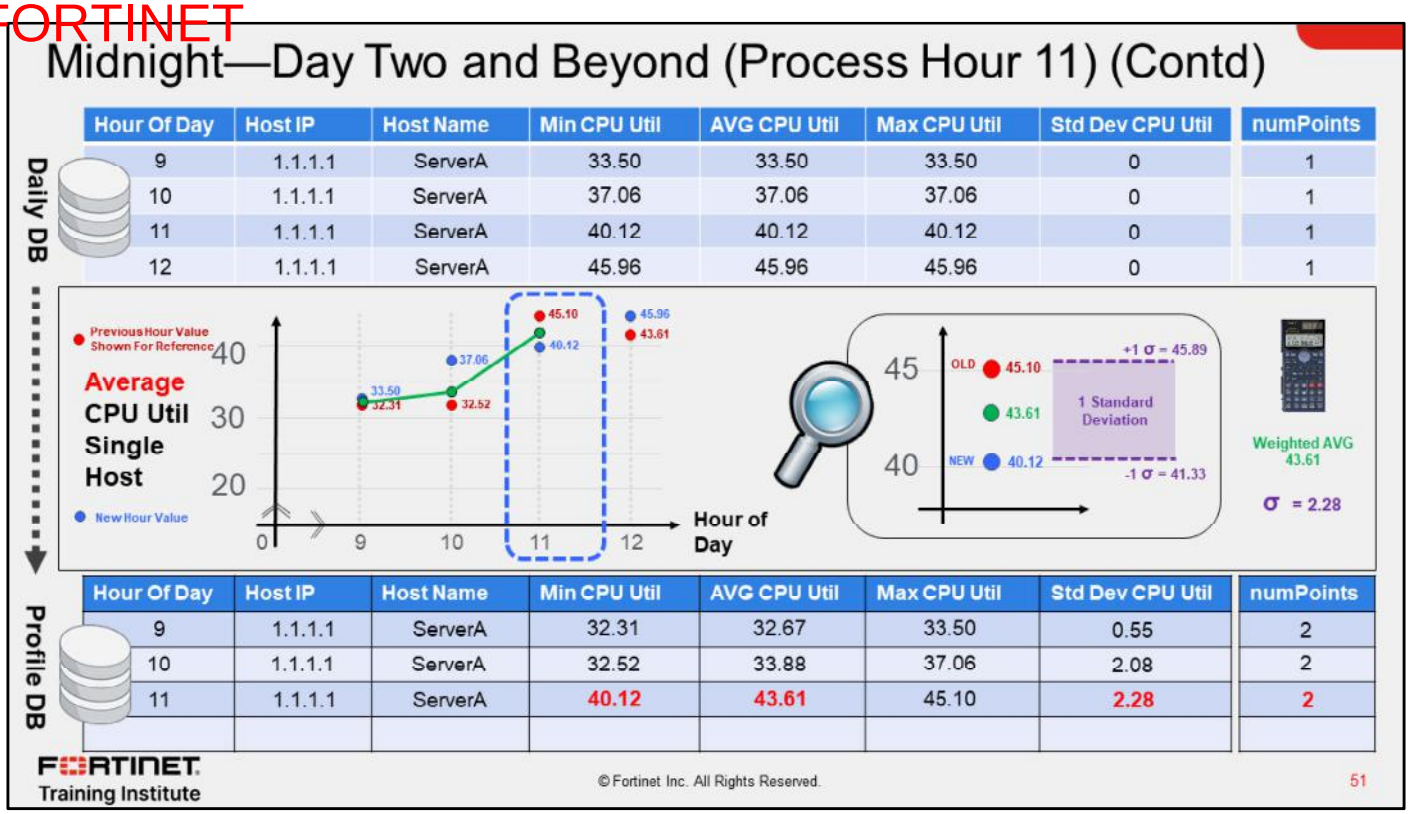
The maximum CPU utilization value for hour 10 needs to be updated in the profile database. It will update the maximum value in the profile database for hour 10 to 33.88. It will also update the average CPU utilization value for that hour to 37.06. The numPoints column will be updated to 2 because the calculation is based on two data points. It will then calculate the standard deviation between those two data points and update that to 2.08. That will be the new baseline for hour 10.

DO NOT REPRINT
© FORTINET



FortiSIEM will perform similar calculations for hour 11. The daily database contains the new CPU utilization data for day two, which is 40.12. The profile database has the data for the same hour from the previous day, which is 45.10. The weighted average between day one and day two is 43.61. FortiSIEM compares the minimum and maximum value between the daily database and the profile database. It concludes that the maximum value does not need to be changed and it will keep that value as 45.10.

DO NOT REPRINT
© FORTINET

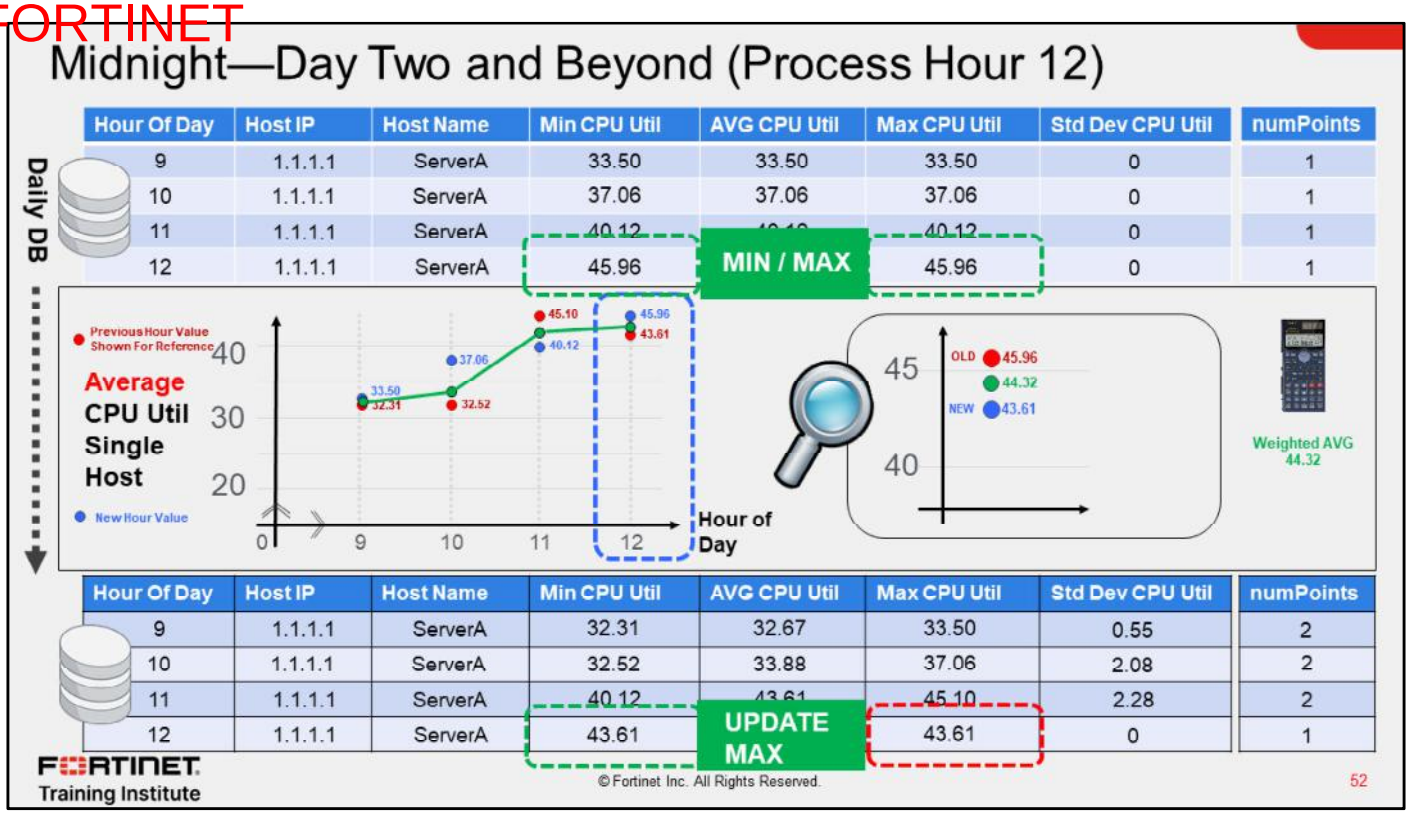


Hour Of Day	Host IP	Host Name	Min CPU Util	AVG CPU Util	Max CPU Util	Std Dev CPU Util	numPoints
9	1.1.1.1	ServerA	33.50	33.50	33.50	0	1
10	1.1.1.1	ServerA	37.06	37.06	37.06	0	1
11	1.1.1.1	ServerA	40.12	40.12	40.12	0	1
12	1.1.1.1	ServerA	45.96	45.96	45.96	0	1

Hour Of Day	Host IP	Host Name	Min CPU Util	AVG CPU Util	Max CPU Util	Std Dev CPU Util	numPoints
9	1.1.1.1	ServerA	32.31	32.67	33.50	0.55	2
10	1.1.1.1	ServerA	32.52	33.88	37.06	2.08	2
11	1.1.1.1	ServerA	40.12	43.61	45.10	2.28	2

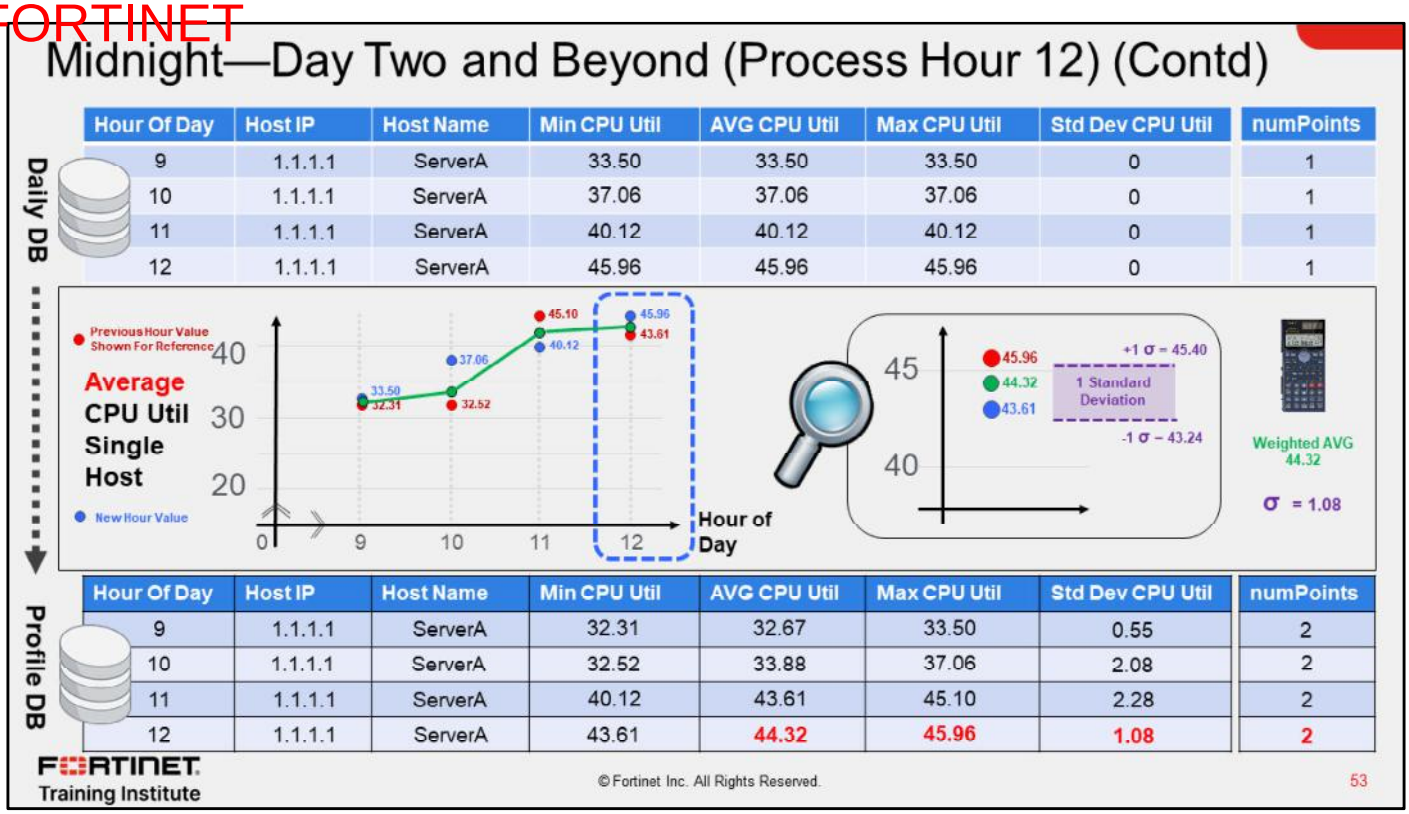
The minimum CPU utilization value for hour 11 needs to be updated in the profile database. It will update the minimum value in the profile database for hour 11 to 40.12. It will also update the average CPU utilization value for that hour to 43.61. The numPoints column will be updated to 2 because the calculation is based on two data points. It will then calculate the standard deviation between those two data points and update that to 2.28. That will be the new baseline for hour 11.

DO NOT REPRINT
© FORTINET



FortiSIEM will perform similar calculations for hour 12. The daily database contains the new CPU utilization data for day two, which is 43.61. The profile database has the data for the same hour from the previous day, which is 45.96. The weighted average between day one and day two is 44.32. FortiSIEM compares the minimum and maximum value between the daily database and the profile database. It concludes that the minimum value does not need to be changed and it will keep that value as 43.61.

DO NOT REPRINT
© FORTINET

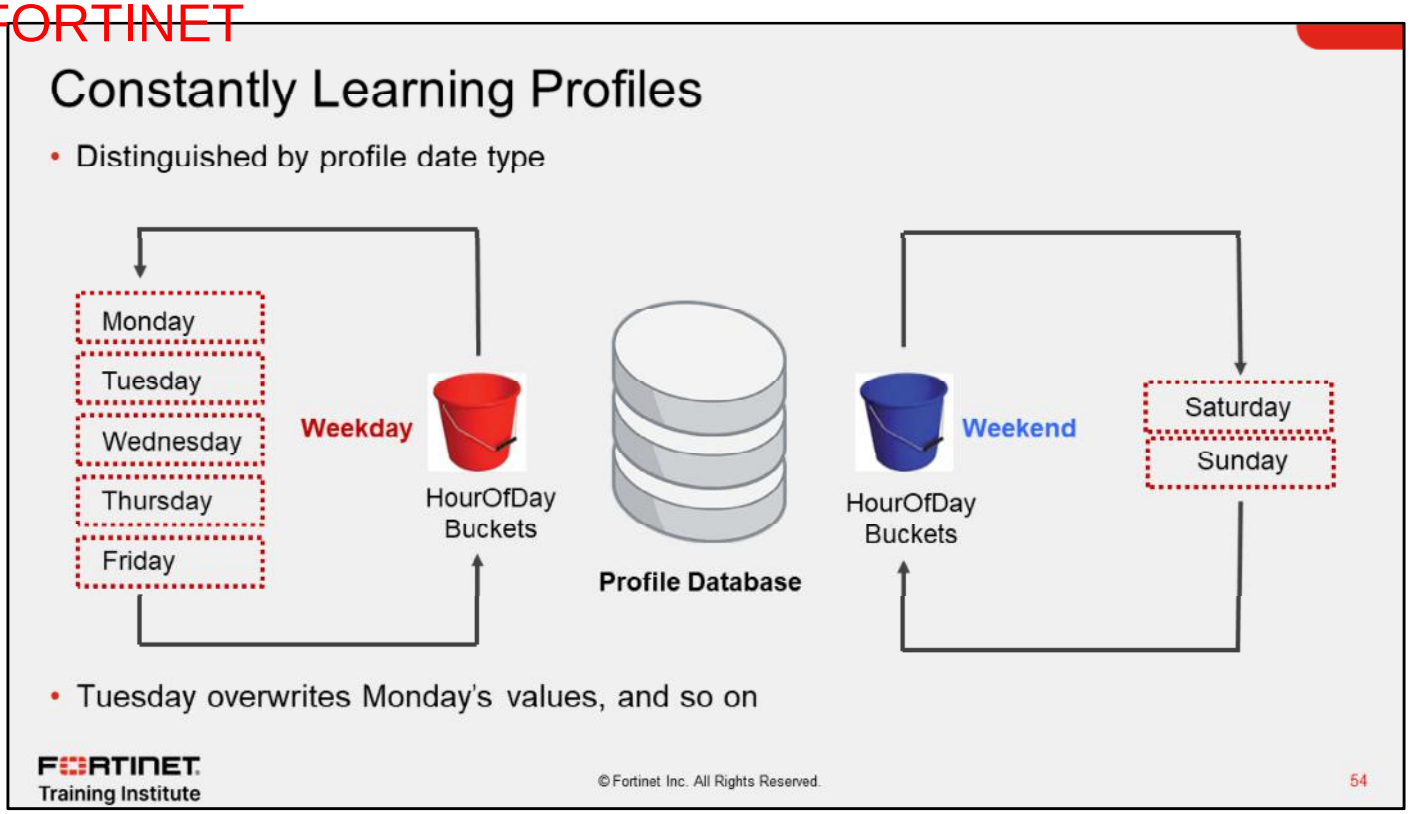


Hour Of Day	Host IP	Host Name	Min CPU Util	AVG CPU Util	Max CPU Util	Std Dev CPU Util	numPoints
9	1.1.1.1	ServerA	33.50	33.50	33.50	0	1
10	1.1.1.1	ServerA	37.06	37.06	37.06	0	1
11	1.1.1.1	ServerA	40.12	40.12	40.12	0	1
12	1.1.1.1	ServerA	45.96	45.96	45.96	0	1

Hour Of Day	Host IP	Host Name	Min CPU Util	AVG CPU Util	Max CPU Util	Std Dev CPU Util	numPoints
9	1.1.1.1	ServerA	32.31	32.67	33.50	0.55	2
10	1.1.1.1	ServerA	32.52	33.88	37.06	2.08	2
11	1.1.1.1	ServerA	40.12	43.61	45.10	2.28	2
12	1.1.1.1	ServerA	43.61	44.32	45.96	1.08	2

The maximum CPU utilization value for hour 12 needs to be updated in the profile database. It will update the maximum value in the profile database for hour 12 to 45.96. It will also update the average CPU utilization value for that hour to 44.32. The numPoints column will be updated to 2 because the calculation is based on two data points. It will then calculate the standard deviation between those two data points and update that to 1.08. That will be the new baseline for hour 12.

DO NOT REPRINT
© FORTINET

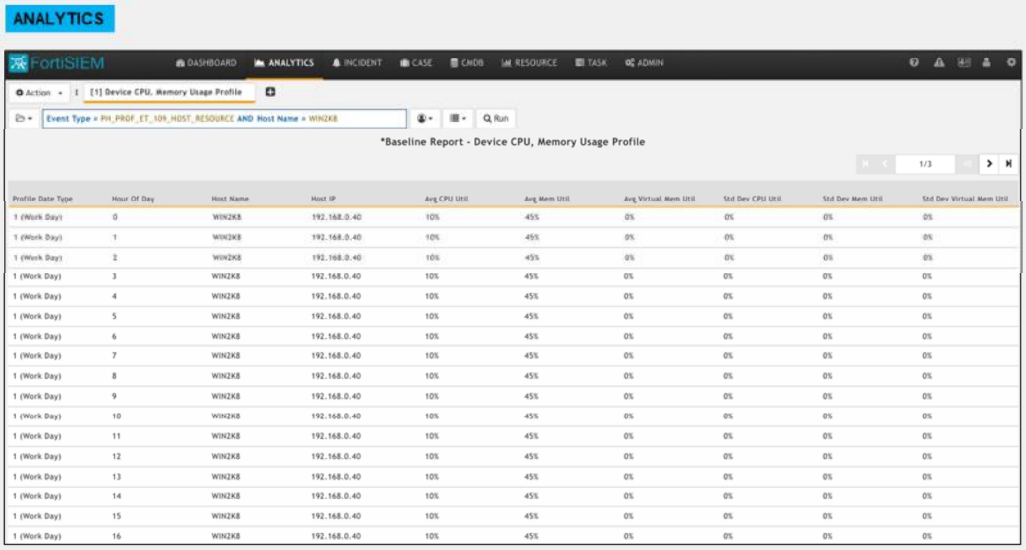


FortiSIEM is constantly learning new baselines every day. For weekdays, the data from Tuesday overwrites the data from Monday, and so on. For weekends, the data from Sunday overwrites the data from Saturday.

DO NOT REPRINT
© FORTINET

Baseline Report Sample—Device CPU Profile

- The baseline reports return moving data per hour of day



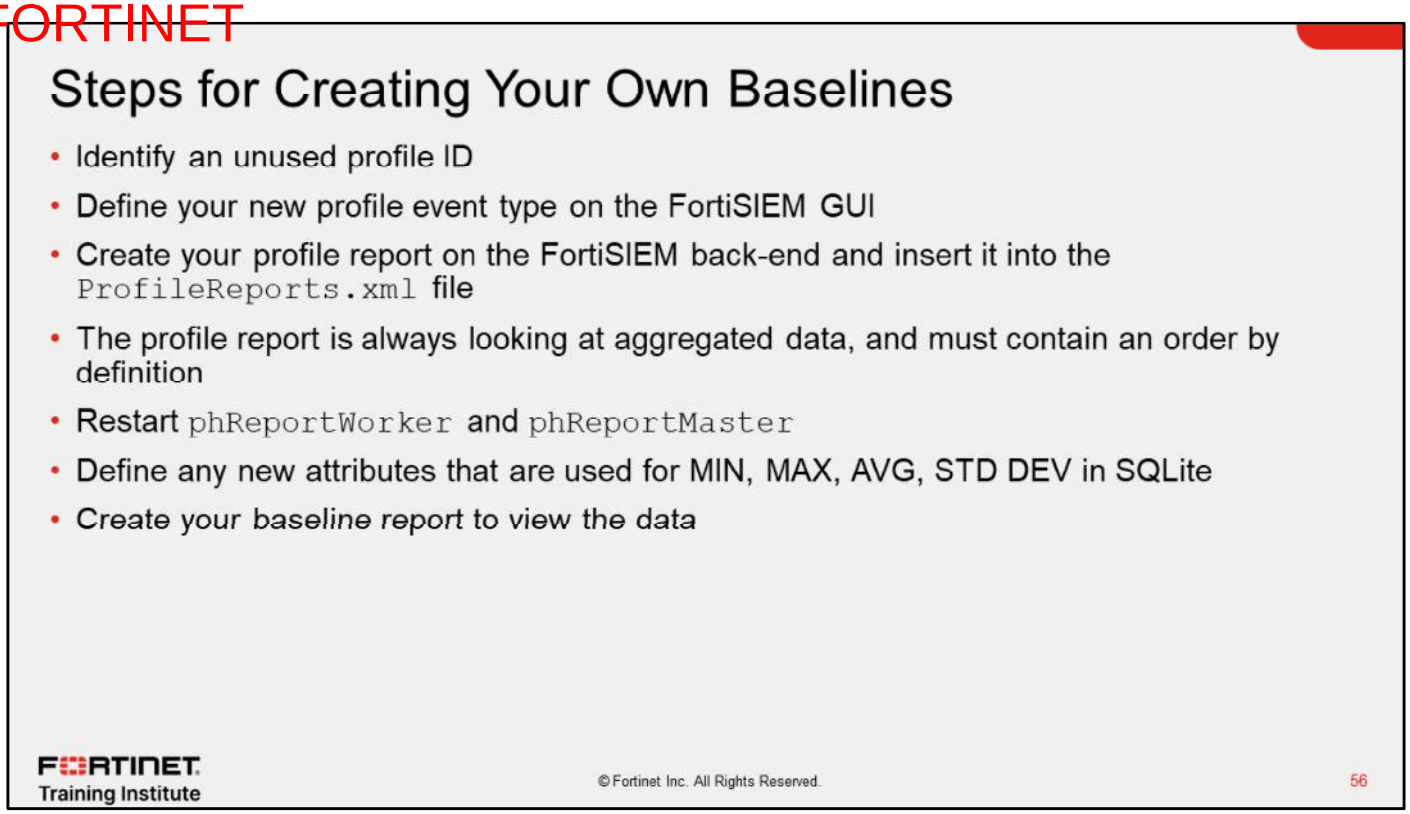
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

55

You can view the baseline data from the profile database for individual profiles by running the baseline profile report. You can view the current baseline of minimum, maximum, average, and standard deviation values for individual profiles.

DO NOT REPRINT
© FORTINET

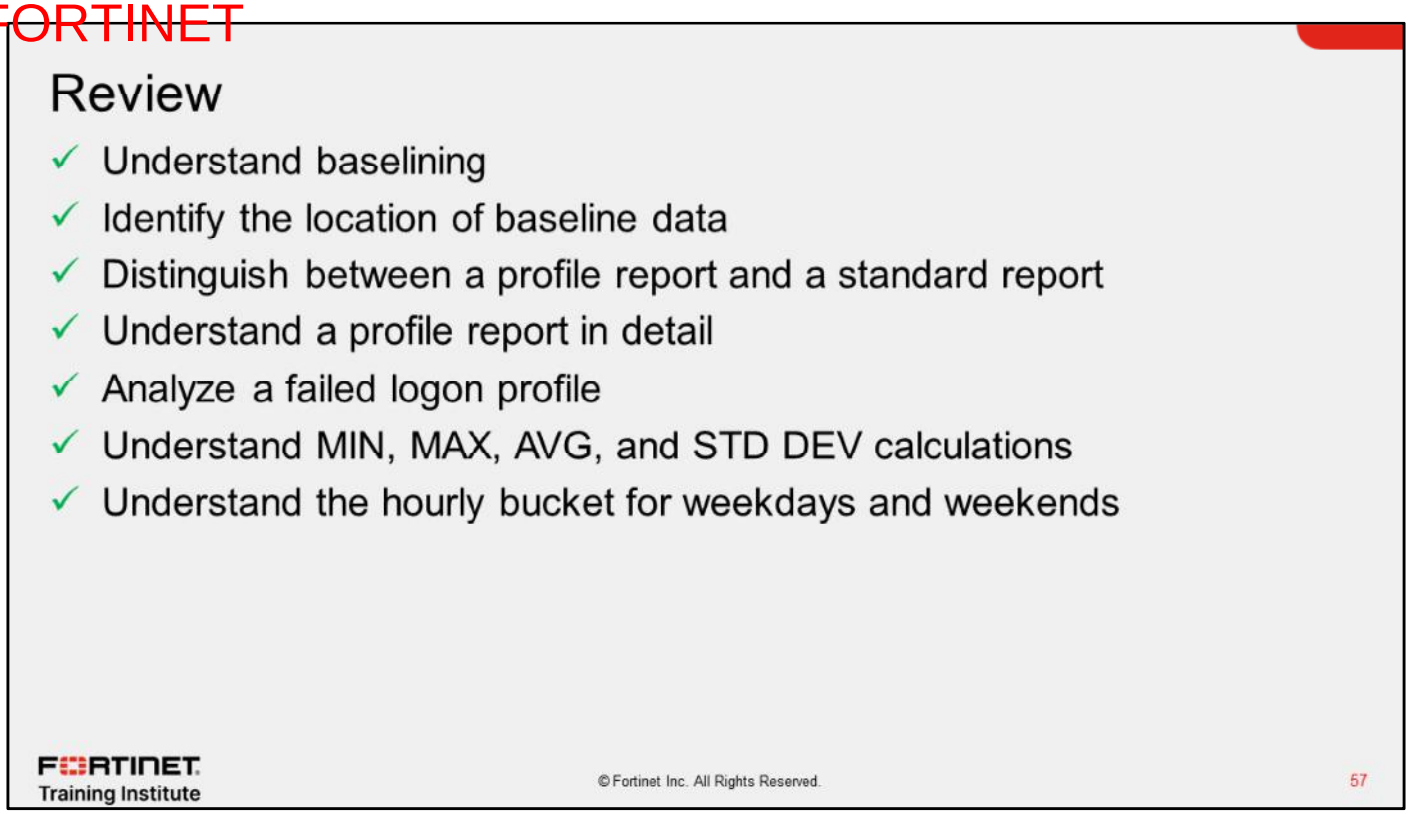


To create your own baseline profile, you have to identify an unused profile ID that you can use. Define your new profile event type on the FortiSIEM GUI, and then create your profile report on the FortiSIEM back end and insert it into the `ProfileReports.xml` file.

The profile report is always looking at aggregated data and must contain an order by definition.

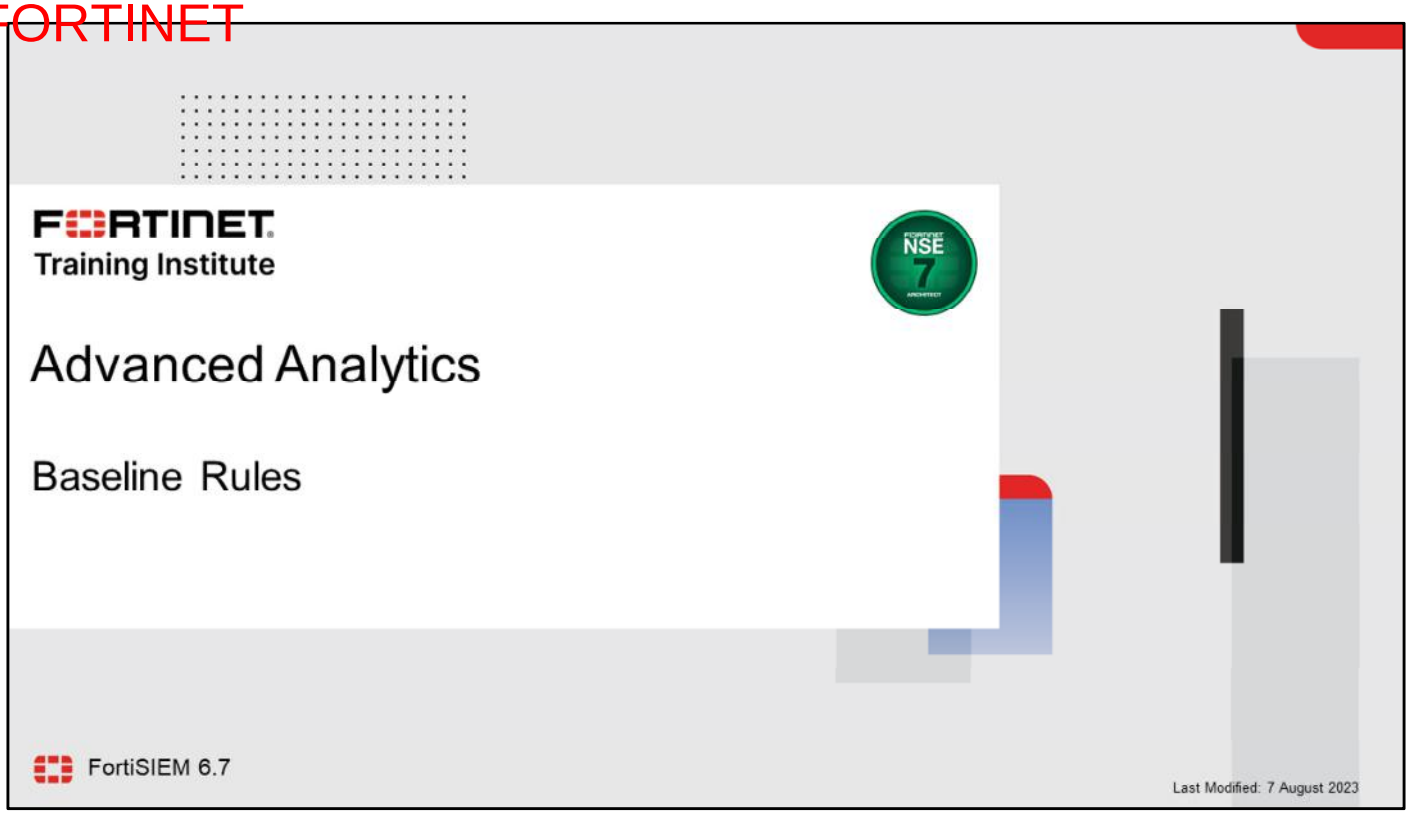
Finally, restart `phReportWorker` and `phReportMaster`, and define any new attributes that are used for min, max, average, and standard deviation in SQLite.

DO NOT REPRINT
© FORTINET



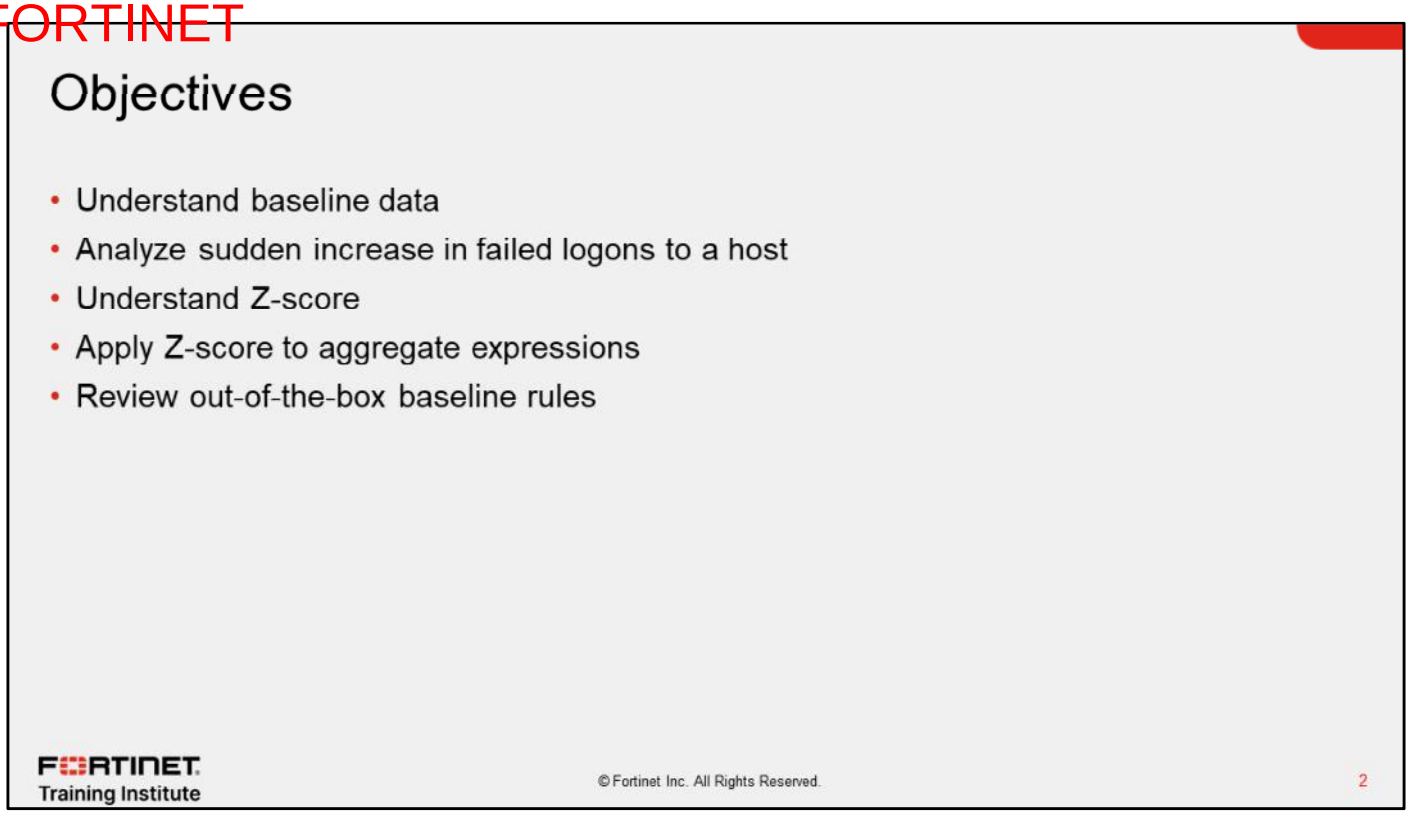
By mastering the objectives covered in this lesson, you learned how to understand baselines on FortiSIEM and analyze the back-end profile report that is responsible for tracking baseline data.

DO NOT REPRINT
© FORTINET



In this lesson, you will learn about baseline rules and how they can evaluate aggregate functions by querying data from the baseline profile.

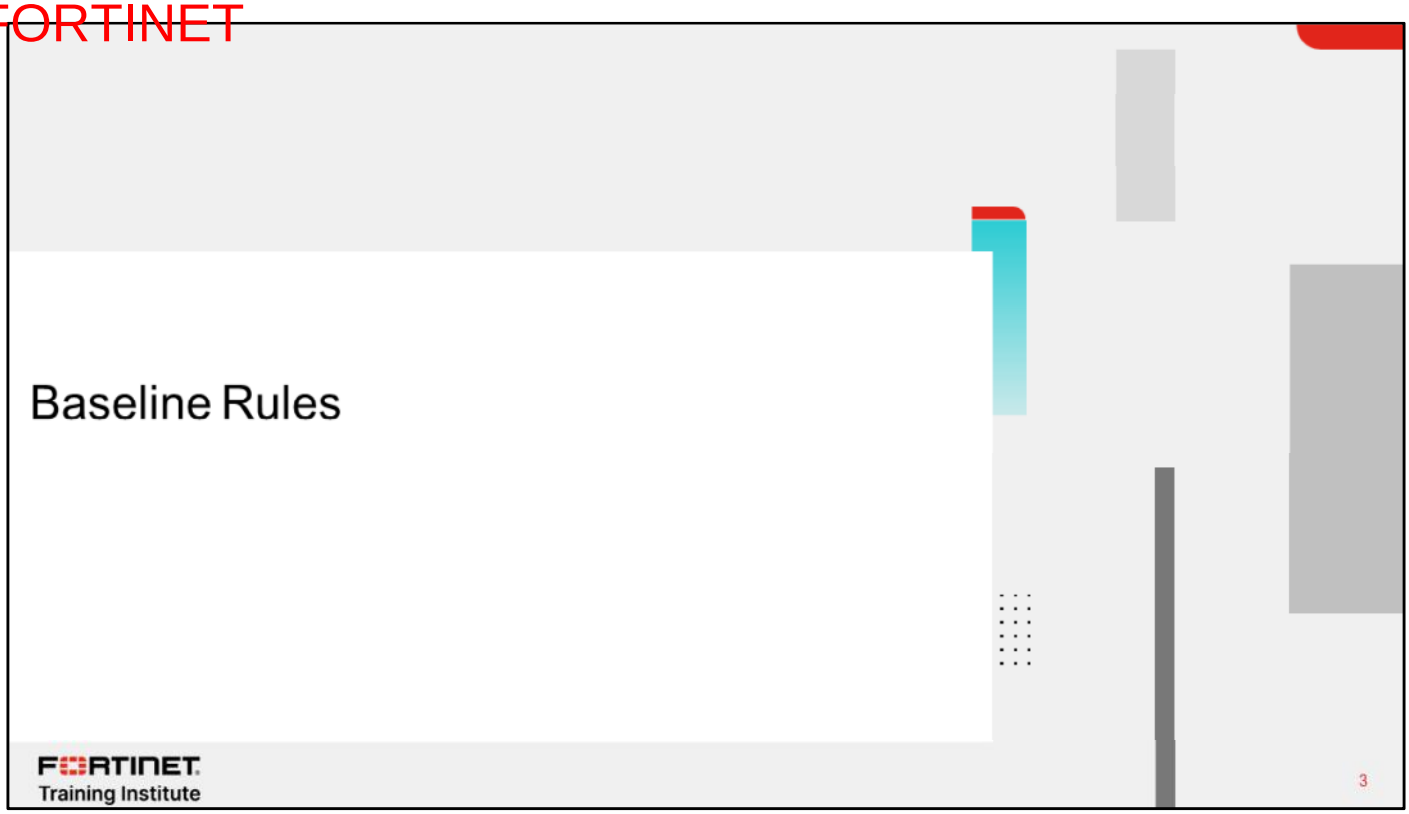
DO NOT REPRINT
© FORTINET



After completing this lesson, you should be able to achieve the objectives shown on this slide.

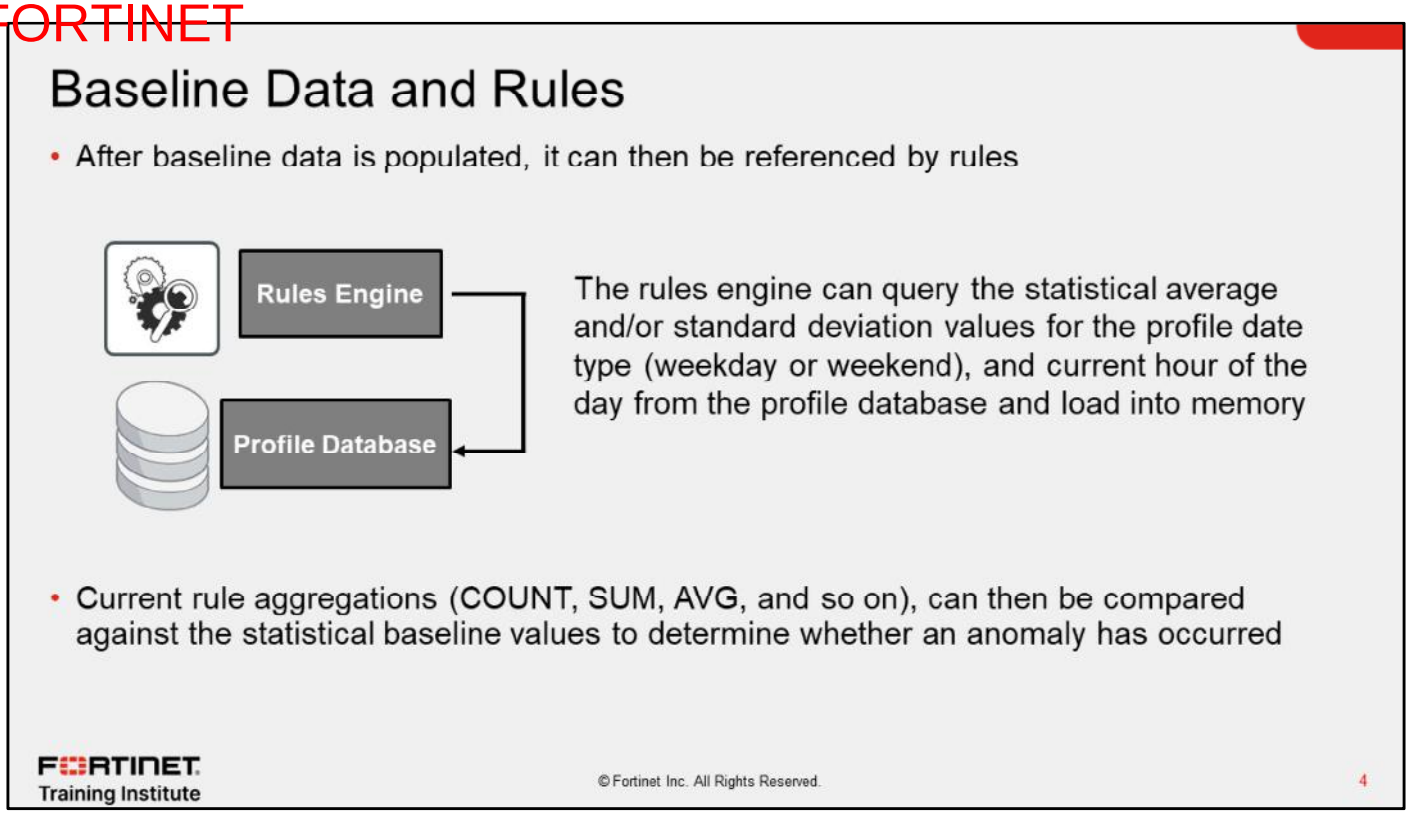
By demonstrating competence in understanding baseline data and Z-score, you will be able to analyze out-of-the-box baseline rules, and create your own baseline rules with statistical expressions.

DO NOT REPRINT
© FORTINET



In this section, you will learn about baseline data and the type of data that is fetched from the profile database by the rule engine.

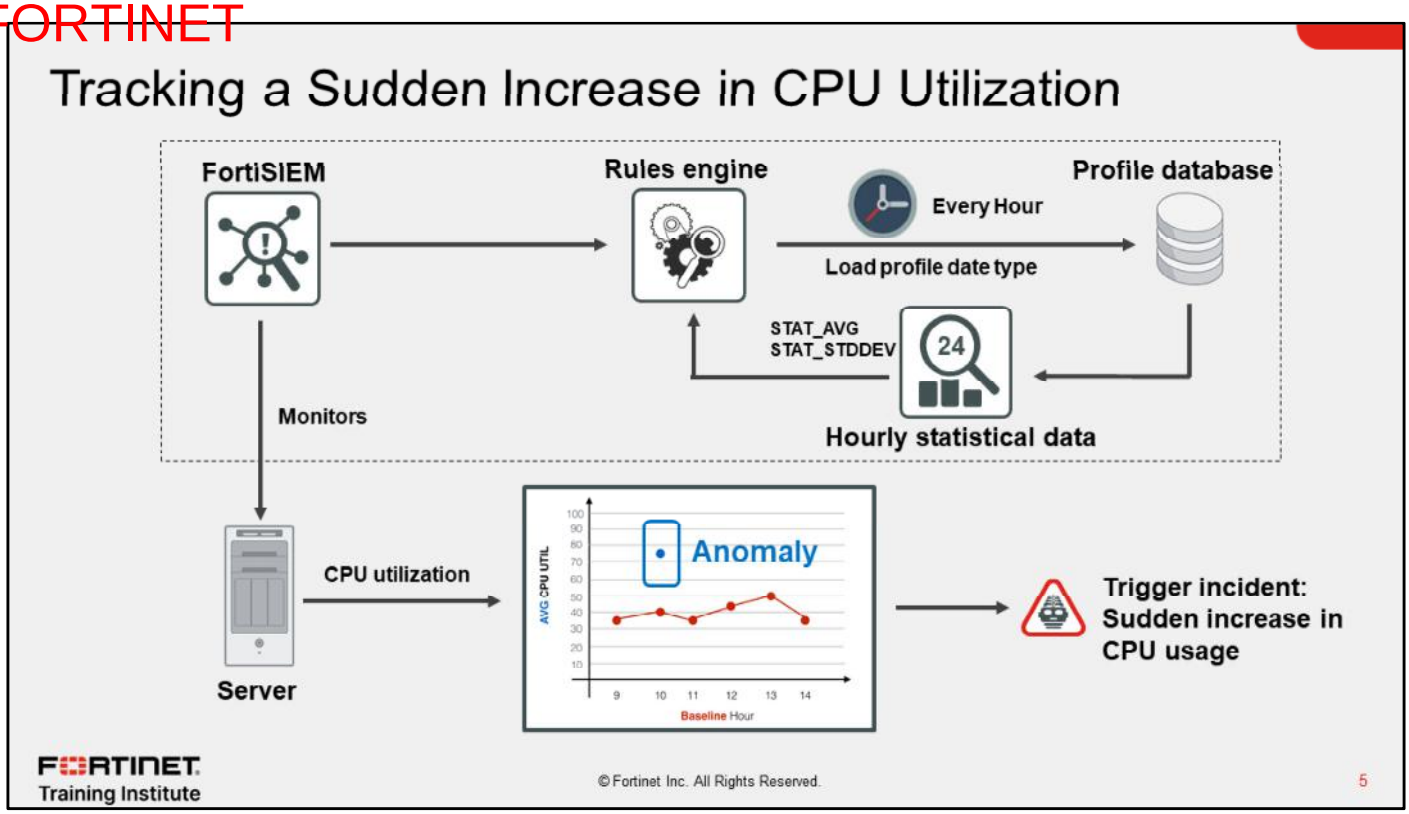
DO NOT REPRINT
© FORTINET



Once baseline data is populated in the profile database, the rule engine can query the statistical average or standard deviation values from any profile in the database. The rule engine can query the current hour of the day from a weekday or from a weekend and load that value into memory. After it loads, the statistical value into memory, the engine compares that value against the current rule aggregations, such as count, sum, average, and so on. Based on that comparison, the rule may or may not trigger an incident.

The value, once cached, stays in the memory for that hour. If the rule needs to process more events during that hour, then it will use the value that is already cached. If, for some reason, the cached value is lost, then it retrieves them from the SQLite database and caches them for future use during the hour.

DO NOT REPRINT
© FORTINET



This slide shows an example of a sudden increase in CPU utilization on a server.

The rule engine will query the profile database every hour for the statistical average and standard deviation values for that hour. The values from the profile database are cached by the rule engine. FortiSIEM then keeps monitoring the server, and periodically compares the actual CPU utilization against the statistical average and statistical standard deviation value that was cached in the memory from the profile database. If the calculations return a value that indicates an anomaly, then an incident is triggered.

DO NOT REPRINT
© FORTINET

Out-of-the-Box System Resource Baseline Rules

Rule	Profile ID	Description
Sudden Increase In System CPU Usage	109	Detects sudden increase in CPU usage. The current average over a 30-minute interval is more than 3 standard deviations away from the mean and the current average is at least 50 percent.
Sudden Increase in System Memory Usage	109	Detects a sudden increase in system memory usage. Over a 30-minute interval. Either the physical or virtual memory is 25% more than the statistical average over that same time period and the current physical memory usage is at least 50%.
Sudden Increase in Ping Response Times	127	Detects a sudden 100% increase of Ping Response Times over a 30-minute time window.
Sudden Increase in Disk I/O	121	Detects anomaly in disk I/O usage for servers, VMs, and ESX hosts. Over a 30-minute interval. Either the read or write volume is more than 3 standard deviations away from the mean over that same time period and the read/write volume is at least 1 Mbps.
Sudden Increase in Server Process Count	117	Detects that a server is suddenly running 25% more processes than the average.

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

6

This slide shows a few examples of the out-of-the-box system resource baseline rules.

The **Sudden Increase in System CPU Usage** rule queries profile ID 109. It detects the increase in CPU usage by calculating the current average over a 30-minute interval. The rule triggers when this value is more than 3 standard deviations away from the mean and the current average is at least 50 percent.

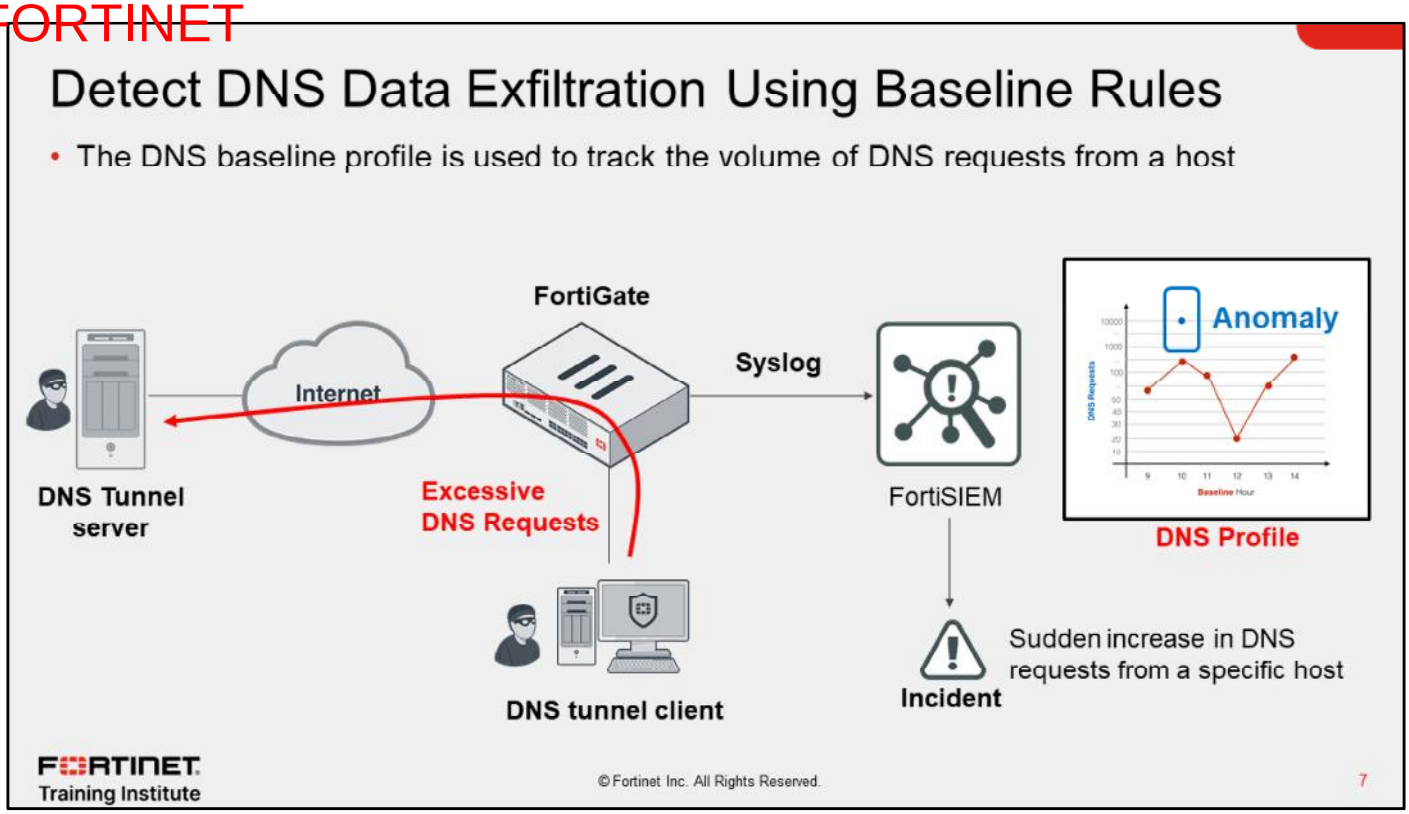
The **Sudden Increase in System Memory Usage** rule queries profile ID 109. It detects the increase in system memory usage over a 30-minute interval. This rule triggers when either the physical or virtual memory usage is 25% more than the statistical average over that same time period and the current physical memory usage is at least 50%.

The **Sudden Increase in Ping Response Times** rule queries profile ID 127. It triggers on a 100% increase of ping response times over a 30-minute period.

The **Sudden Increase in Disk I/O** rule queries profile ID 121. It detects anomalies in disk I/O usage for servers, VMs, and ESXi hosts over a 30-minute interval. The rule triggers when the read or write volume is more than 3 standard deviations away from the mean over that same time period and the read/write volume is at least 1 Mbps.

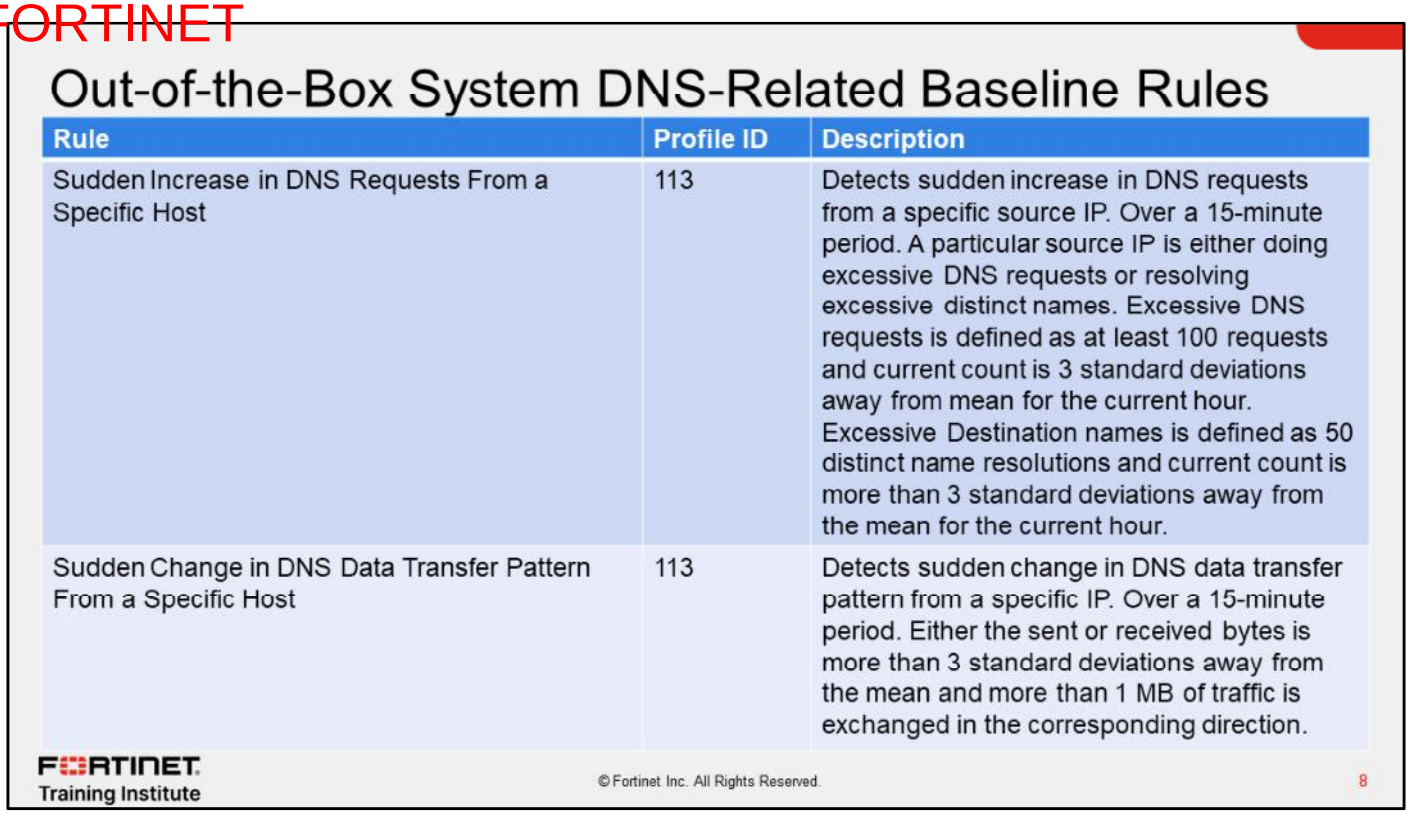
The **Sudden Increase in Server Process Count** rule queries profile ID 117. It detects a sudden increase in running processes by 25% or more than the average on a server.

DO NOT REPRINT
© FORTINET



The DNS profile can track the volume of DNS requests from a host. Excess DNS requests from a host could indicate malware running on an end host that could be tunneling data on port 53 to an attacker. In such situations, the end user is compromised and is unaware of the attack because DNS traffic looks legitimate on the firewall and is usually not blocked. However, suspicious behavior of this type can be detected by FortiSIEM by comparing the statistical DNS traffic from the end user and the current DNS traffic. If there is an anomaly, then an incident will be triggered.

DO NOT REPRINT
© FORTINET

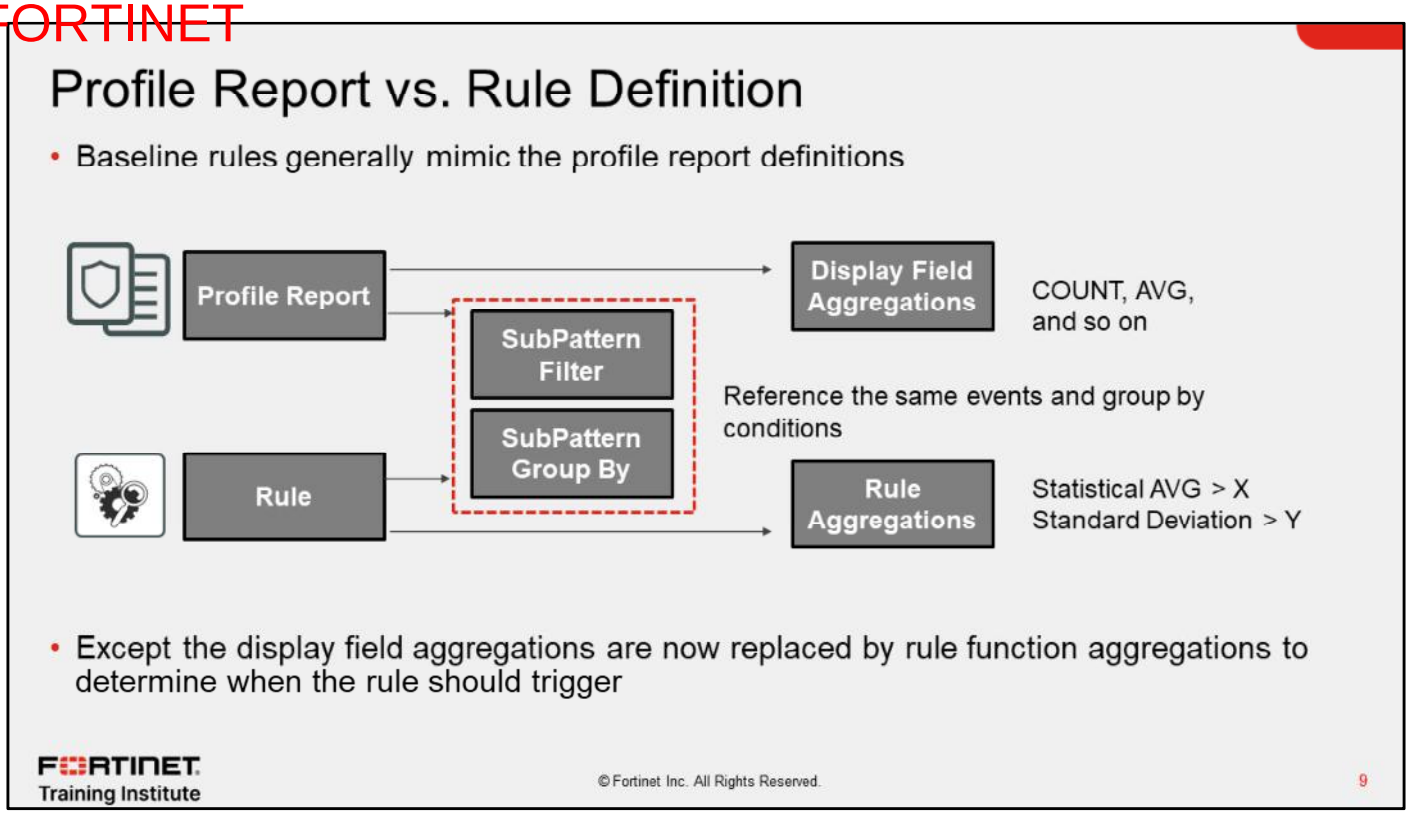


There are two baseline rules that are built into the system for detecting DNS anomaly traffic.

The Sudden Increase in DNS Requests From a Specific Host rule is detected over a 15-minute period. This rule will trigger if a particular source IP is either generating excessive DNS requests or resolving excessive distinct names. Excessive DNS requests is defined as at least 100 requests and the current count is 3 standard deviations away from mean for the current hour. Excessive destination names is defined as 50 distinct name resolutions and the current count is more than 3 standard deviations away from the mean for the current hour. This rule queries the profile ID 113.

The Sudden Change in DNS Data Transfer Pattern From a Specific Host rule is detected over a 15-minute time period. This rule will trigger when either the sent or received bytes is more than 3 standard deviations away from the mean and more than 1 MB of traffic is exchanged in the corresponding direction. This rule queries the profile ID 113.

DO NOT REPRINT
© FORTINET



Baseline rules generally mimic the profile report definition. The subpattern filter and group by definitions for the baseline rule and the profile report are the same. The baseline rule has its own definition of aggregation function where it compares the current value of an attribute against the statistical value from the profile database. However, the aggregation in the profile report is related to the display fields.

DO NOT REPRINT
© FORTINET

Profile Database Tables

- The rules engine is using the weighted AVG and SDEV values specifically from each profile table when the statistical rule functions are referenced

```
sqlite> select * from profile_122;
custId|profDateType|hourOfDay|reptDevName|reptDevIpAddr|minMatchedEv
ents|maxMatchedEvents|avgMatchedEvents|sdevMatchedEvents|minDistinct
User|maxDistinctUser|avgDistinctUser|sdevDistinctUser|minDistinctSrc
IpAddr|maxDistinctSrcIpAddr|avgDistinctSrcIpAddr|sdevDistinctSrcIpAd
dr|numPoints
```

- Min and max values are not used by the rules engine

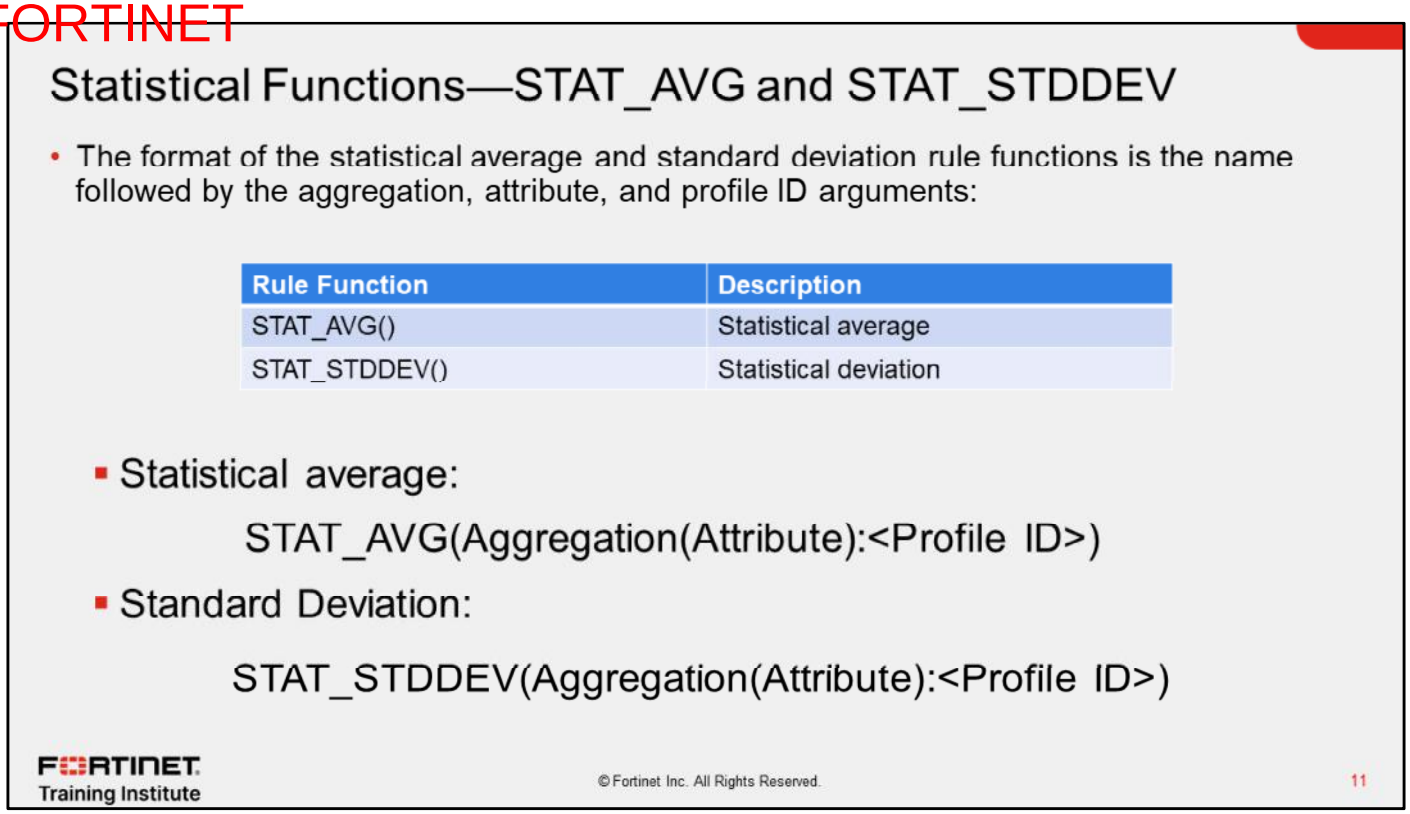
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

10

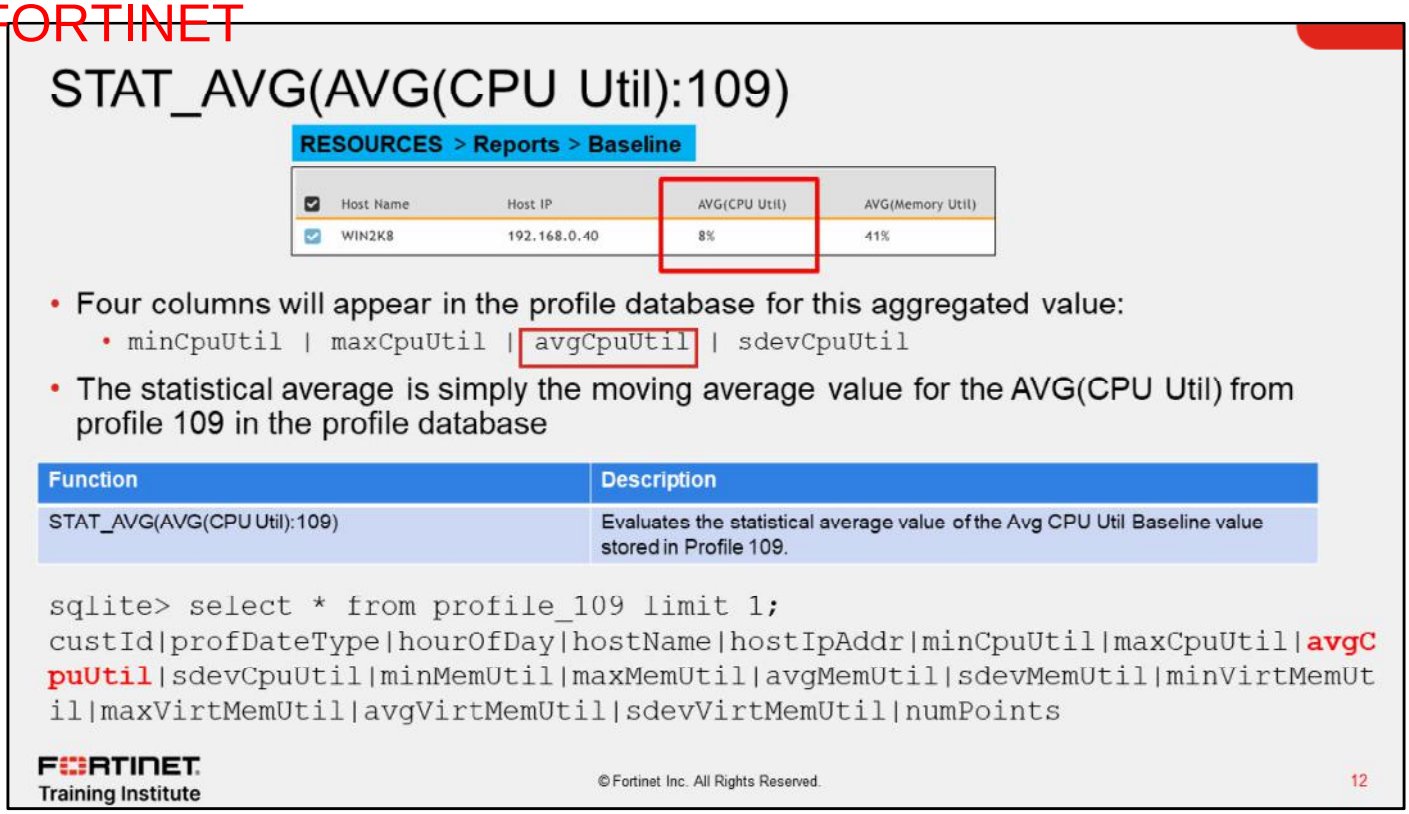
The rules engine is using the weighted average and standard deviation values, specifically from each profile table when the statistical rule functions are referenced. The `min` and `max` values are not used by the rules engine.

DO NOT REPRINT
© FORTINET



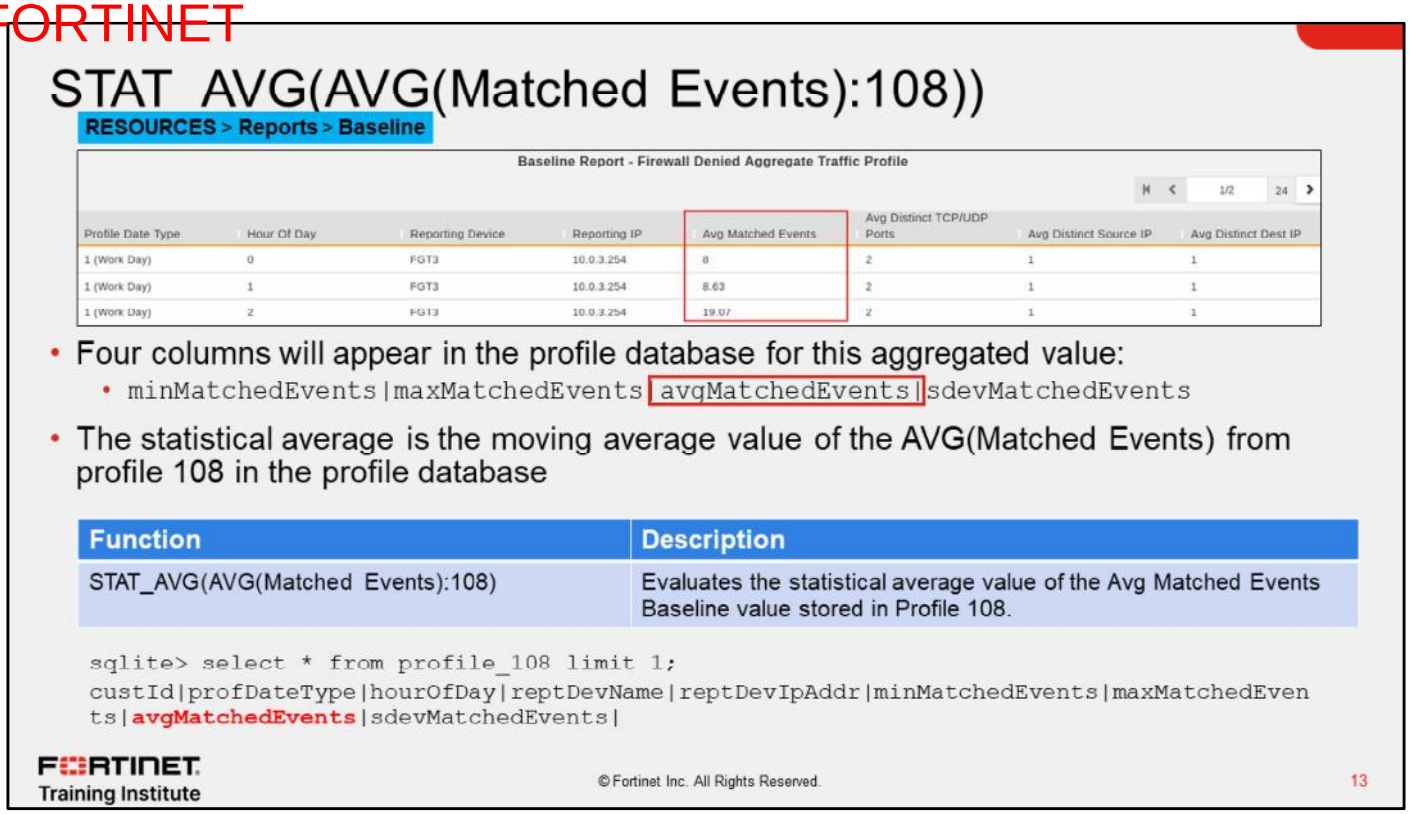
The format of the statistical average and standard deviation rule functions is the name followed by the aggregation, attribute, and profile ID arguments.

DO NOT REPRINT
© FORTINET



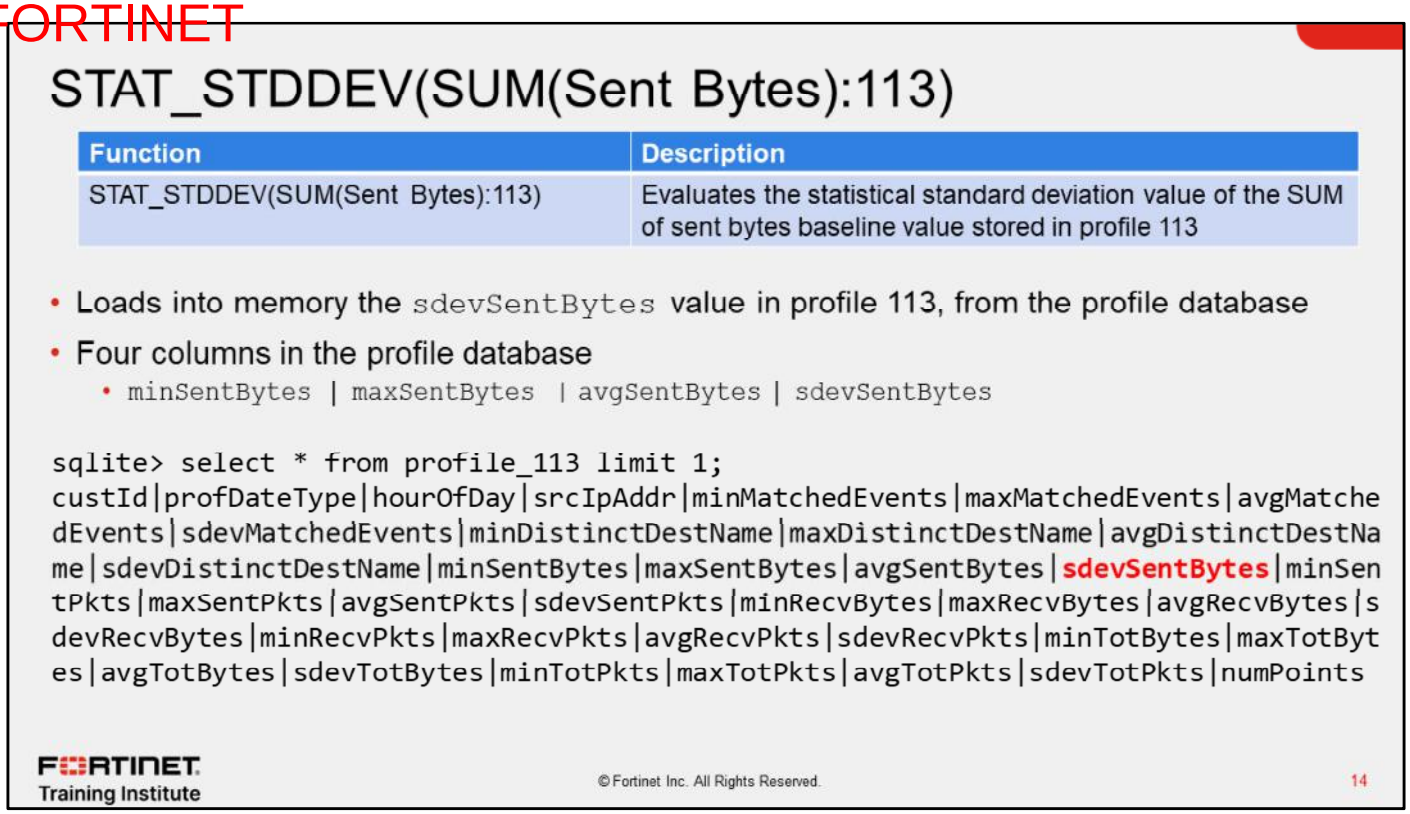
Now analyze the statistical average CPU utilization. Average CPU utilization is a very common performance attribute when you are monitoring any system in your organization. In the profile database, there will be four columns for this aggregated value, as shown on this slide. The statistical average is simply the moving average value of the average CPU utilization column from profile 109 in the profile DB.

DO NOT REPRINT
© FORTINET



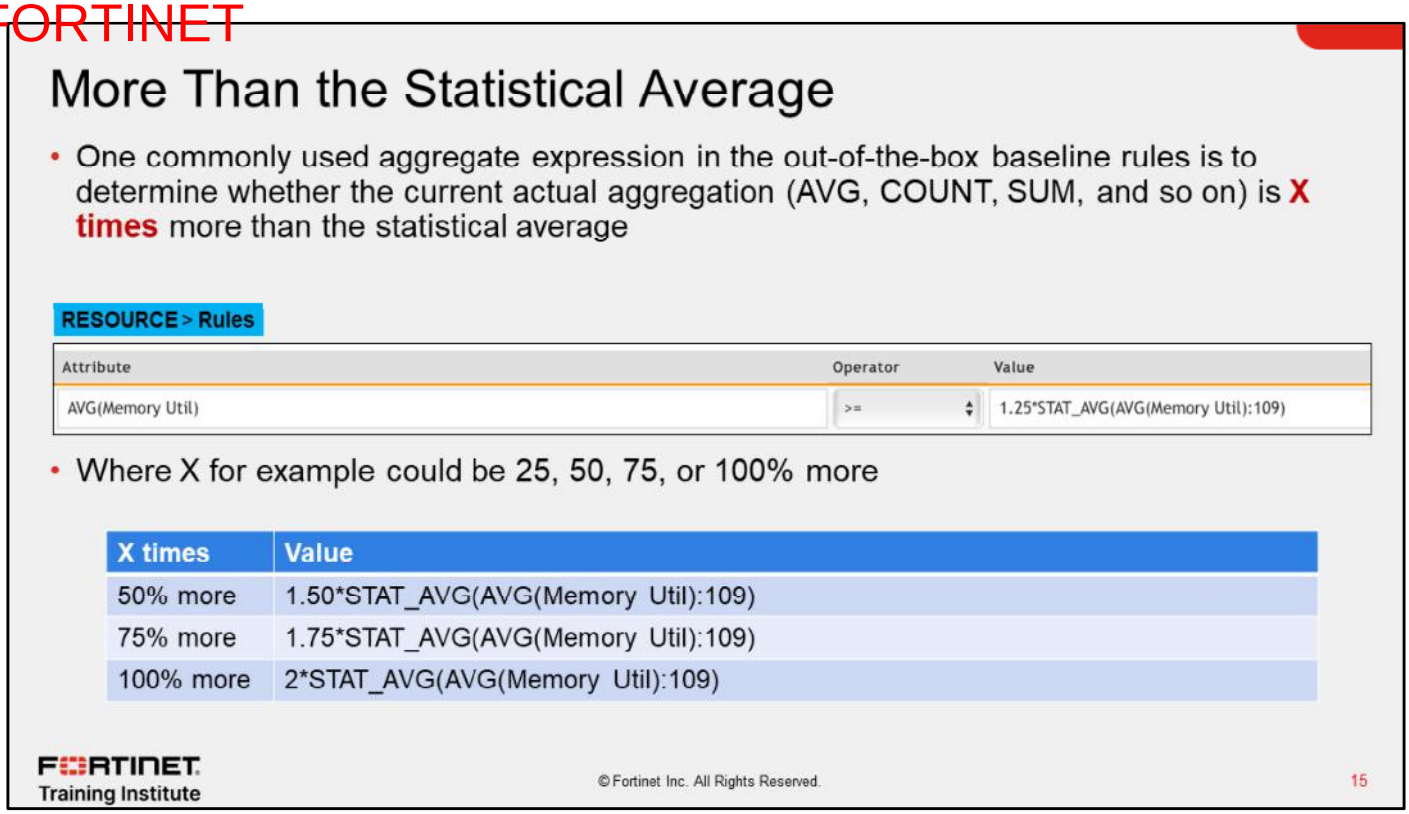
This slide shows an analysis of the statistical average of matched events for the firewall denied aggregate traffic profile. In the profile database, there will be four columns for this aggregated value as shown on this slide. The statistical average is the moving average value of the average matched events column from profile 108 in the profile database.

DO NOT REPRINT
© FORTINET



This slide shows an example of the statistical standard deviation of the sum of sent bytes for the baseline profile 113. The sent bytes will also have four columns in the profile database as shown on this slide. If a baseline rule needs to evaluate an aggregate function and requires the standard deviation value for an attribute, then that value will be loaded into the memory. This is the case for any baseline values stored in any profile.

DO NOT REPRINT
© FORTINET



One commonly used aggregate expression in the out-of-the-box baseline rules is to determine whether the current actual aggregation, such as the average, count, or sum, is X times more than the statistical average.

DO NOT REPRINT
© FORTINET

Num Points Relevance in Rules

- FortiSIEM needs to learn behavior over time, to prevent premature triggering of a rule before a baseline is set and becomes active
- The rules engine will therefore fetch only values from the profile database that have a numPoints value equal to two or more
- This value is defined in the `phoenix_config.txt` file, and it can be changed:

```
[BEGIN phRuleMaster]
min_stat_points=2 #minimum number of profiled data points before
triggering incidents
```

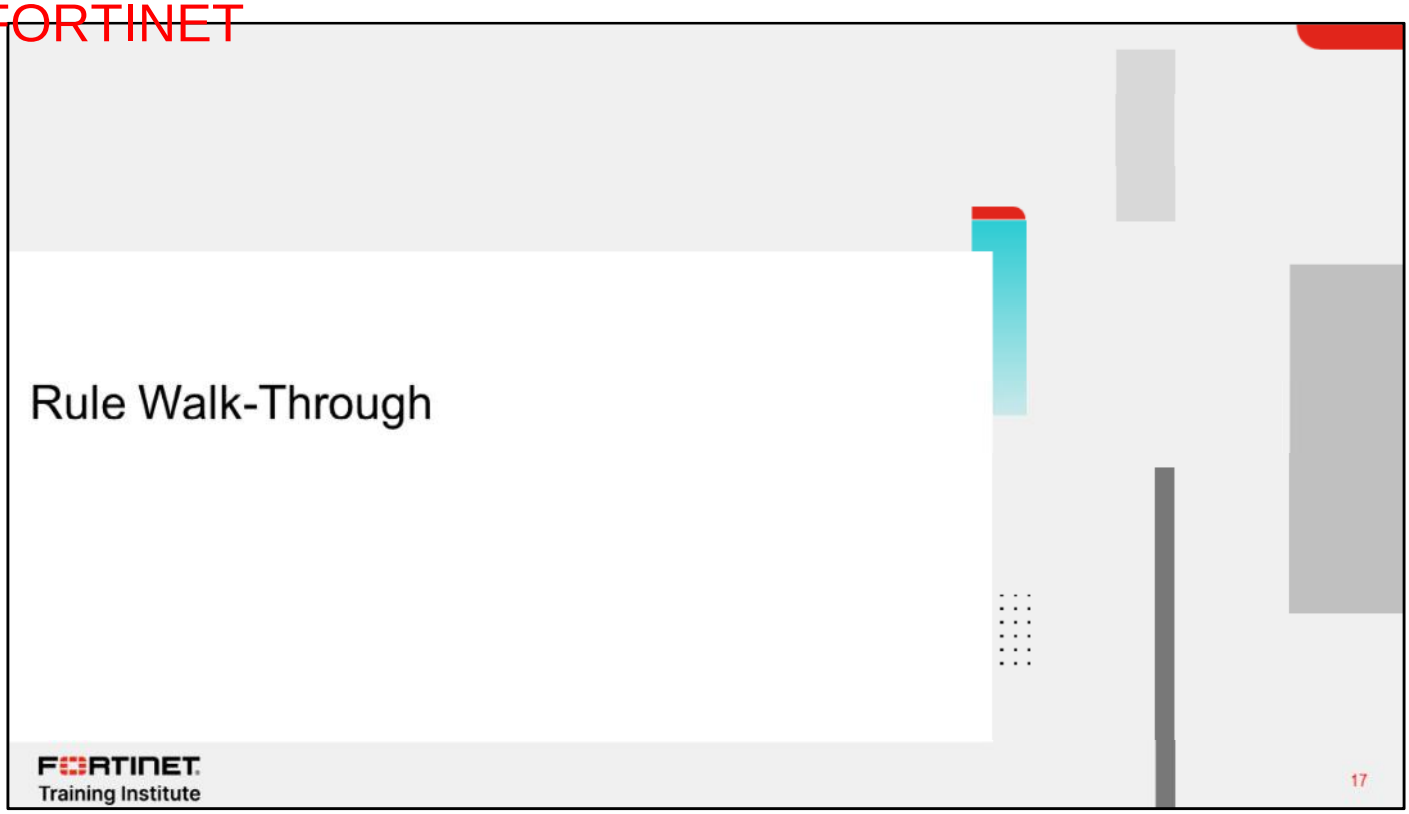
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

16

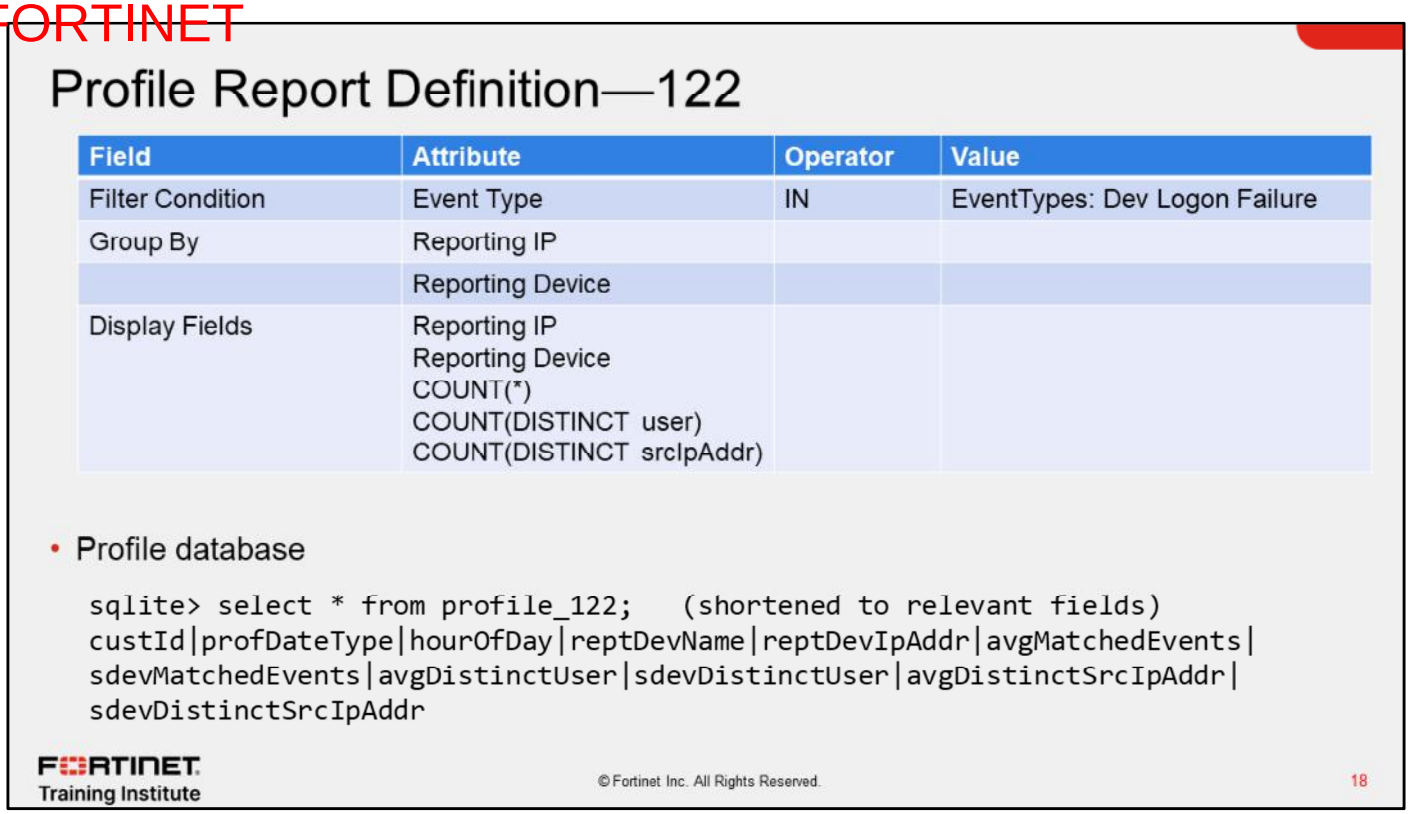
The numpoints column value in the profile database plays an important role when rules evaluate any attribute. The importance of the numpoint is to prevent premature triggering of a rule before a baseline is set and becomes active. The rules engine will therefore fetch only values from the profile database that have a numpoints value equal to two or more. This value is defined in the `phoenix_config.txt` file, and it can be changed by the administrator.

DO NOT REPRINT
© FORTINET



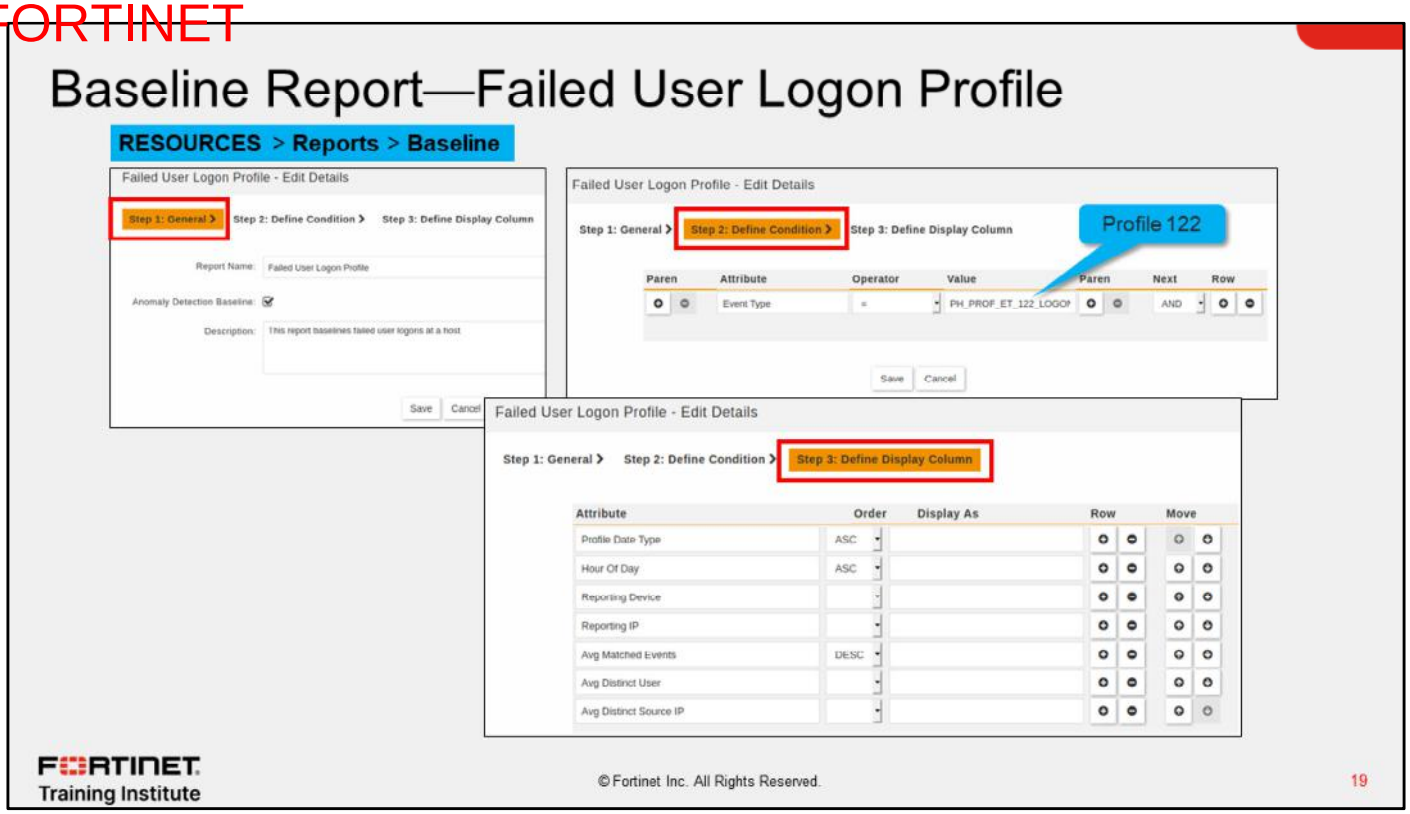
In this section, you will learn in detail about an out-of-the-box rule—**Sudden Increase in Failed Logons To A Host**. This is a baseline rule that requires baseline data to evaluate the aggregate function.

DO NOT REPRINT
© FORTINET



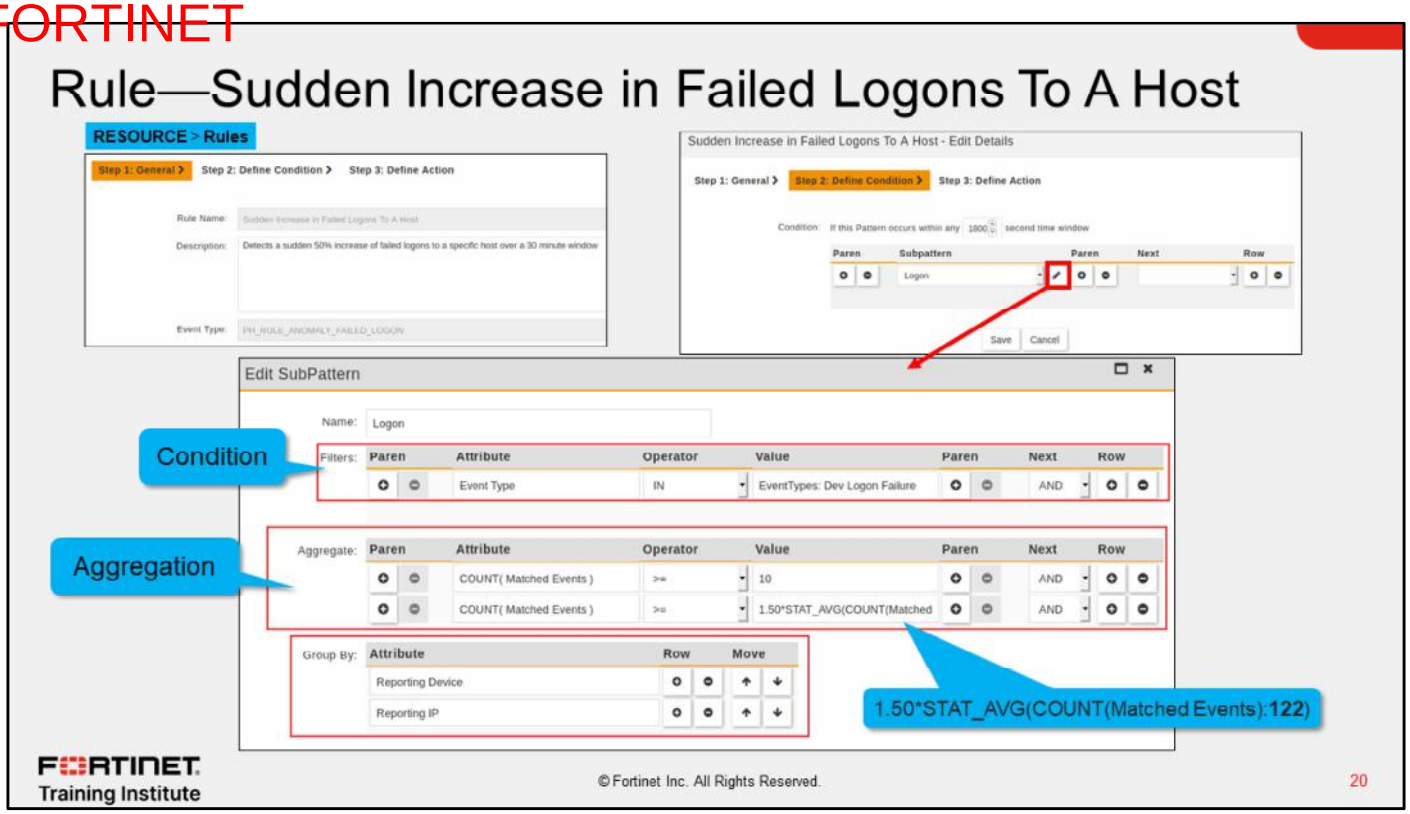
The profile report for 122 is defined in the back-end XML file. The table summarizes the filter, group by attributes, and display fields attributes that are defined in the back-end XML file.

DO NOT REPRINT
© FORTINET



The display column of the baseline report returns the selected SQLite database columns from profile 122. **Profile Date Type**, **Hour of Day**, **Reporting Device**, and **Reporting IP** are the group by attributes. **Avg Matched Events**, **Avg Distinct User**, and **Avg Distinct Source IP** are the aggregate functions.

DO NOT REPRINT
© FORTINET

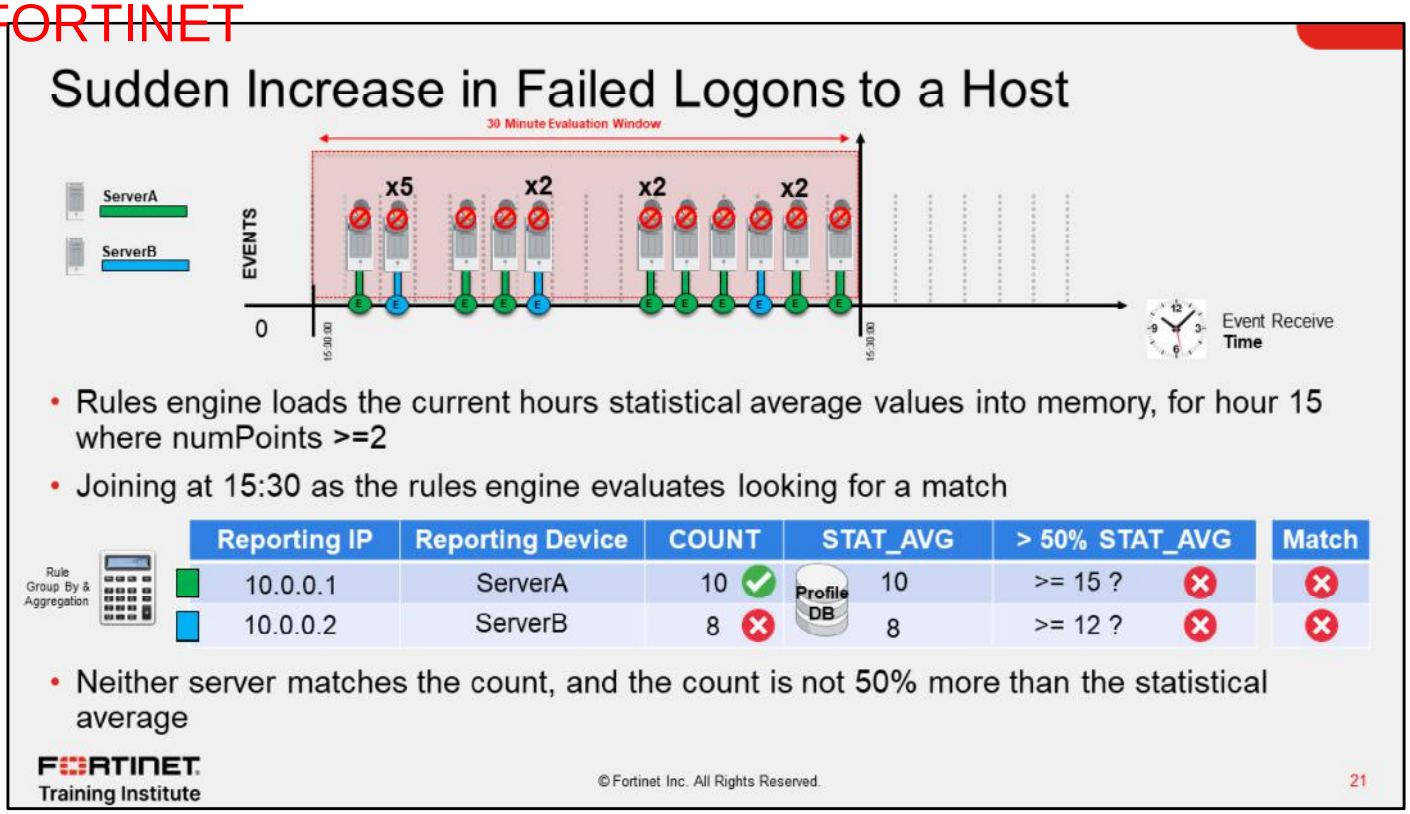


The baseline rule that is built into the system to monitor a sudden increase in failed logons to a host monitors for a sudden 50% increase of failed logons to a specific host over a 30-minute period.

The subpattern logon for the rule is monitoring for event types that belongs to the **Dev Logon Failure** group. The events of logon failure types are grouped by **Reporting Device** and **Reporting IP**.

If the matched events count is greater than or equal to 10, and the matched events count is greater than or equal to 1.5 times the statistical average of profile 122, then an incident will be triggered.

DO NOT REPRINT
© FORTINET

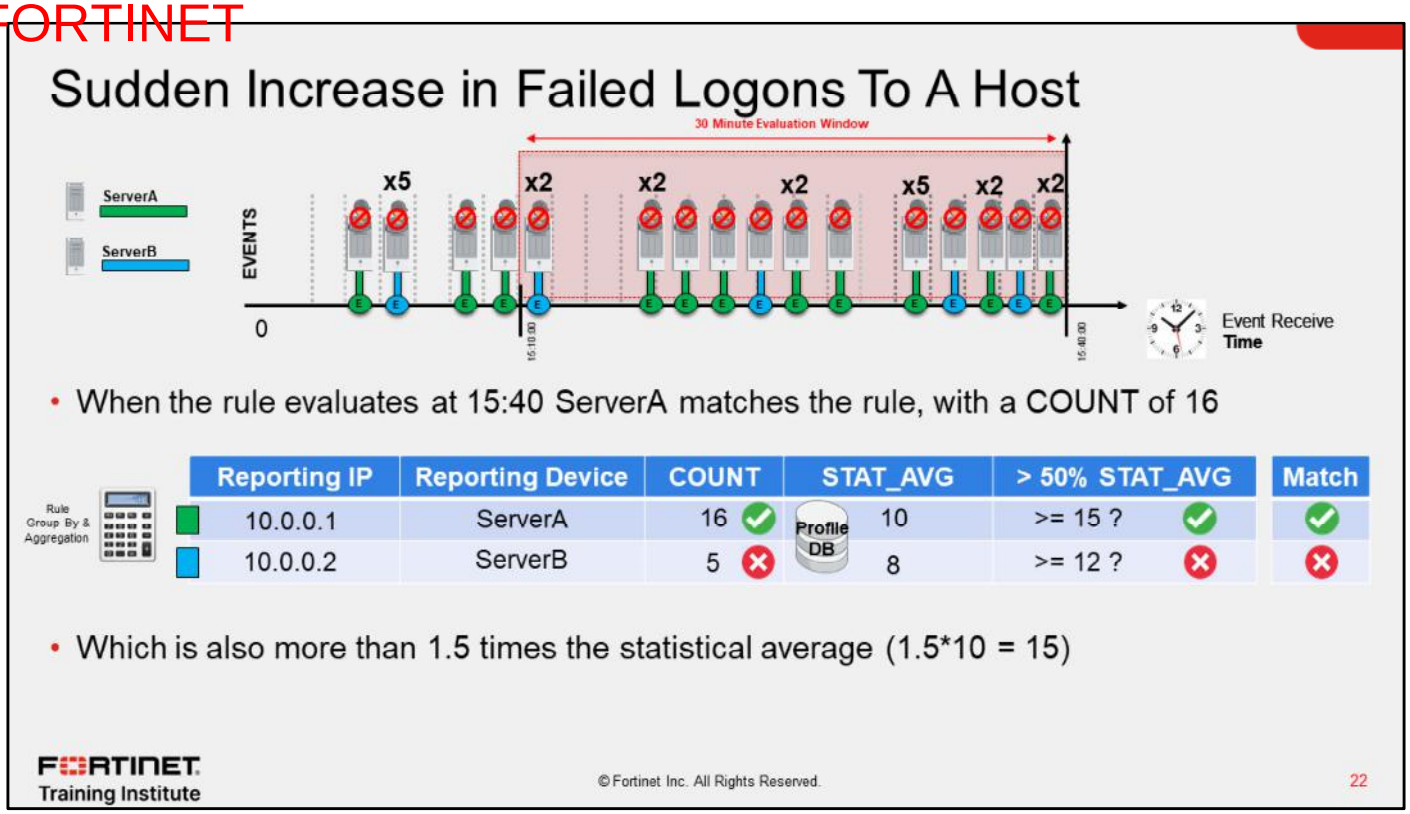


Take a look at the baseline for a sudden increase in failed logons to a host.

At the start of hour 15, the rule engine will load the statistical average value for the matched events count for profile 122 from the profile DB. It will load only the statistical average value if the `numPoints` for that hour is greater than or equal to 2. In the example shown on this slide, it will load 10 for ServerA and 8 for ServerB from the profile database into the memory.

The rule engine wakes up at 15:30 and evaluates the failed logon events for the past 30 minutes. Neither server matches the count and the count is not 50% more than the statistical average.

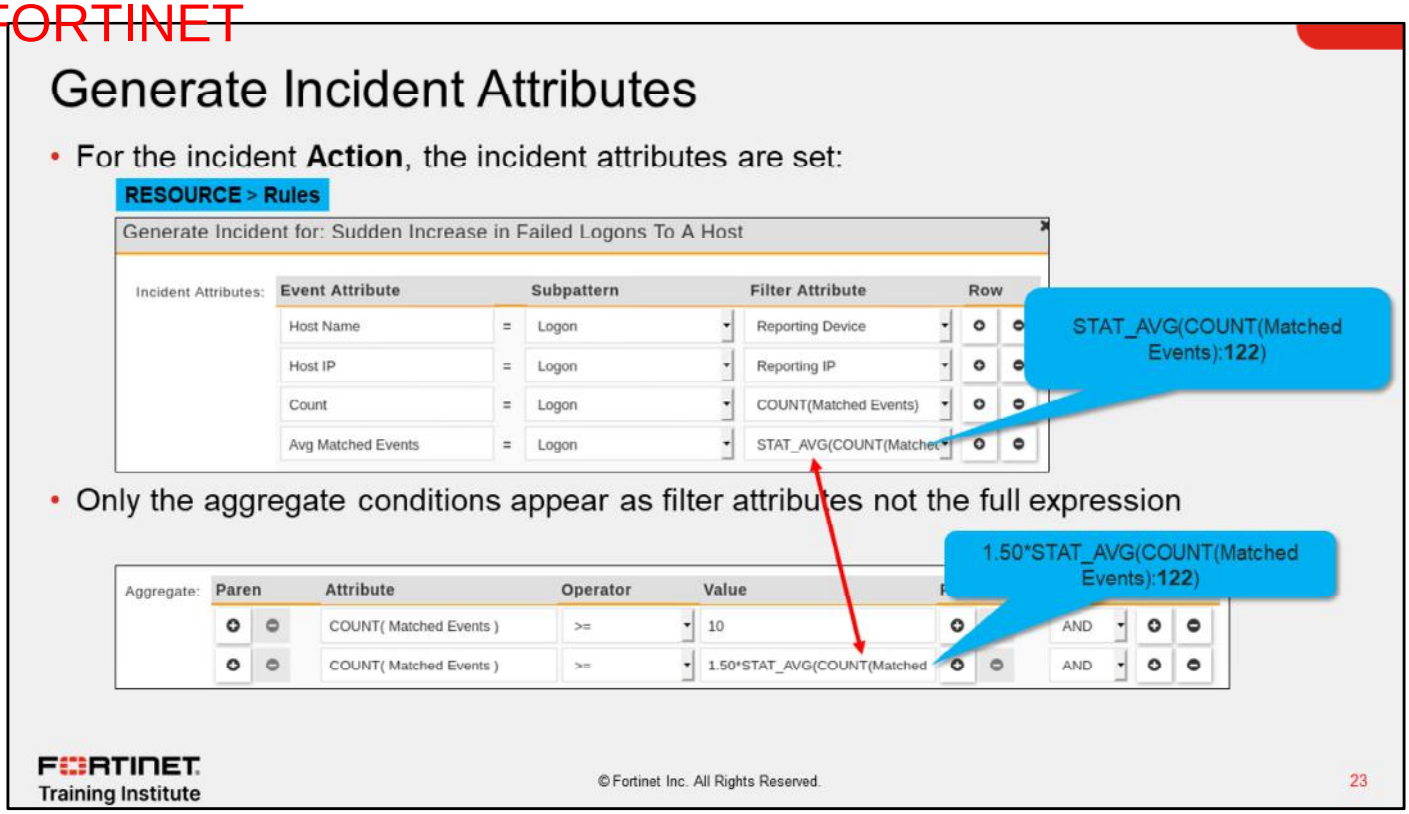
DO NOT REPRINT
© FORTINET



The rule engine again evaluates the events for a 30-minute period at 15:40. This time it does not need to load the statistical average values from the profile database to memory because they are already cached.

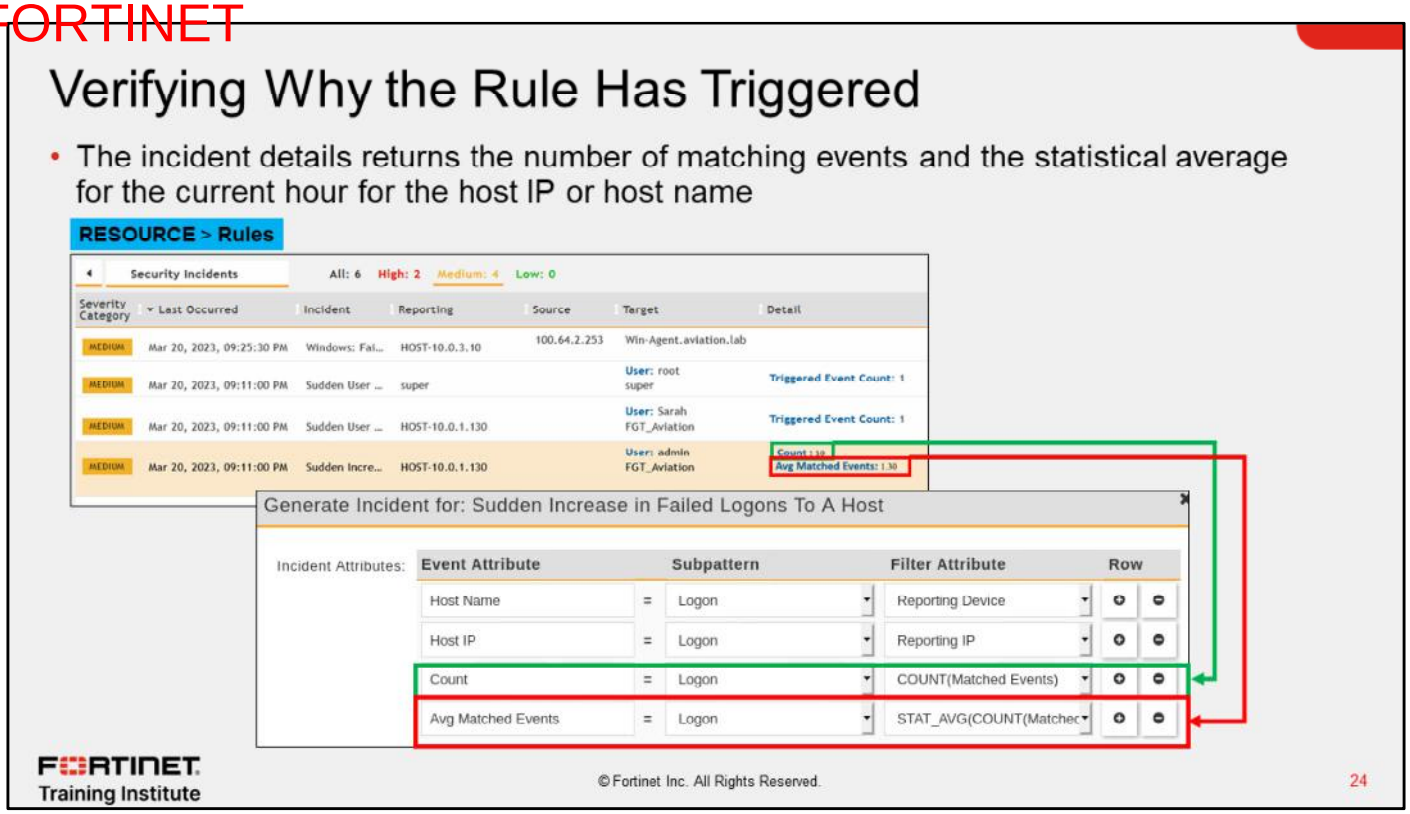
This time ServerA matches the rule, with an event count of 16 which is greater than 10 as defined in the rule, and it is also greater than 1.5 times the statistical average value that was loaded into the memory from profile 122.

DO NOT REPRINT
© FORTINET



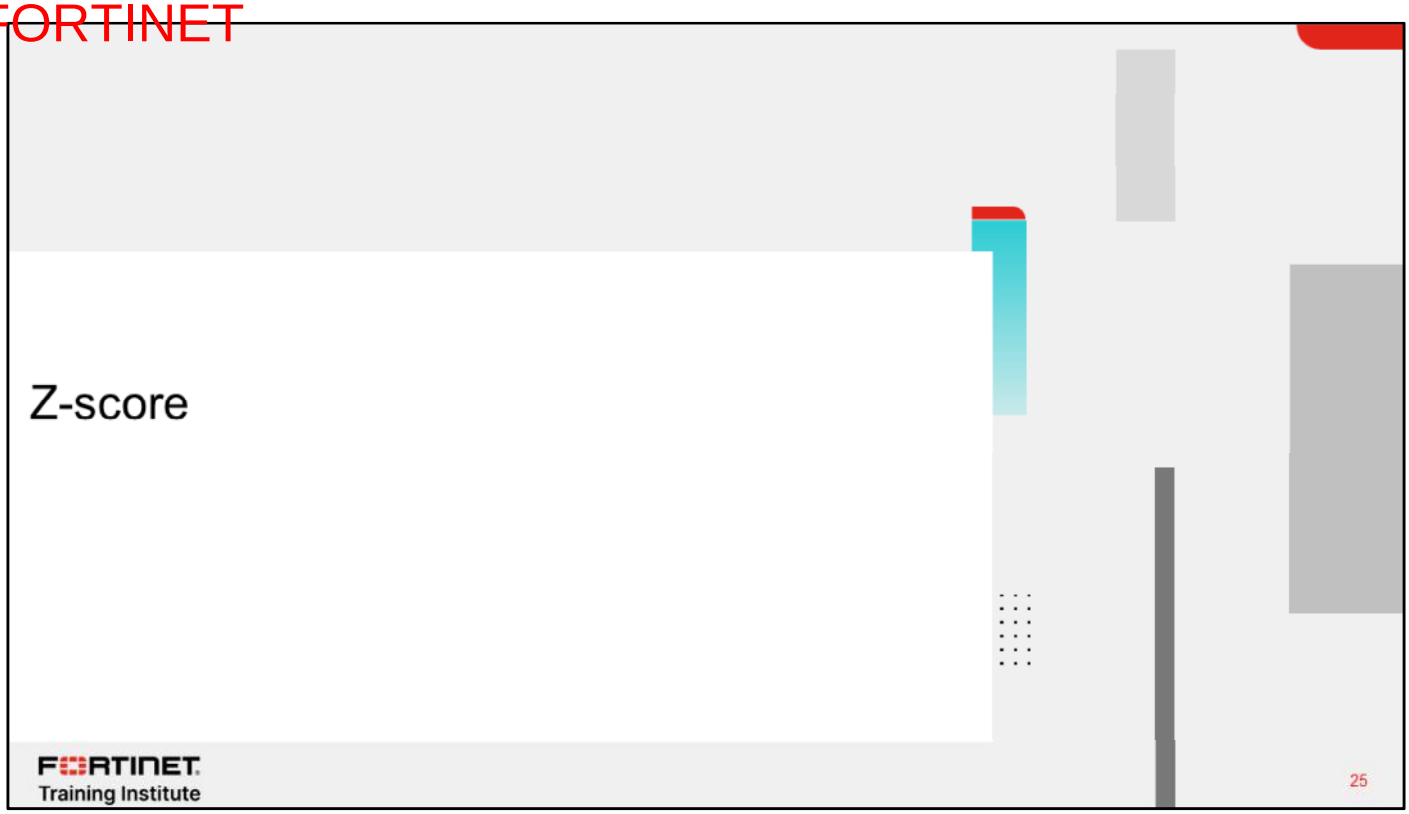
When a baseline rule triggers, sometimes you will need to review the rule definition to identify why it triggered. In the example shown on this slide, only the aggregate condition appears as the filter attribute, but not the full expression.

DO NOT REPRINT
© FORTINET



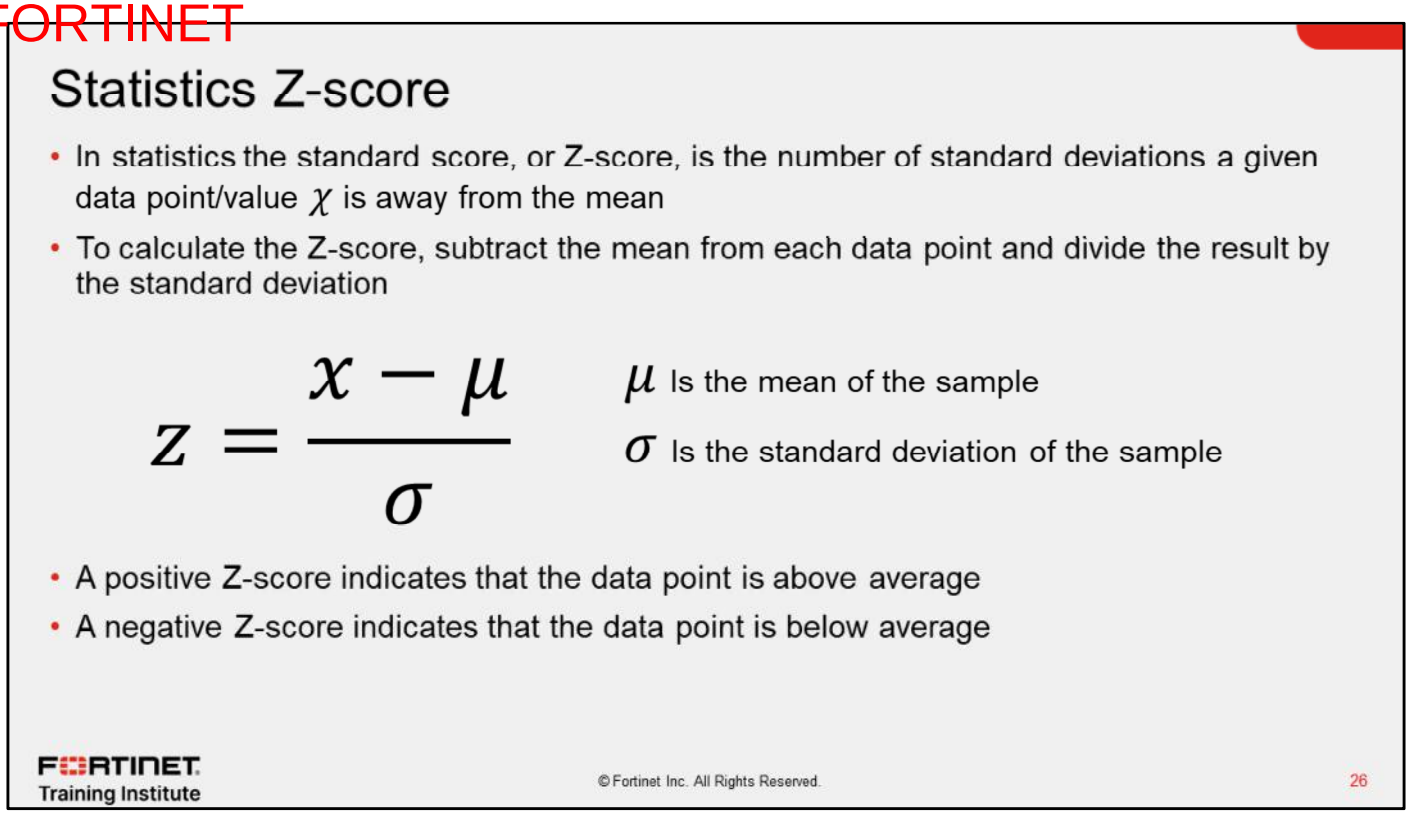
The incident details returns the number of matching events and the statistical average for the current hour for the host IP or host name that was cached from the profile database. However, the **Detail** column does not display the actual results of the aggregation expression because it is defined in the rule.

DO NOT REPRINT
© FORTINET



In this section, you will learn about the Z-score and which rules use the Z-score to evaluate the aggregate function.

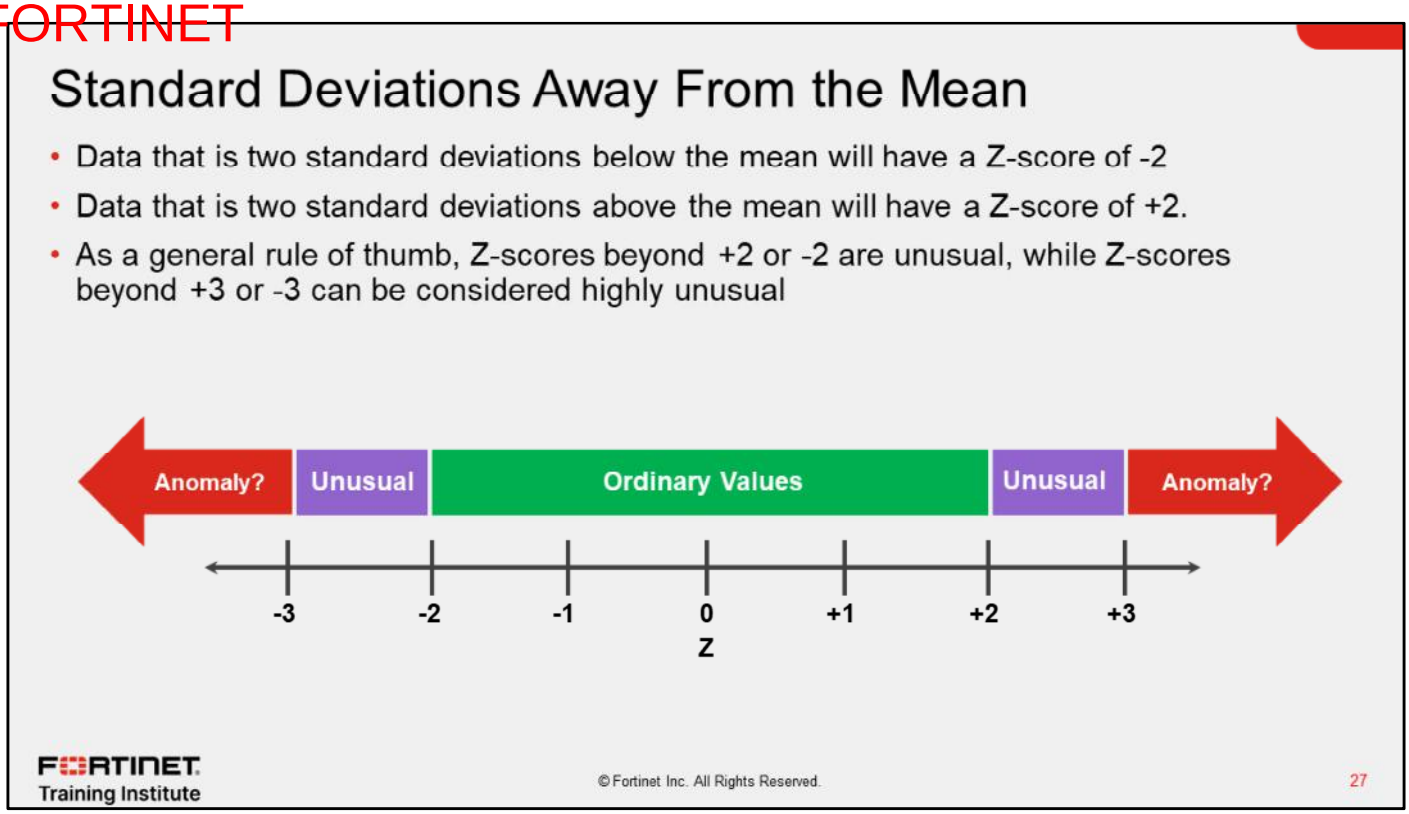
DO NOT REPRINT
© FORTINET



In statistics, the standard score or Z-score is the number of standard deviations a given data point or value is away from the mean. To calculate the Z-score, subtract the mean from each data point and divide the result by the standard deviation.

A positive Z-score indicates the data point is above average and a negative Z-score indicates the data point is below average.

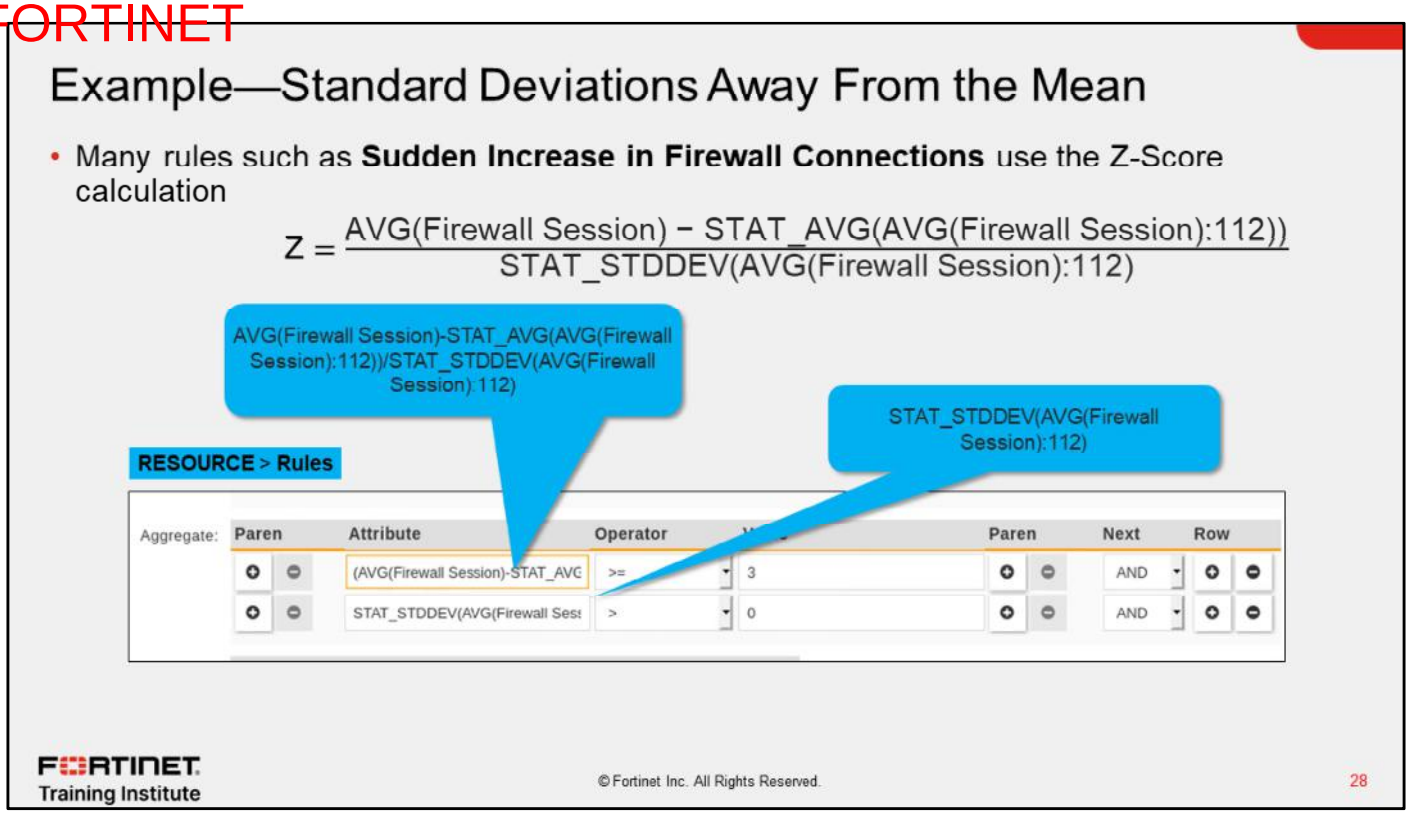
DO NOT REPRINT
© FORTINET



Data that is two standard deviations below the mean will have a Z-score of -2. Data that is two standard deviations above the mean will have a Z-score of +2.

As a general rule of thumb, Z-scores beyond +2 or -2 are unusual, while Z-scores beyond +3 or -3 can be considered highly unusual.

DO NOT REPRINT
© FORTINET



There are many rules on FortiSIEM that use the Z-score for aggregate functions. One such example is the rule that detects a sudden increase in firewall connections.

The statistical average and statistical standard deviation firewall session values are fetched from profile 112 and loaded into the memory for the hour of the day that is under evaluation. The statistical average firewall session is subtracted from the average firewall session, which is a real value, and divided by the statistical standard deviation value of the firewall session.

If the result is greater than or equal to 3 and the statistical standard deviation value of the firewall session is greater than 0, then the rule will trigger an incident.

DO NOT REPRINT
© FORTINET

Combining Aggregate Functions

- Rule: **Sudden Decrease in Reported Events From a Host**
  - Detects that a reporting device is suddenly reporting less events. The current average over the 1-hour period is less than three times the standard deviation and also 50% less than the statistical mean

Attribute	Operator	Value
(AVG(Event Rate)-STAT_AVG(AVG(Event Rate):116))/STAT_STDDEV(AVG(Event Rate):116)	<=	-3
AVG(Event Rate)	<=	0.5*STAT_AVG(AVG(Event Rate):116)
STAT_STDDEV(AVG(Event Rate):116)	>	0

RESOURCE > Rules



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

29

The **Sudden Decrease in Reported Events From a Host** rule detects that a reporting device is suddenly reporting fewer events. The current average over the 1-hour period is less than three times the standard deviation and also 50% less than the statistical mean.

DO NOT REPRINT
© FORTINET

Traffic-Related Baseline Rules

Rule	Profile ID	Description
Sudden Increase in Firewall Permitted Inbound Traffic to a Specific TCP/UDP Port	104	Detects traffic anomaly on inbound TCP/UDP port usage on a firewall. This means that over a 30-minute period, the permitted inbound traffic (both number of flows and total bytes) to a specific TCP/UDP port is more than 3 standard deviations away than the average for the current hour.
Sudden Increase In Firewall Permitted Outbound Traffic to a Specific TCP/UDP Port	105	Detects traffic anomaly on outbound TCP/UDP port usage. Detects traffic anomaly on outbound TCP/UDP port usage on a firewall. This means that over a 30-minute period, the permitted outbound traffic (both number of flows and total bytes) to a specific TCP/UDP port is more than 3 standard deviations away than the average for the current hour.
Sudden Increase In Firewall Denied Inbound Traffic to a Specific TCP/UDP Port	106	Detects anomalous denied inbound traffic profile on a specific TCP/UDP port. Over a 30-minute period, both the total number of denies and the number of unique source IP addresses are at least 3 standard deviations away from the mean for the current hour.
Sudden Increase In Firewall Denied Outbound Traffic to a Specific TCP/UDP Port	107	Detects anomalous denied outbound traffic profile on a specific TCP/UDP port. Over a 30-minute period, both the total number of denies or the number of unique destination IP addresses are at least 3 standard deviations away from the mean for the current hour.

This slide shows some of the traffic–related baseline rules.

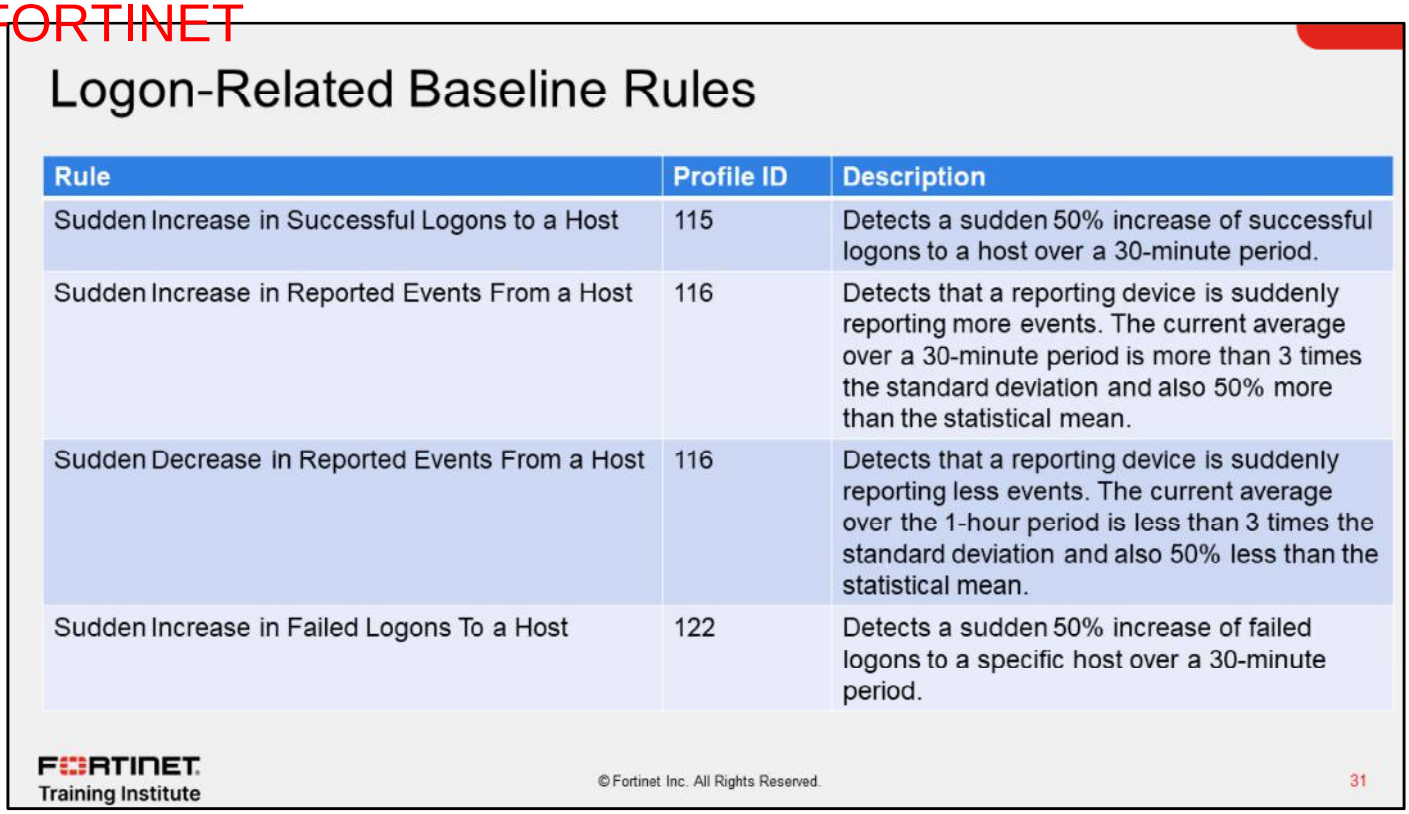
The **Sudden Increase in Firewall Permitted Inbound Traffic to a Specific TCP/UDP Port** rule detects traffic anomalies on inbound TCP/UDP port usage on a firewall. This means that over a 30-minute period, the permitted inbound traffic—both number of flows and total bytes—to a specific TCP/UDP port is more than 3 standard deviations away than the average for the current hour.

The **Sudden Increase in Firewall Permitted Outbound Traffic to a Specific TCP/UDP Port** rule detects traffic anomalies on outbound TCP/UDP port usage on a firewall. This means that over a 30-minute period, the permitted outbound traffic—both number of flows and total bytes—to a specific TCP/UDP port is more than 3 standard deviations away than the average for the current hour.

The **Sudden Increase in Firewall Denied Inbound Traffic to a Specific TCP/UDP Port** rule detects anomalous denied inbound traffic profiled on a specific TCP/UDP port over a 30-minute period. This rule will detect both the total number of denies and the number of unique source IP addresses are at least 3 standard deviations away from the mean for the current hour.

The **Sudden Increase in Firewall Denied Outbound Traffic to a Specific TCP/UDP Port** rule detects anomalous denied outbound traffic profiled on a specific TCP/UDP port over a 30-minute period. This rule will detect both the total number of denies or the number of unique destination IP addresses are at least 3 standard deviations away from the mean for the current hour.

DO NOT REPRINT
© FORTINET



This slide shows some of the logon related baseline rules.

The **Sudden Increase in Successful Logons to a Host** rule detects a 50% increase in successful logons to a host over a 30-minute period.

The **Sudden Increase in Reported Events From a Host** rule monitors for reporting devices that are suddenly reporting more events. The rule will trigger if the current average over a 30-minute period is more than 3 times the standard deviation and also 50% more than the statistical mean.

The **Sudden Decrease in Reported Events From a Host** rule monitors for reporting devices that are suddenly reporting fewer events. The rule will trigger if the current average over the 1-hour time period is less than 3 times the standard deviation and also 50% less than the statistical mean.

The **Sudden Increase in Failed Logons to a Host** rule monitors for a sudden 50% increase in failed logons to a specific host over a 30-minute period.

DO NOT REPRINT
© FORTINET

Monitoring Response Times Baseline Rules

Rule	Profile ID	Description
Sudden Increase in Ping Response Times	127	Detects a sudden 100% increase in ping response times over a 30-minute period.
Sudden Increase in SNMP Response Times	128	Detects a sudden 100% increase in SNMP response times over a 30-minute period.
Sudden Increase in WMI Response Times	129	Detects a sudden 50% increase in WMI response times over a 30-minute period.
Sudden Increase in ICMP Requests From A Host	114	Detects that a host is making excessive ICMP requests. Both the volume and the number of distinct destinations is more than twice the statistical average and at least 100 ICMP requests.

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

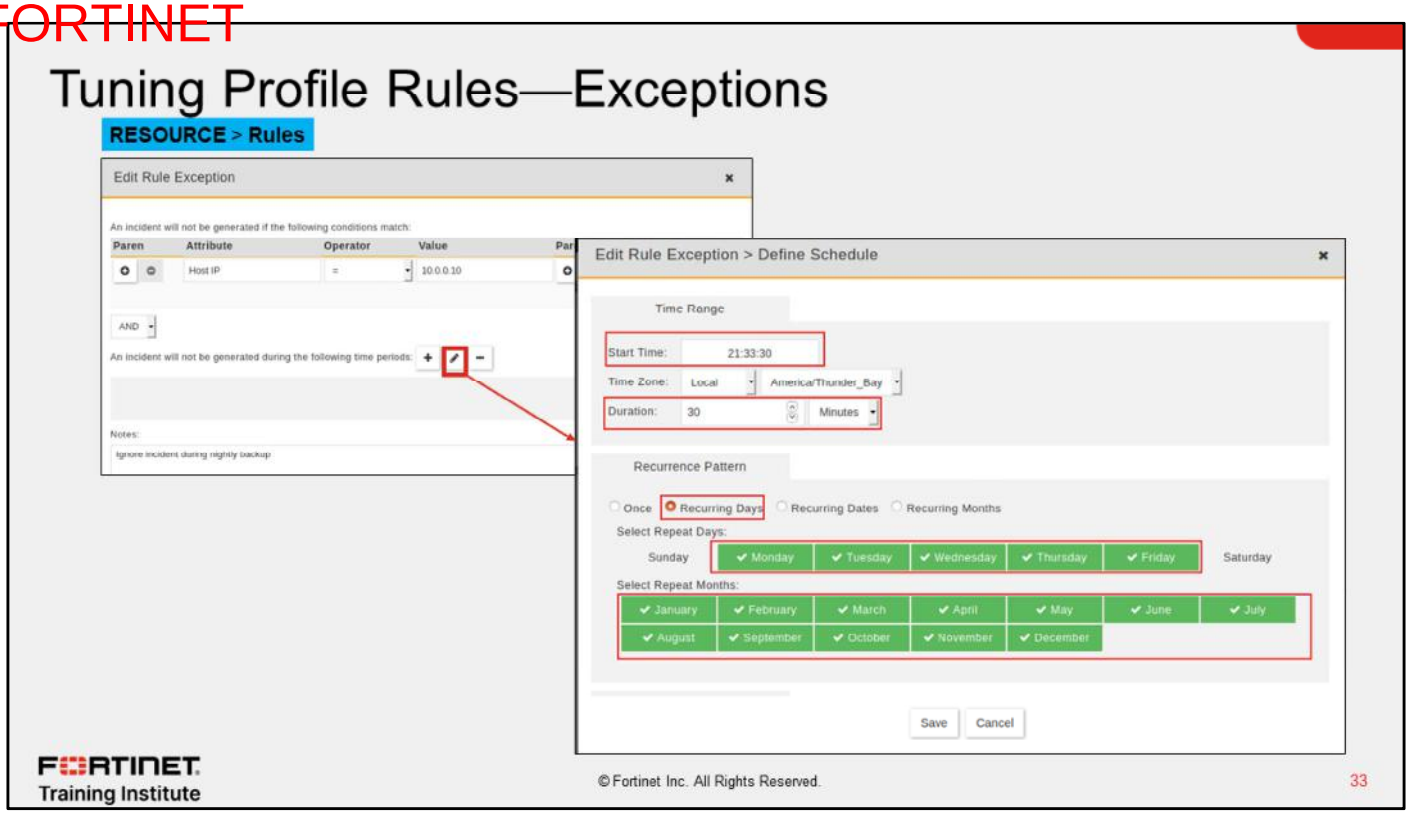
32

This slide shows some of the monitoring response times baseline rules.

The ping, SNMP, and WMI response times rules detect for a sudden increase in the response times of these protocols over a 30-minute period.

The **Sudden Increase in ICMP Requests From a Host** rule monitors for hosts making excessive ICMP requests—both the volume and the number of distinct destinations. The rule triggers when the requests are more than twice the statistical average and at least 100 ICMP requests.

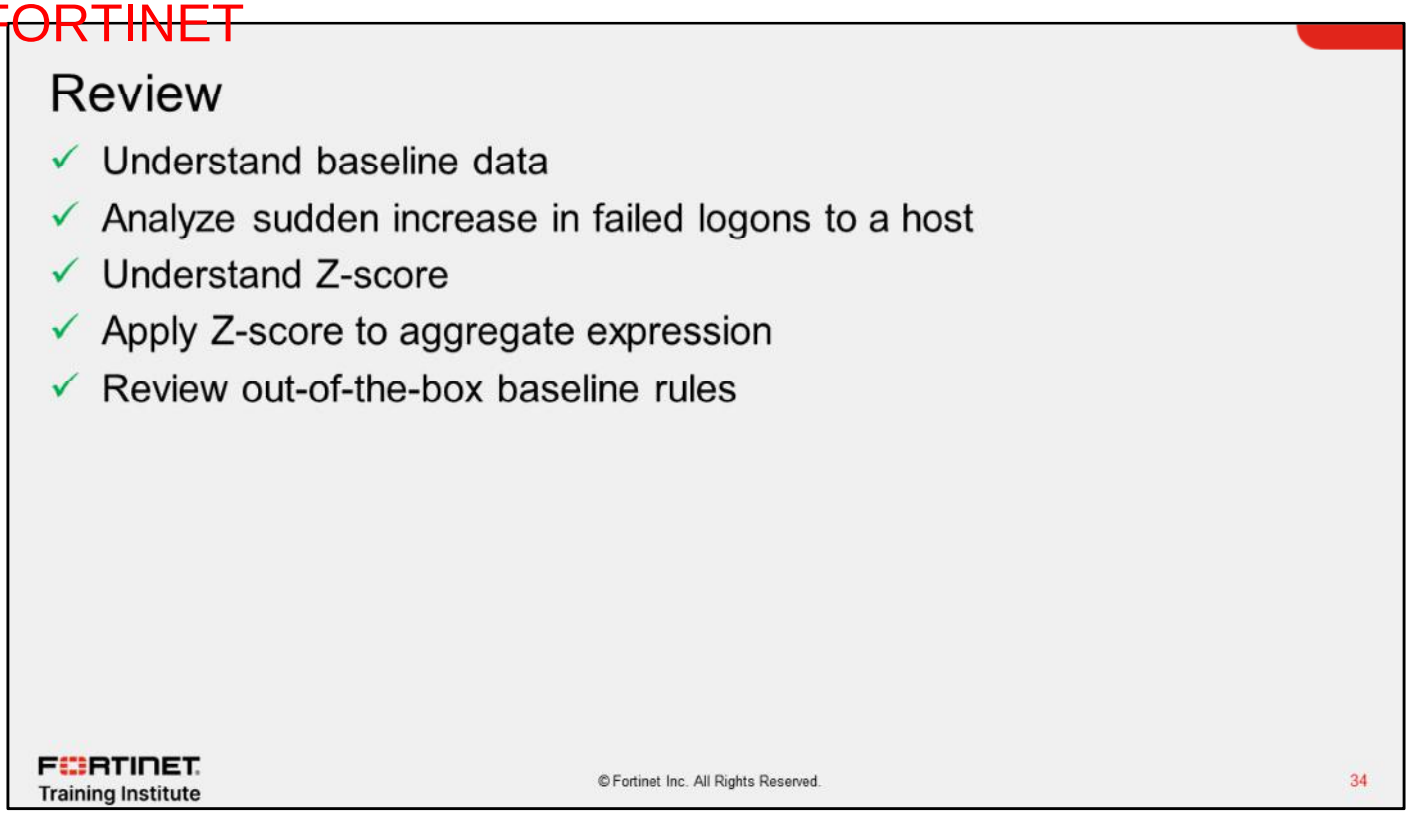
DO NOT REPRINT
© FORTINET



You can tune the rules to avoid generating incidents during certain time periods. You configure this in the rule exceptions. You can define the attributes and the values for which you do not want to trigger an incident.

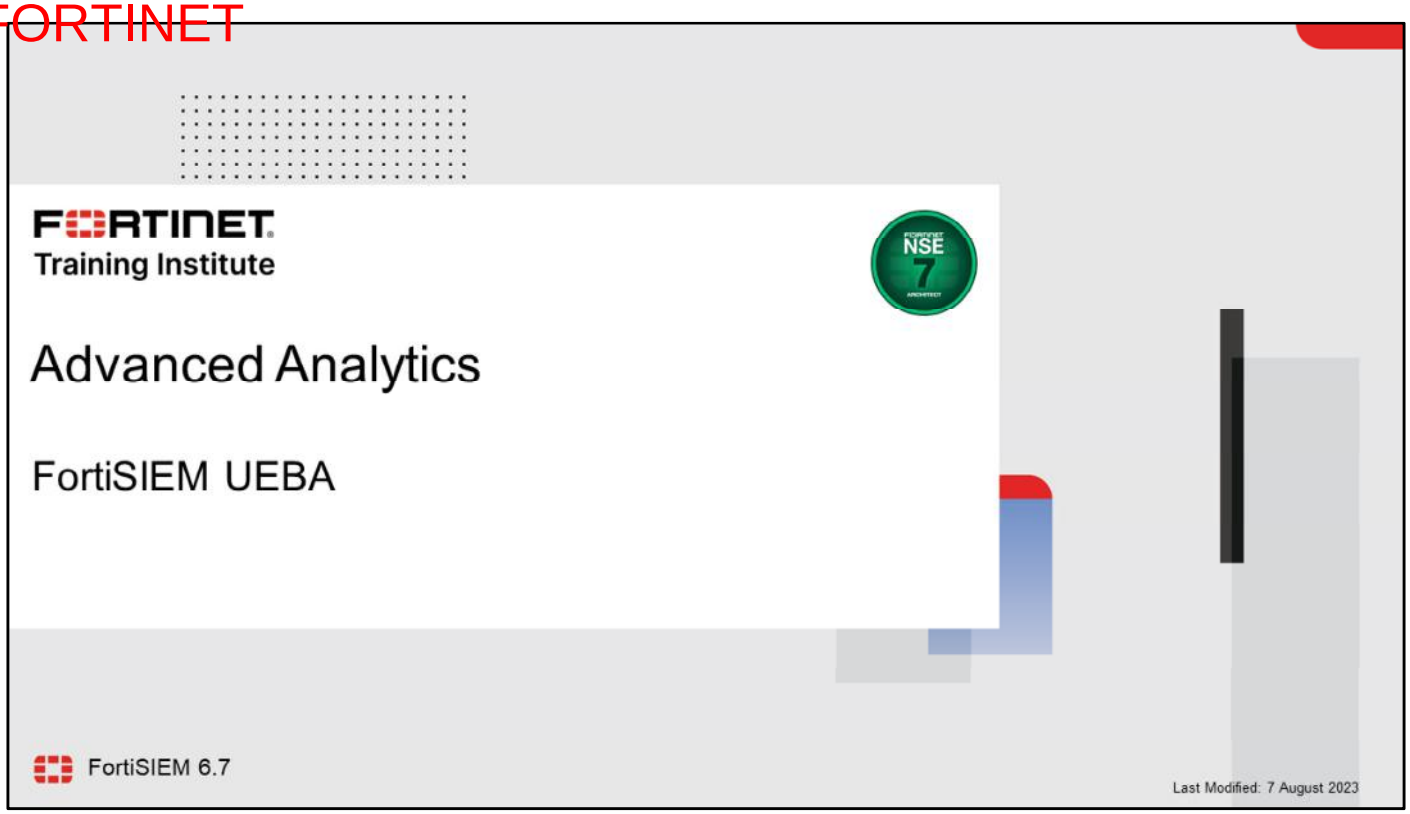
In the example shown on this slide, the incident will not be triggered for a 30-minute period at night when a backup is performed on the host `10.0.0.10` during weekdays.

DO NOT REPRINT
© FORTINET



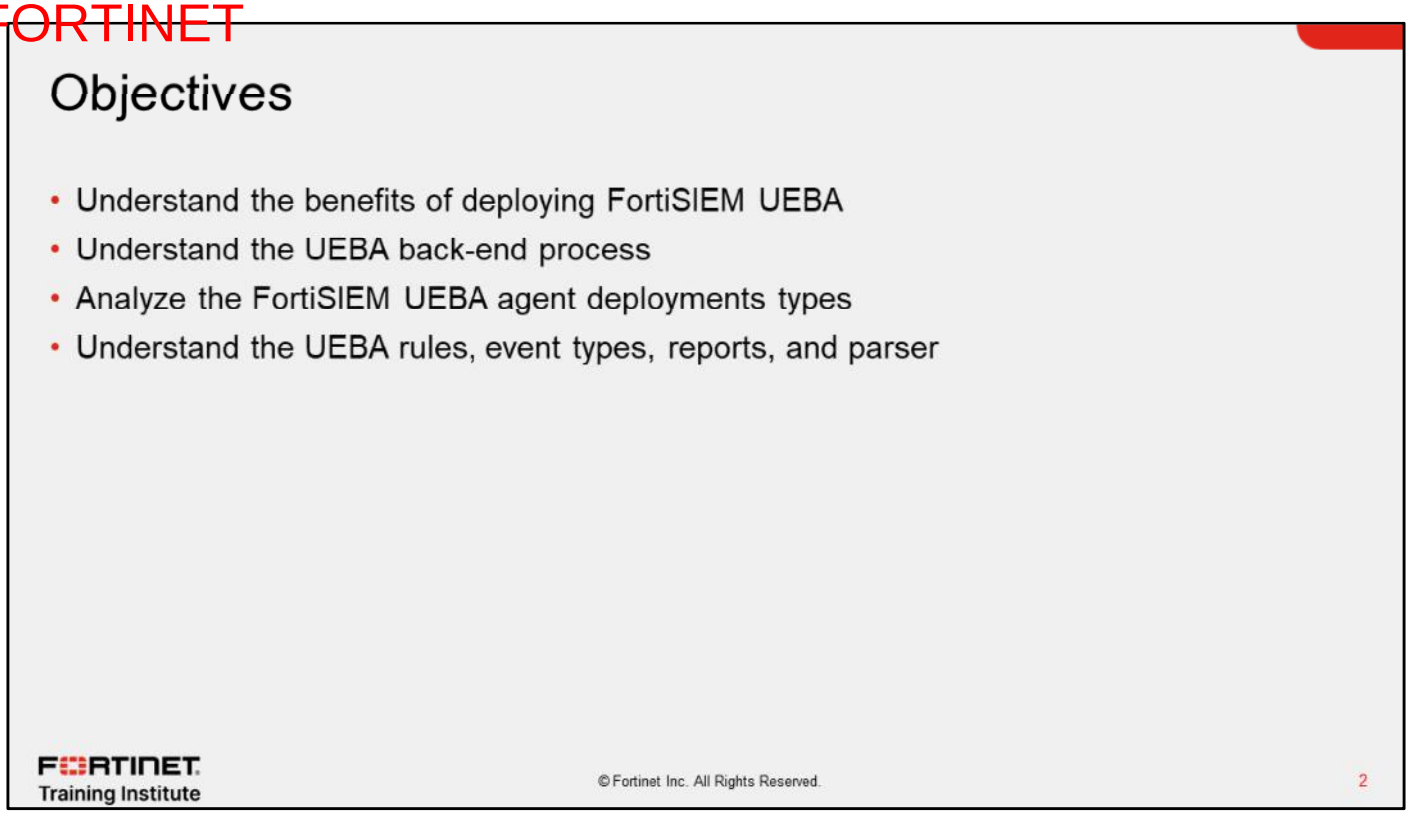
By mastering the objectives covered in this lesson, you learned how to analyze out-of-the-box baseline rules, and create your own baseline rules with statistical expressions.

DO NOT REPRINT
© FORTINET



In this lesson, you will learn about UEBA on FortiSIEM. You will also learn about the benefits of using UEBA, and deployment modes of the UEBA agent.

DO NOT REPRINT
© FORTINET



After completing this lesson, you should be able to achieve the objectives shown on this slide.

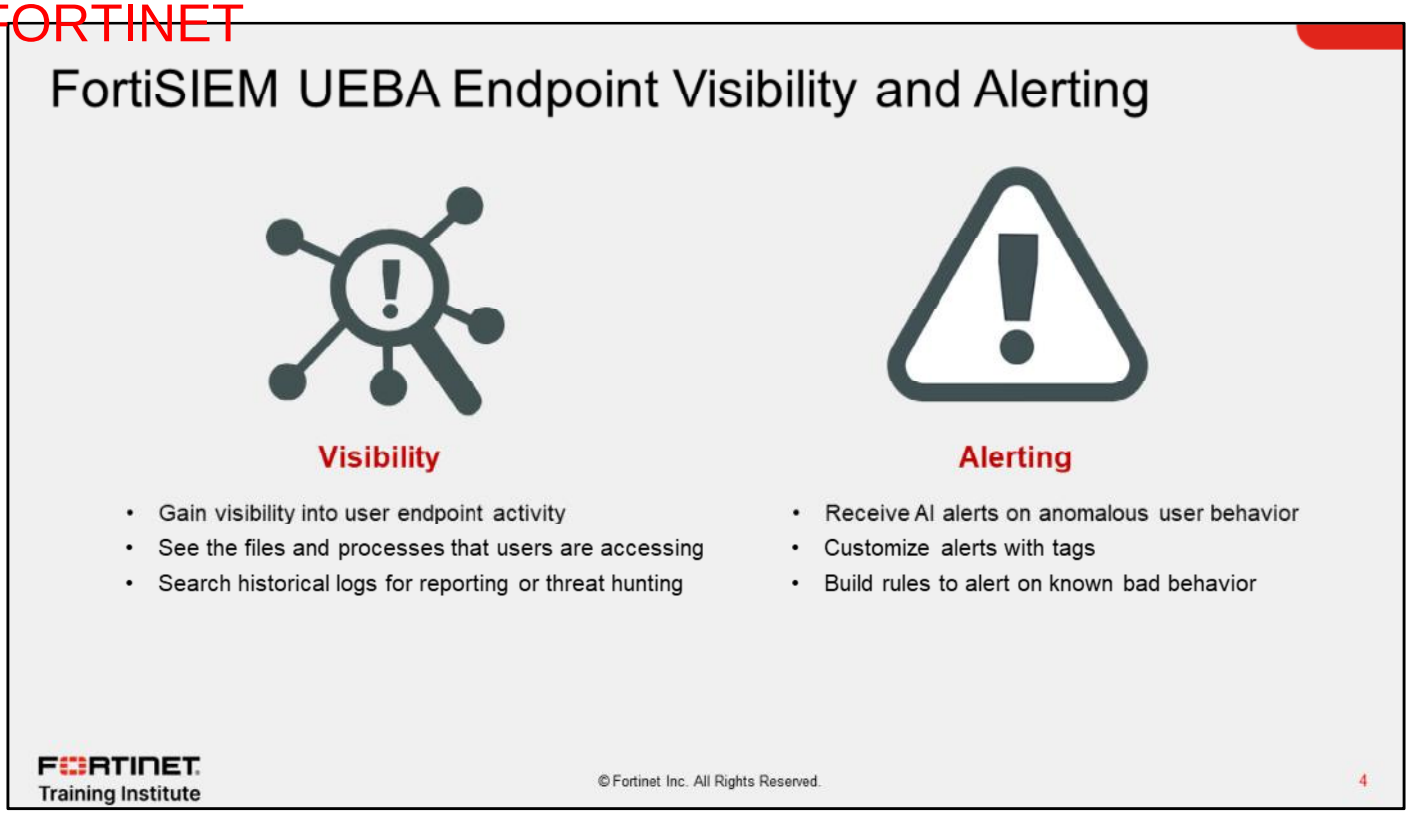
By demonstrating competence in understanding UEBA on FortiSIEM, you will be able to configure UEBA on agents and track UEBA events on FortiSIEM.

DO NOT REPRINT
© FORTINET



In this section, you will learn about UEBA and how you can deploy UEBA in a FortiSIEM service provider environment.

DO NOT REPRINT
© FORTINET



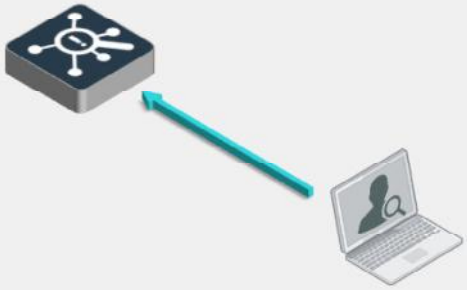
FortiSIEM UEBA provides two important functions: visibility and alerting. You can gain visibility into user endpoint activity, view the files and processes users access on their machines, and search historical logs for reporting or threat hunting. If there is anomalous user behavior, the FortiSIEM AI generates an incident. You can customize alerts with tags, and build rules to generate an incident when a known bad behavior is observed on endpoint devices.

The integration with FortiInsight brings user and entity behavior analytics (UEBA) to FortiSIEM. Previous versions provided the integration using an API, which required two separate installations. This integration requires a single installation only, with no overlapping functionality. The AI module runs on super and worker nodes. All agent activity is routed to one node in a sticky manner. If a worker is down, agent events are routed to another worker. If a worker is added, new agents are routed to that worker. Additionally, AI models are now persistent across AI module restarts.

DO NOT REPRINT
© FORTINET

Agent-Based UEBA

- Lightweight Windows endpoint agent
- Captures five key metrics
  - User
  - Process
  - Device
  - Resource (for example, file)
  - Behavior (for example, read, write, create, delete, process stop)
- Uploads to FortiSIEM for analysis
  - Minimal endpoint overhead
- Agent logs gathered directly from the kernel
  - Accurate and reliable
  - No reliance on audit policy configuration
- Encrypted TLS upload



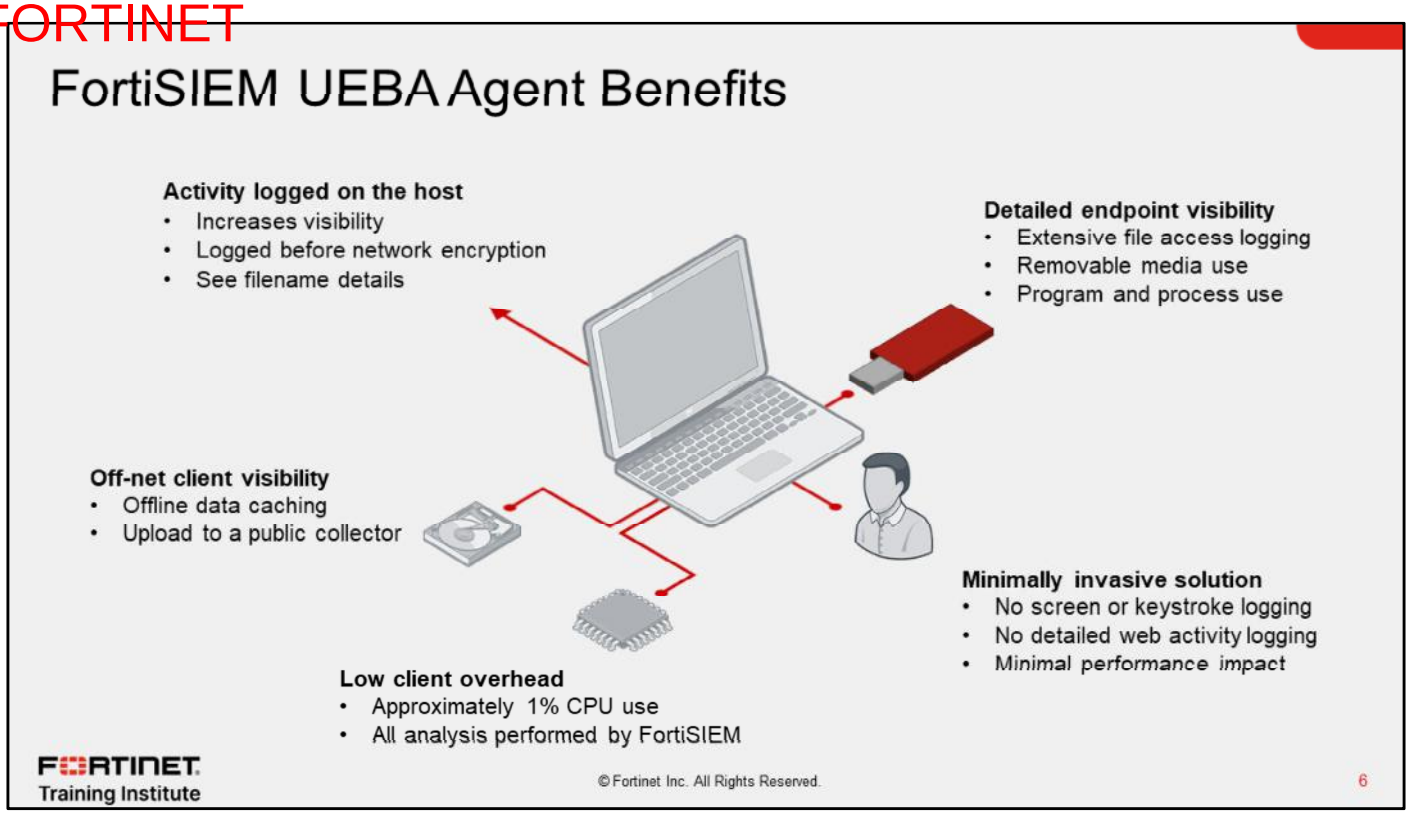
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

5

The agent-based UEBA is lightweight, and captures five key metrics: user, process, device, resource, and behavior. To reduce endpoint overhead, UEBA events are uploaded to FortiSIEM for analysis. The UEBA events are collected directly from the kernel of the device, therefore they are accurate and reliable. UEBA agents do not rely on the audit policy configuration of windows, and the events are uploaded through a secure channel.

DO NOT REPRINT
© FORTINET



Using the UEBA agent provides the following benefits:

Activity is logged on the host
- Increases visibility
- Logged before network encryption
- Can see filename details

Off-net client visibility
- Offline data caching
- Upload to the public collector

Low client overhead
- Approximately 1% CPU use
- All analysis is performed by FortiSIEM

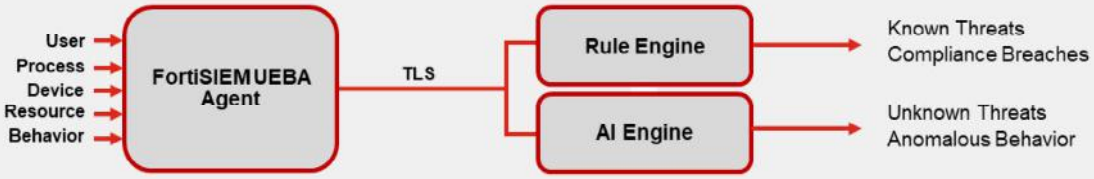
Minimally invasive solution
- No screen or keystroke logging
- No detailed web activity logging
- Minimal performance impact

Detailed endpoint visibility
- Extensive file access logging
- Removable media use
- Program and process use

DO NOT REPRINT
© FORTINET

FortiSIEM UEBA—AI and Rules

- Combined AI and rule-based detection
  - AI detects unknown threats and anomalies
  - Rules detects known threats and compliance breaches
- Increased effectiveness
- Additional functionality



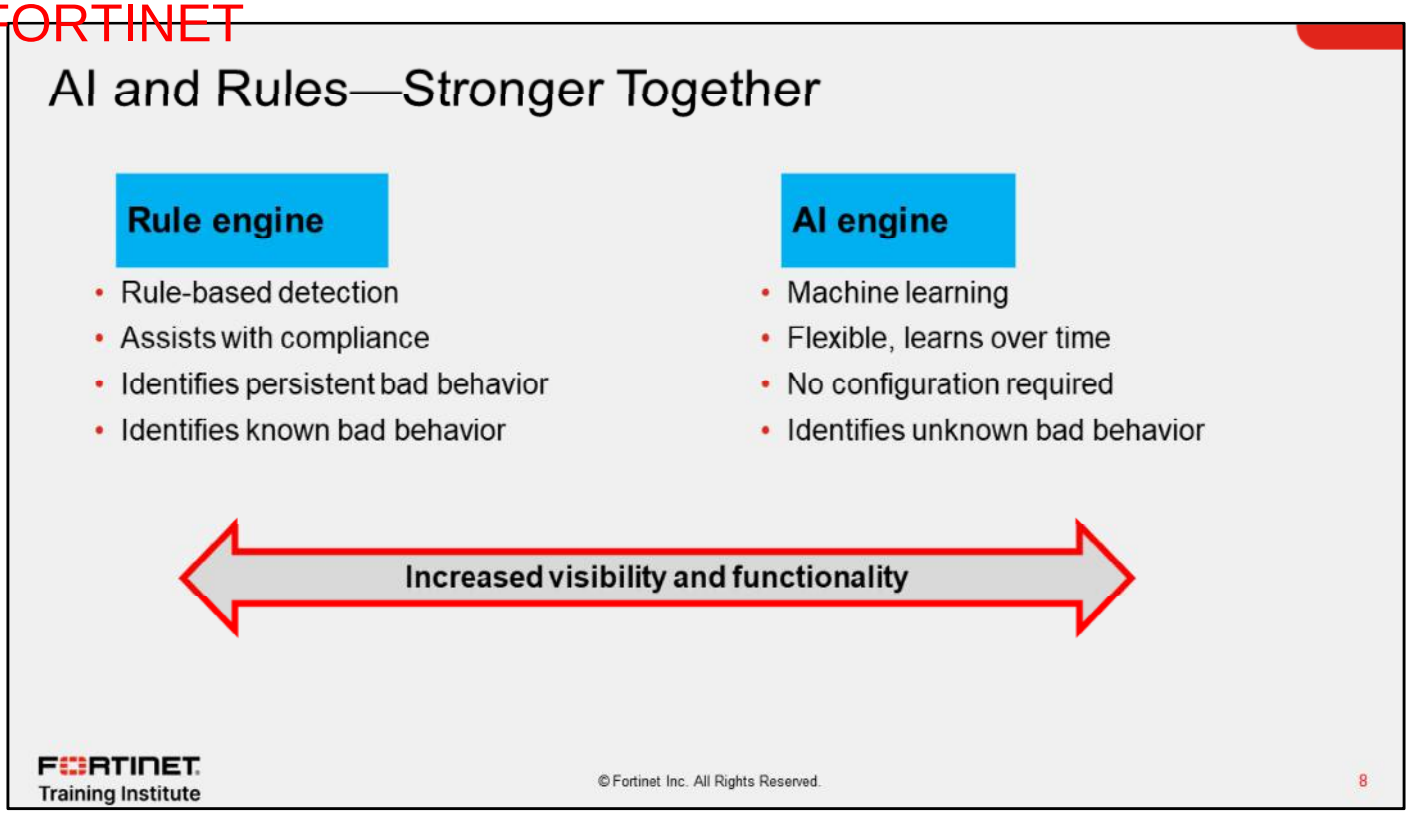
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

7

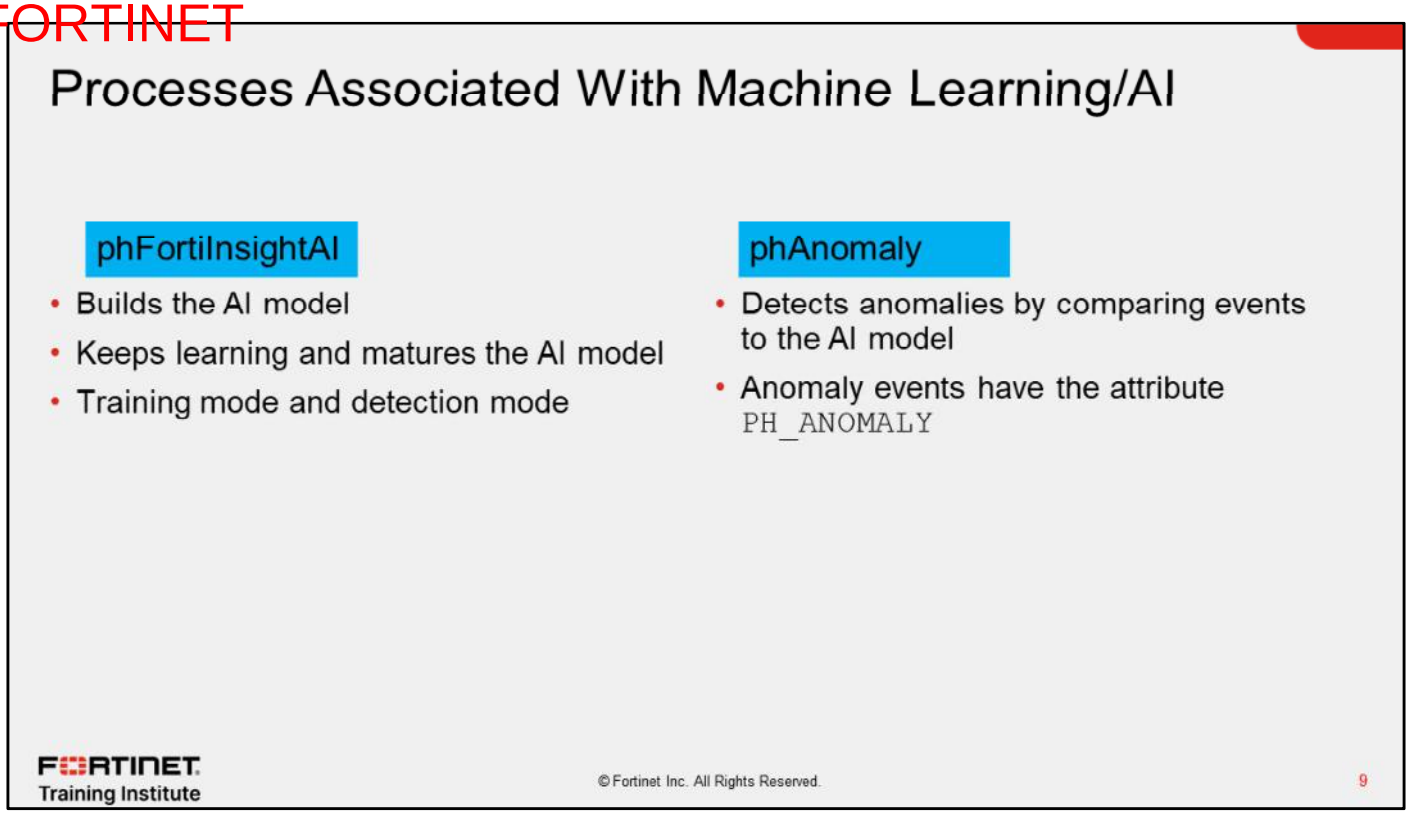
The AI module of FortiSIEM, together with the rule engine, can detect known and unknown threats. The AI module is primarily responsible for building the AI model and identifying anomalies in that model. The rules detect compliance breaches and known threats.

DO NOT REPRINT
© FORTINET



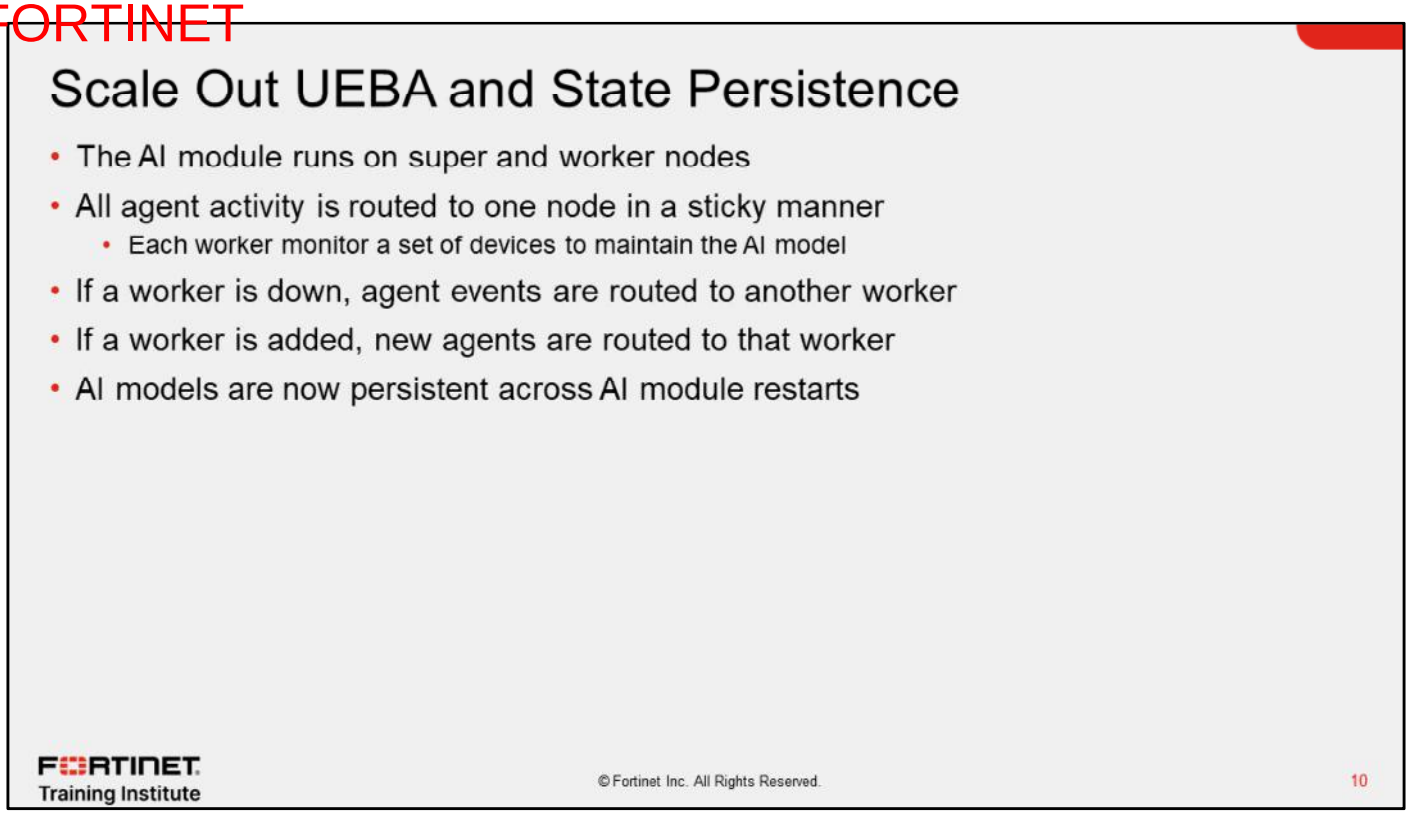
The rule engine detects incidents based on how it is configured to analyze logs. Rules are designed to track and generate incidents for known persistent bad behavior. The AI engine on the other hand learns and builds its model over time, and no configuration is required. When an anomaly is detected by the AI engine an incident is generated.

DO NOT REPRINT
© FORTINET



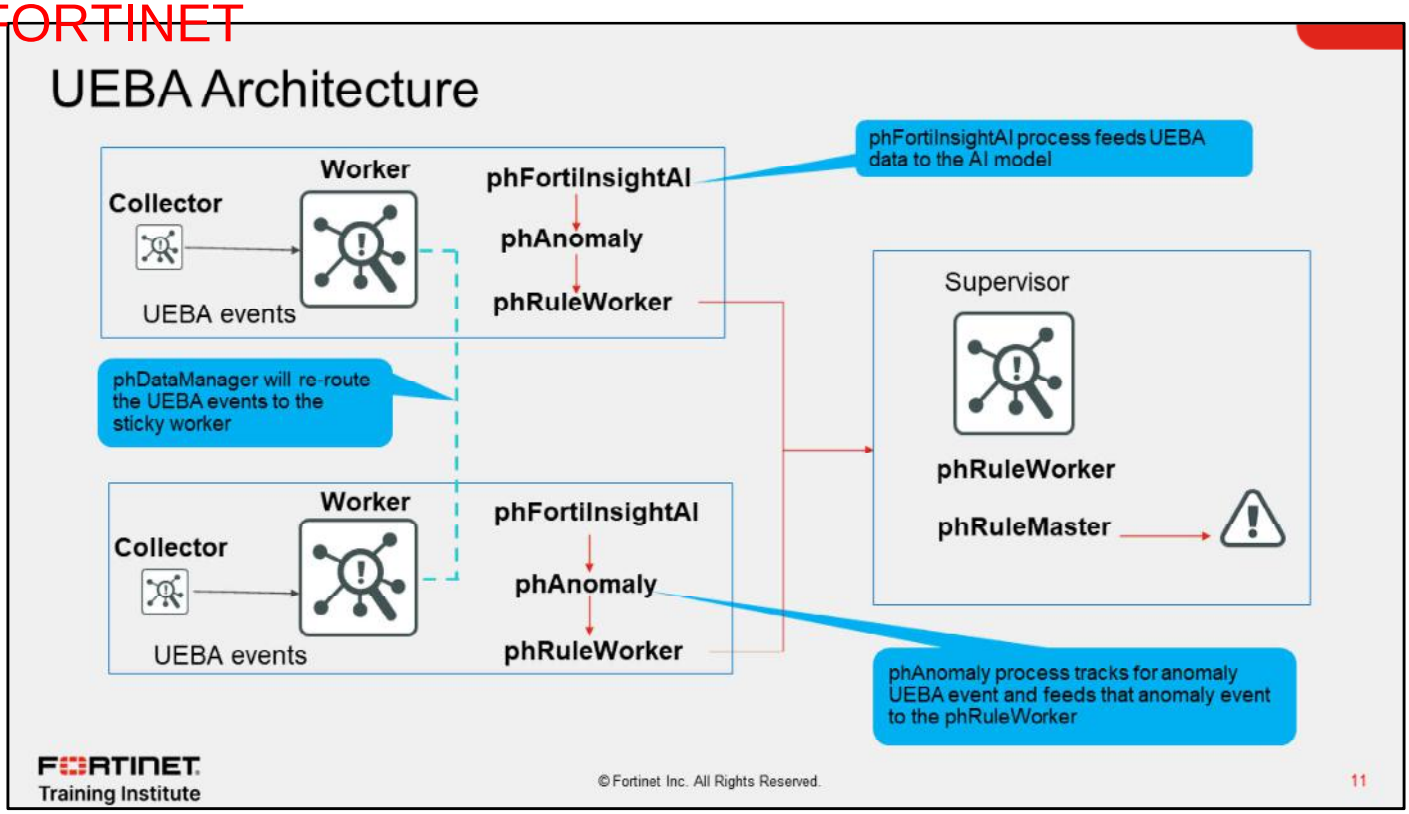
The phFortiInsightAI process is responsible for building the AI model on FortiSIEM. By default it takes 14 days to build the AI model, during this period, it is in training mode. After the AI model is built, it switches to detection mode. When it is in detection mode, the phAnomaly process is responsible for detecting anomalies that arise. The AI model continues to build and mature the model over time. Events that are anomalous, have the attribute `PH_ANOMALY`.

DO NOT REPRINT
© FORTINET



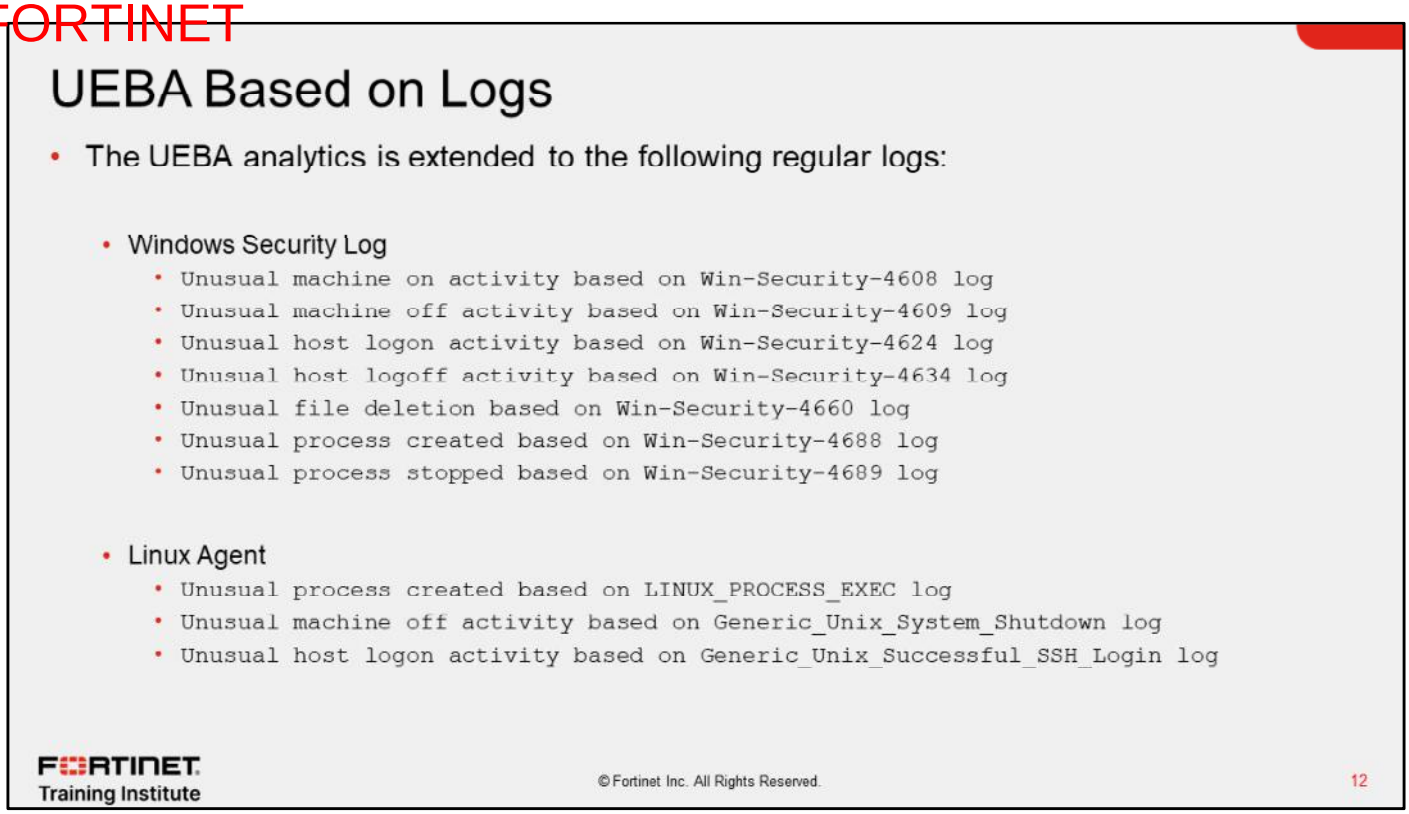
The AI module runs on super and worker nodes. All agent activity is routed to one node in a Sticky manner. Once the events are uploaded to a FortiSIEM cluster then the actual UEBA models are distributed across the workers, but each worker monitors a set of devices to maintain the AI model. If a worker is down, agent events are routed to another worker. If a worker is added, then new agents are routed to that worker. AI models are now persistent across AI module restarts.

DO NOT REPRINT
© FORTINET



Agents deployed at endpoint devices generate UEBA events which are uploaded to a collector. The collector uploads the events to the Supervisor or Worker. The UEBA events are then routed to a worker node in a sticky manner. This stickiness is based on the **Reporting Device Name** which ensures that the same worker receives the same device events to build the complete model, rather than fragmented between worker nodes. The phFortiInsightAI process on the worker continues to build the AI model by extracting various attributes from the UEBA events. The phAnomaly process generates an alert event if it detects an event that deviates from the AI model and is suspicious. The anomaly event is summarized on per UEBA rule basis by the phRuleWorker on the worker and the summarized data is sent to the supervisor. The phRuleMaster on the supervisor analyzes the event and generates an incident based on the active UEBA rules on FortiSIEM.

DO NOT REPRINT
© FORTINET



Previously, UEBA relied on proprietary logs collected by the FortiSIEM Windows UEBA Agent. In this release, the analytics is extended to the regular logs shown on this slide.
Note that these regular logs cover only a subset of the user activities compared to the FortiSIEM UEBA agent.

DO NOT REPRINT
© FORTINET

Windows UEBA vs. Log-Based UEBA

Scenario	Windows UEBA Agent	Windows Security Log	Linux Agent	Linux Log
File create	Yes	No	No	No
File delete	Yes	Yes	No	No
File read	Yes	No	No	No
File write	Yes	No	No	No
File move	Yes	No	No	No
File rename	Yes	No	No	No
File print	Yes	No	No	No
Process stop	Yes	Yes	No	No
Process create	Yes	Yes	Yes	No
File upload	Yes	No	No	No
File download	Yes	No	No	No
Machine on	Yes	Yes	No	No
Machine off	Yes	Yes	No	No
Drive mount	Yes	No	No	No
Drive un-mount	Yes	No	No	No
Host logon	Yes	Yes	No	Yes
Host logoff	Yes	Yes	No	Yes

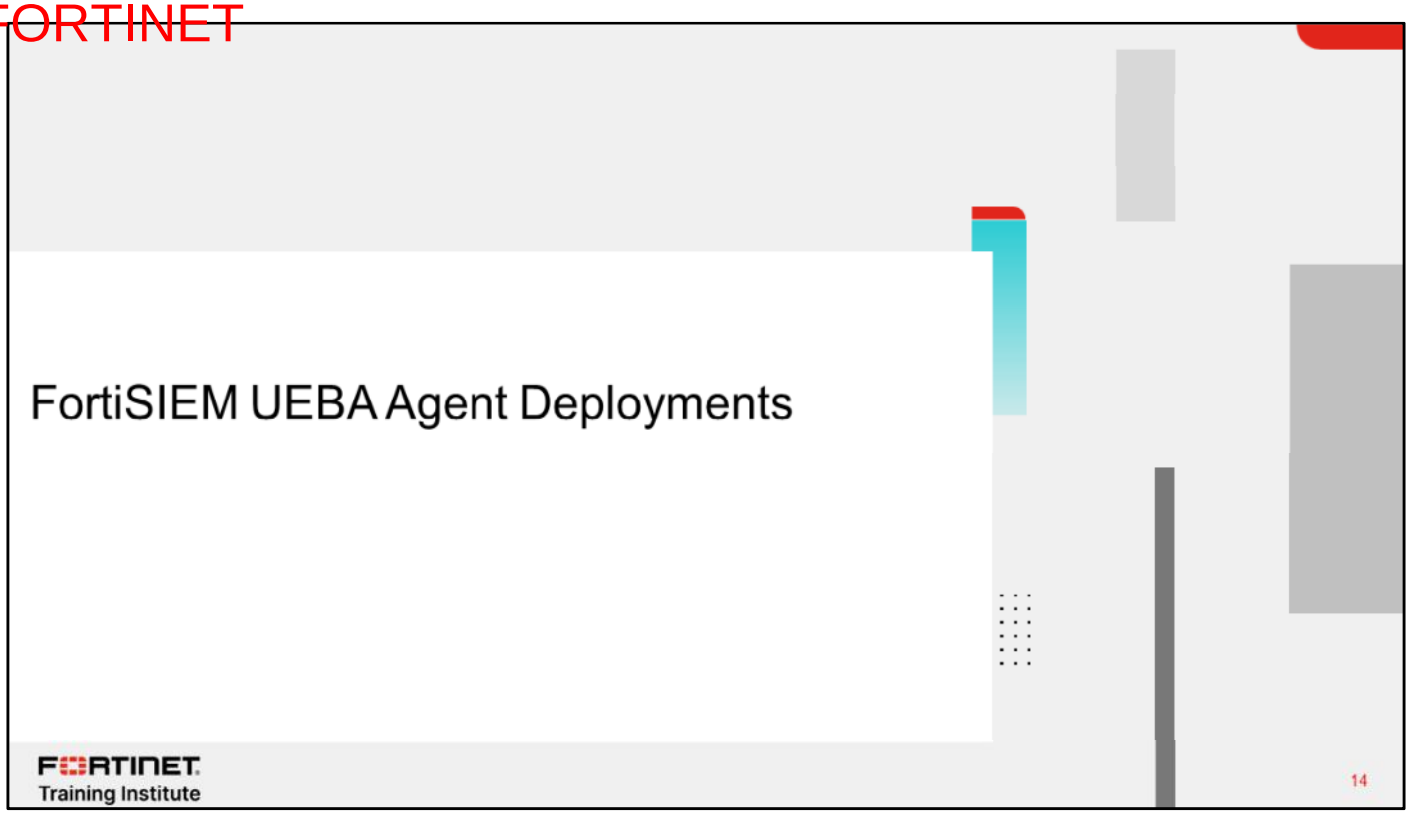
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

13

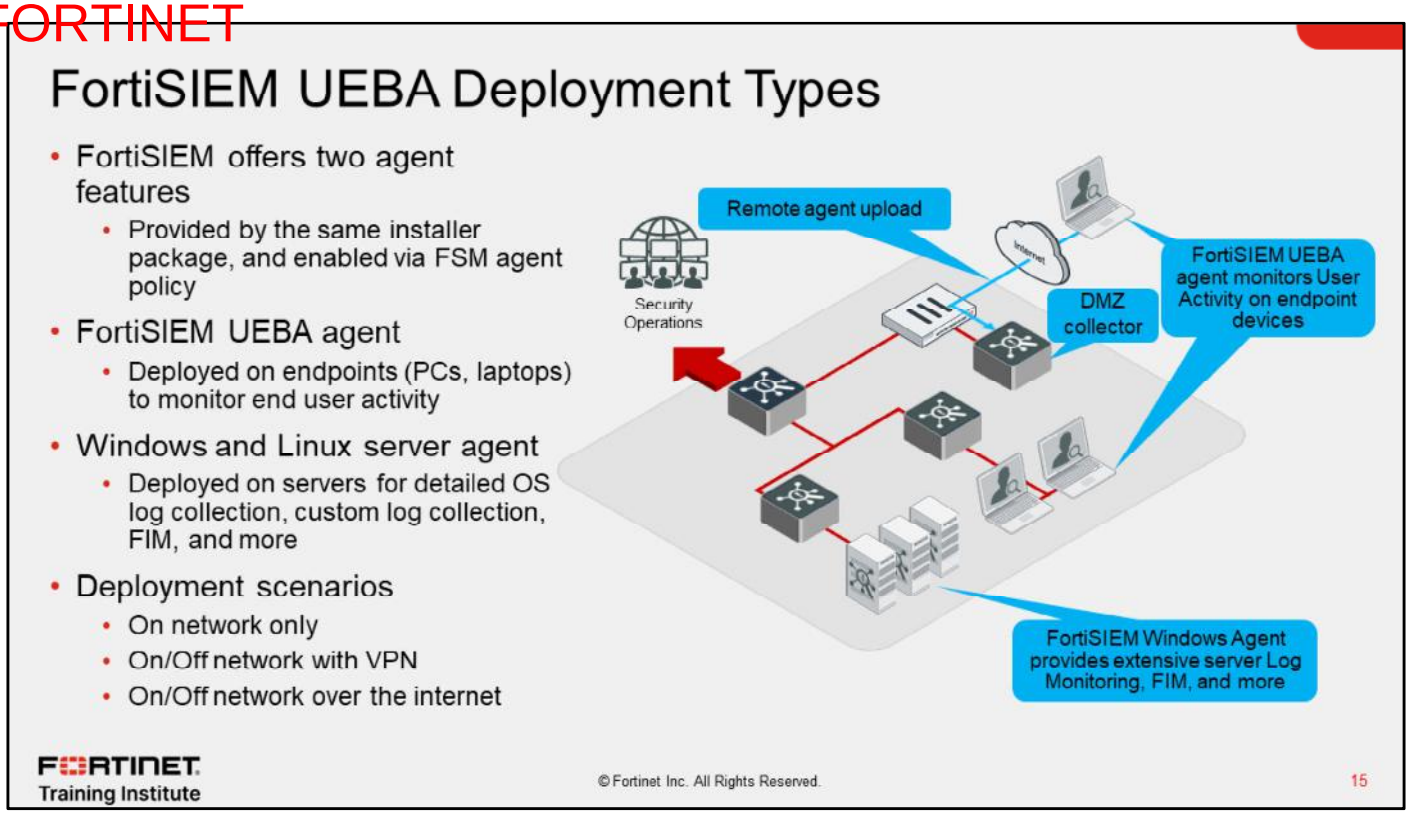
The table shown on this slide provides details on Windows UEBA agent compared to the log-based UEBA.

DO NOT REPRINT
© FORTINET



In this section, you will learn about the deployment options of FortiSIEM UEBA.

DO NOT REPRINT
© FORTINET

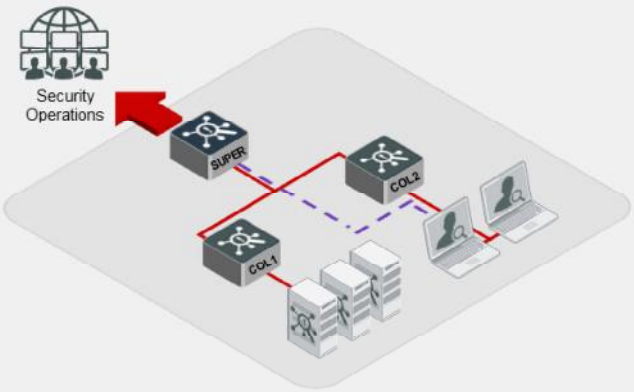


FortiSIEM offers two distinct features for UEBA agents. When the agent is deployed on endpoint devices, such as PCs and laptops, then it can be used to monitor end user activity. If the agent is deployed on servers, it can be used to collect more detailed information regarding the OS, file integrity monitoring, the processes running on the server, and collect only custom logs. The same Windows agent is deployed Windows server and on the Windows endpoint devices. The major difference is that you can enable UEBA for endpoint devices through the Windows agent template and that instructs the agent at the endpoint devices to collect UEBA events.
There are three different types of UEBA deployment scenarios: on network only, on/off network with VPN, and on/off network over the internet.

DO NOT REPRINT
© FORTINET

FortiSIEM Agent Deployments—On-Net Only

- Agents communicate directly with the collectors on the network
- Agents communicate directly to the supervisor on the network or the collectors configured as proxy agent communication to the supervisor
- **Pros**: Simple to deploy and appropriate for some organizations
- **Cons**: Will not monitor off net



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

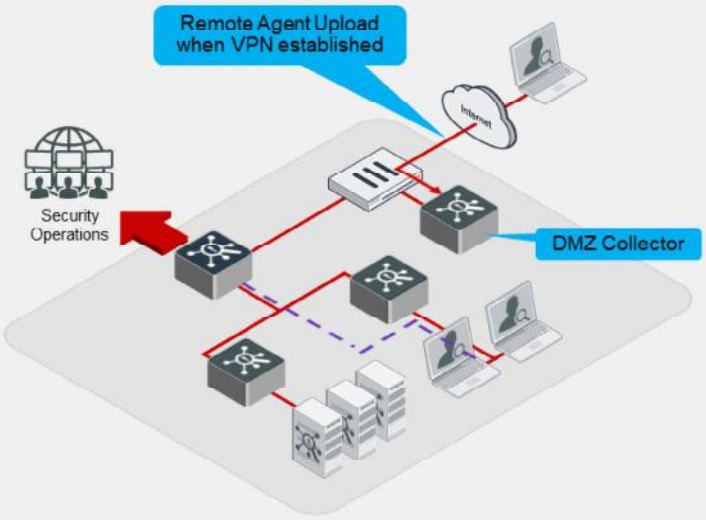
16

For on-net setup, the agents communicate directly with the collectors on the network. The events from the endpoint devices are uploaded directly to the collector or supervisor. The endpoint device is always on the corporate network and has direct access to the supervisor or collector. Health information for the agent is reported directly to the supervisor. If the endpoint devices do not have direct access to the supervisor, it can connect through a proxy to the collector, and then communicate to the supervisor. The collector can be used to proxy the health information to the supervisor.
The advantage of this setup is that it is a simple configuration and it can be appropriate for certain organizations that require endpoint devices to be on the network at all times. The disadvantage is that if the endpoint devices move outside the organization, it becomes difficult to monitor those devices.

DO NOT REPRINT
© FORTINET

FortiSIEM Agent Deployments—On/Off-Net Through VPN

- Agents communicate to on-net collector
- Agents communicate directly to collectors on the network
- Agents communicate directly to the supervisor on the network or the collectors configured to proxy agent communication to the supervisor
- **Pros**: Able to monitor endpoint devices on-net and off-net
- **Cons**: Requires VPN for the agent to upload events



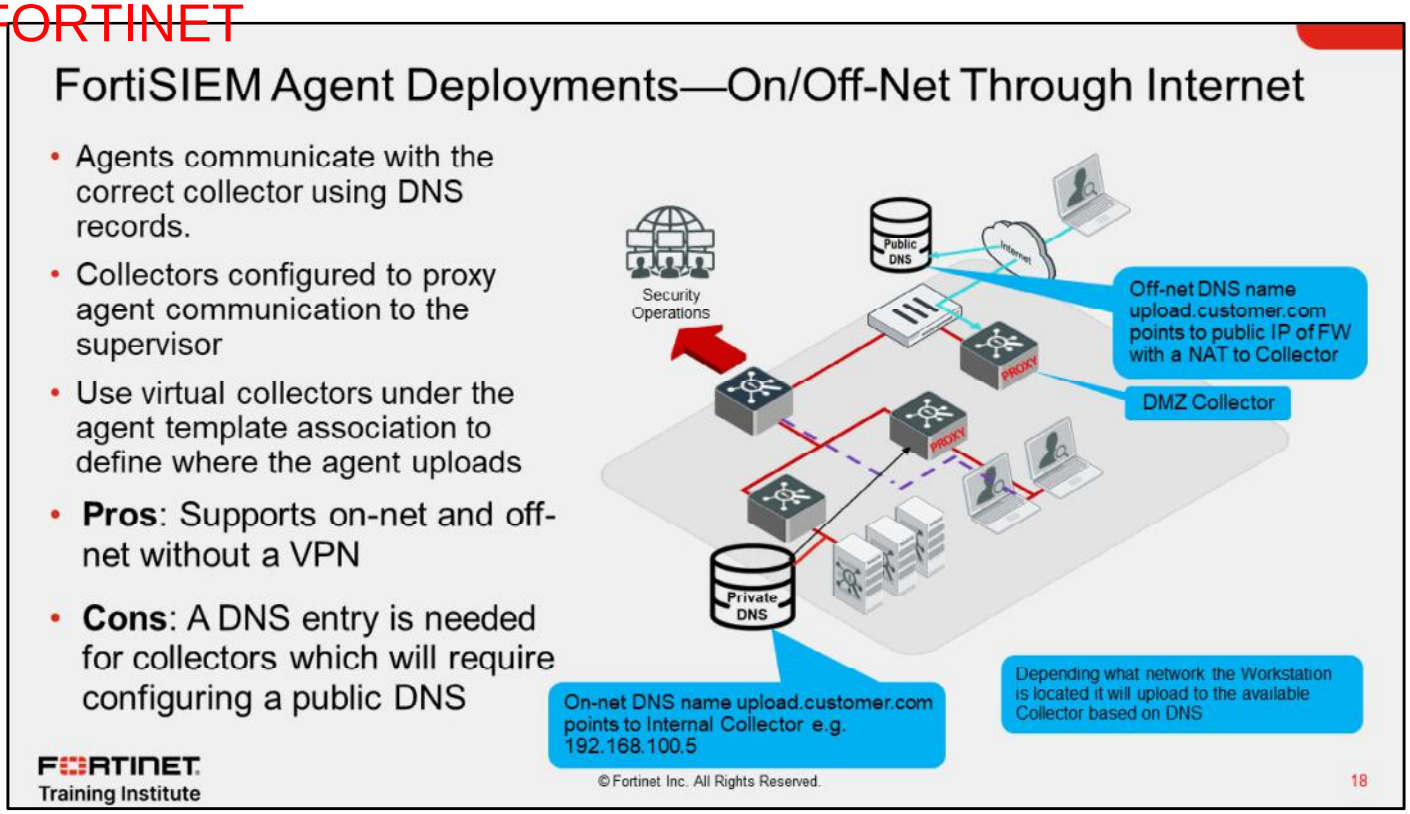
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

17

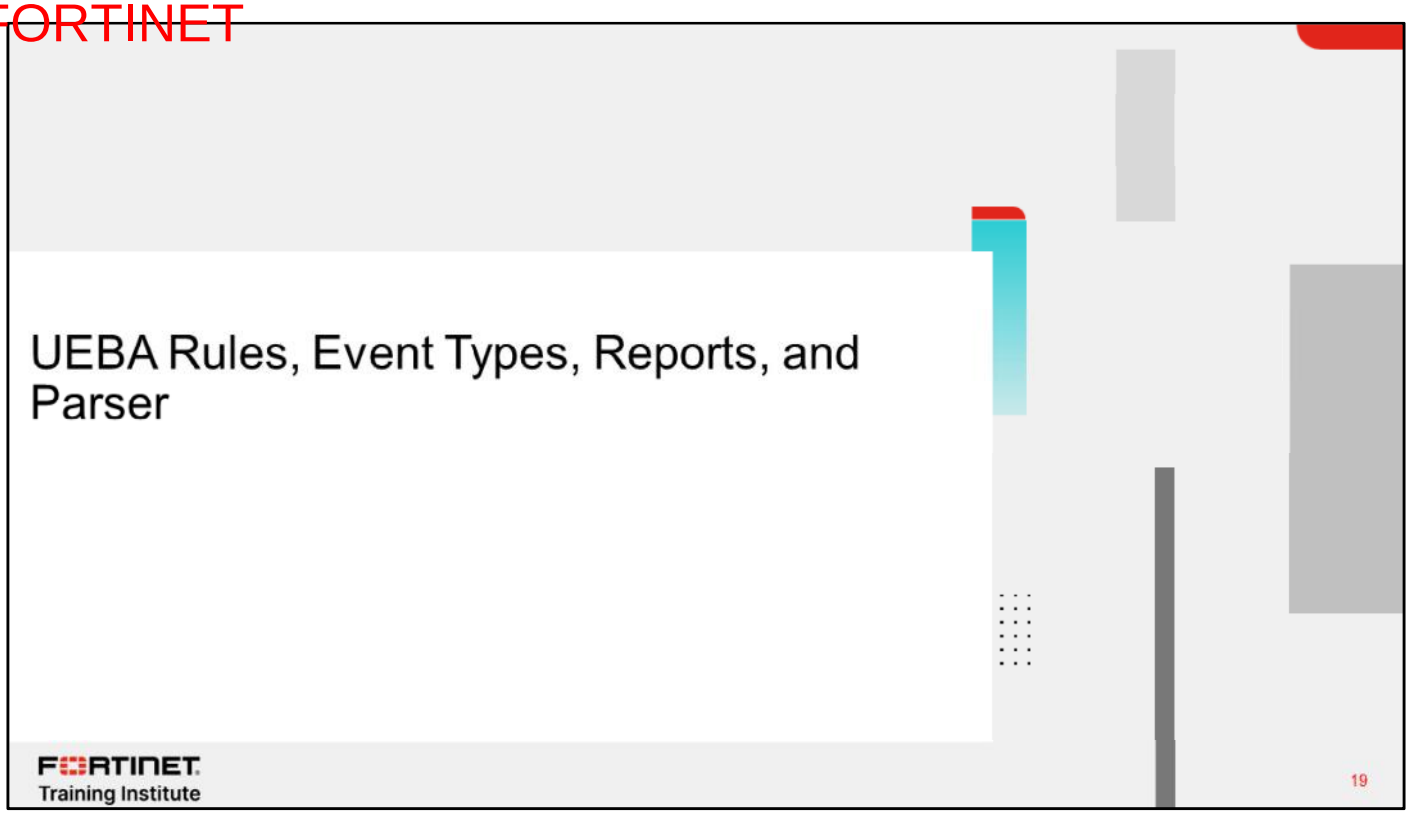
To address the issue of off network, you can deploy the agent to communicate directly with a collector that is deployed on the network when the device is on the corporate network. When the endpoint device is off the network, you can establish a VPN connection to the corporate network, and then upload the events to a separate collector located on a DMZ network. To report the health of the agent, the collector on the DMZ network could proxy the information to the supervisor. If the supervisor can be reached directly through the VPN, the endpoint devices could report the health status directly to the supervisor. The advantage is that you can monitor the endpoint devices when they are on the network and when they are off the network. The disadvantage is that you require the VPN connection to upload events to the collectors.

DO NOT REPRINT
© FORTINET



In this setup, the agents communicate to the collector in an off-network environment through the internet. This model requires some DNS configuration because the agent must communicate with the collector without a VPN connection. When the endpoint device is on the network the DNS record must point to the collector that is located inside the network. In the example on this slide the DNS record would point to the internal collector IP address when the agent is on the network.
For an off-network environment, you should configure your public DNS server with the same FQDN, but it should point to the public IP address of the firewall. Also, you must configure a VIP on the firewall to map the public IP address to the private IP address of the collector that is located on the DMZ network.
Your internal and your public DNS should a resolvable DNS name pointing to the appropriate collector. Then, agents can communicate to the supervisor via a proxy through the collector to report agent health information.

DO NOT REPRINT
© FORTINET



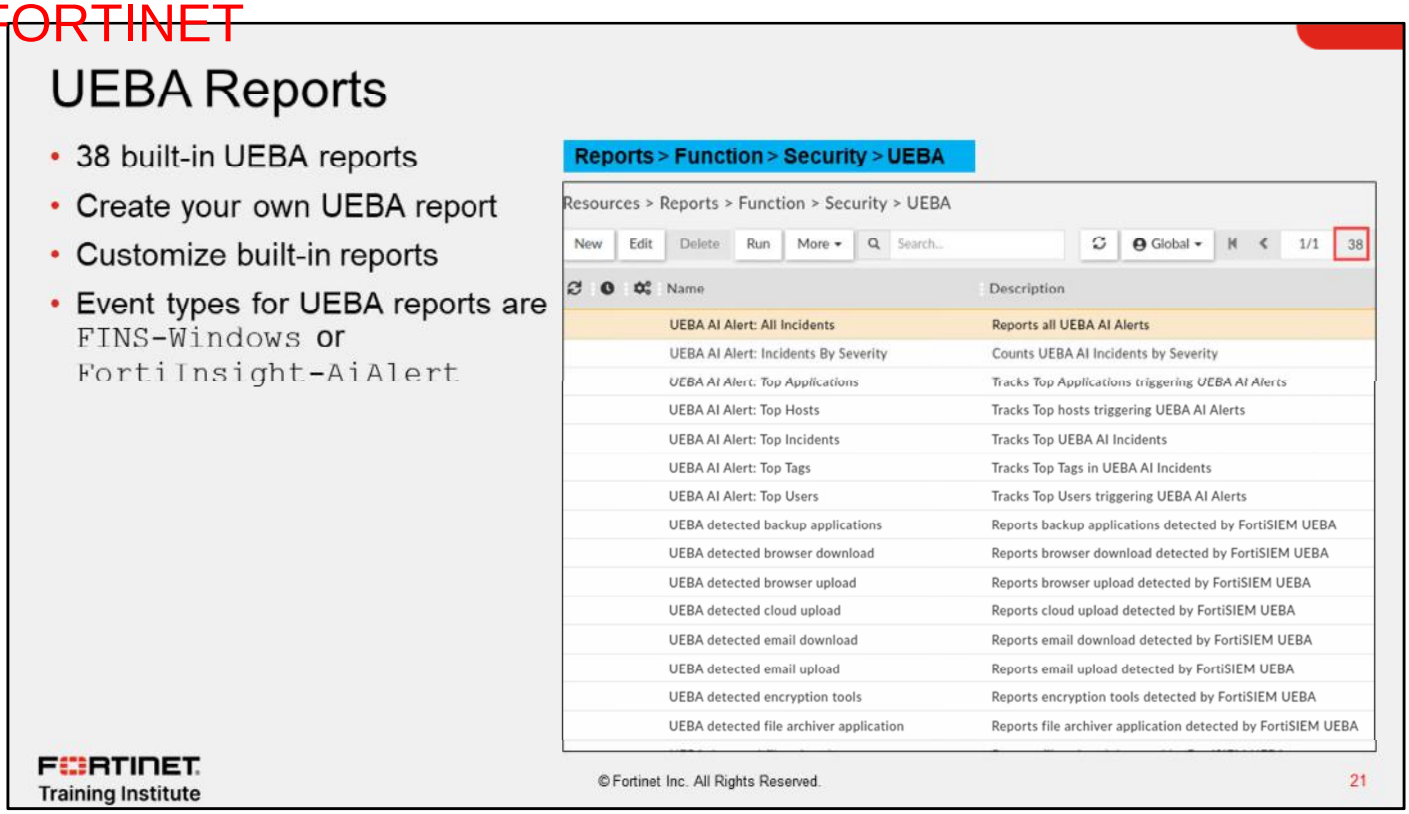
In this section, you will learn about the UEBA rules, event types, parser, dashboard, incident page, report, and watchlist. You will also learn how to configure UEBA tags.

DO NOT REPRINT
© FORTINET



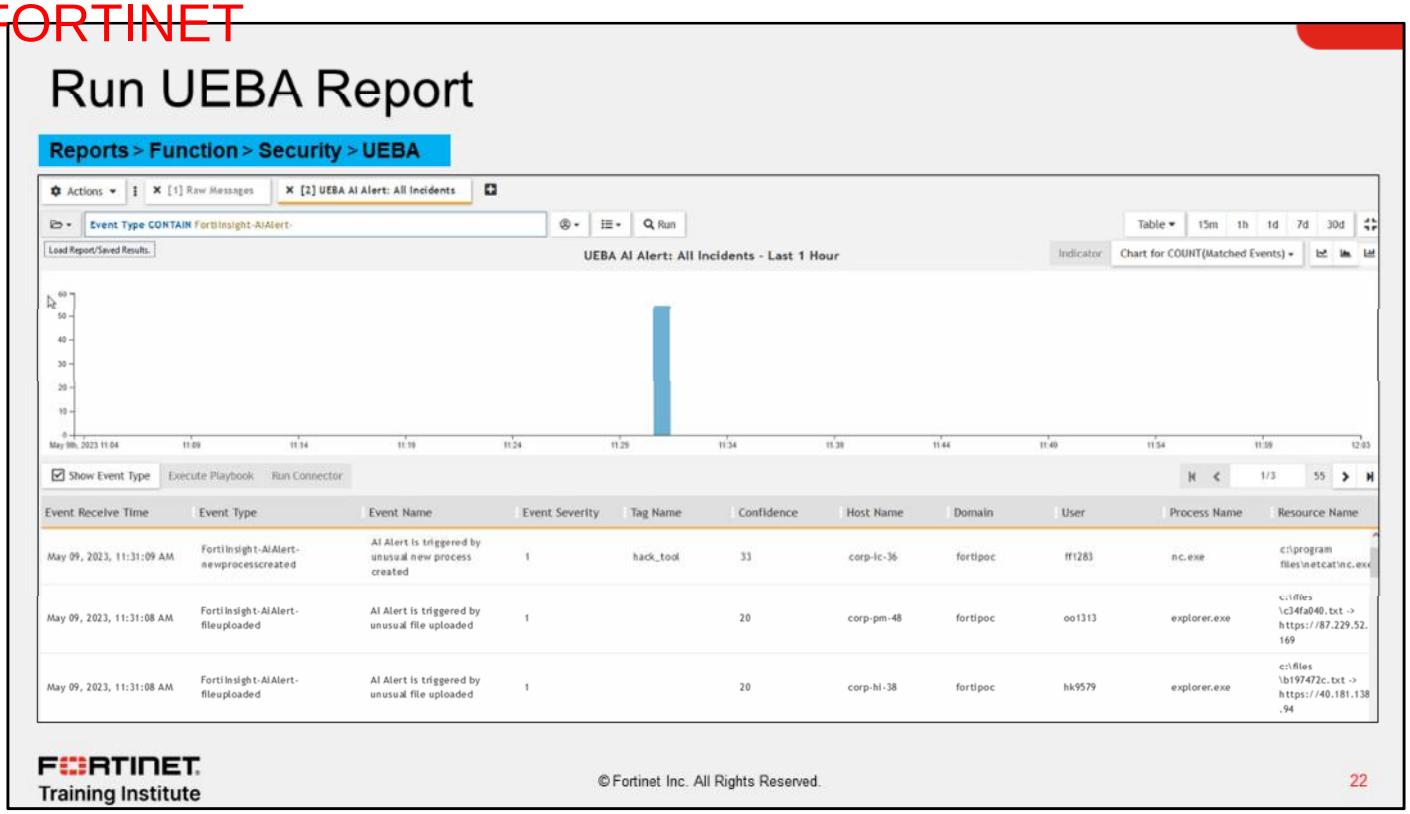
The UEBA alerts dashboard displays UEBA incidents by severity, top UEBA hosts, top UEBA applications, top UEBA tags, top users, and so on. You can also add your own customized UEBA widgets.

DO NOT REPRINT
© FORTINET



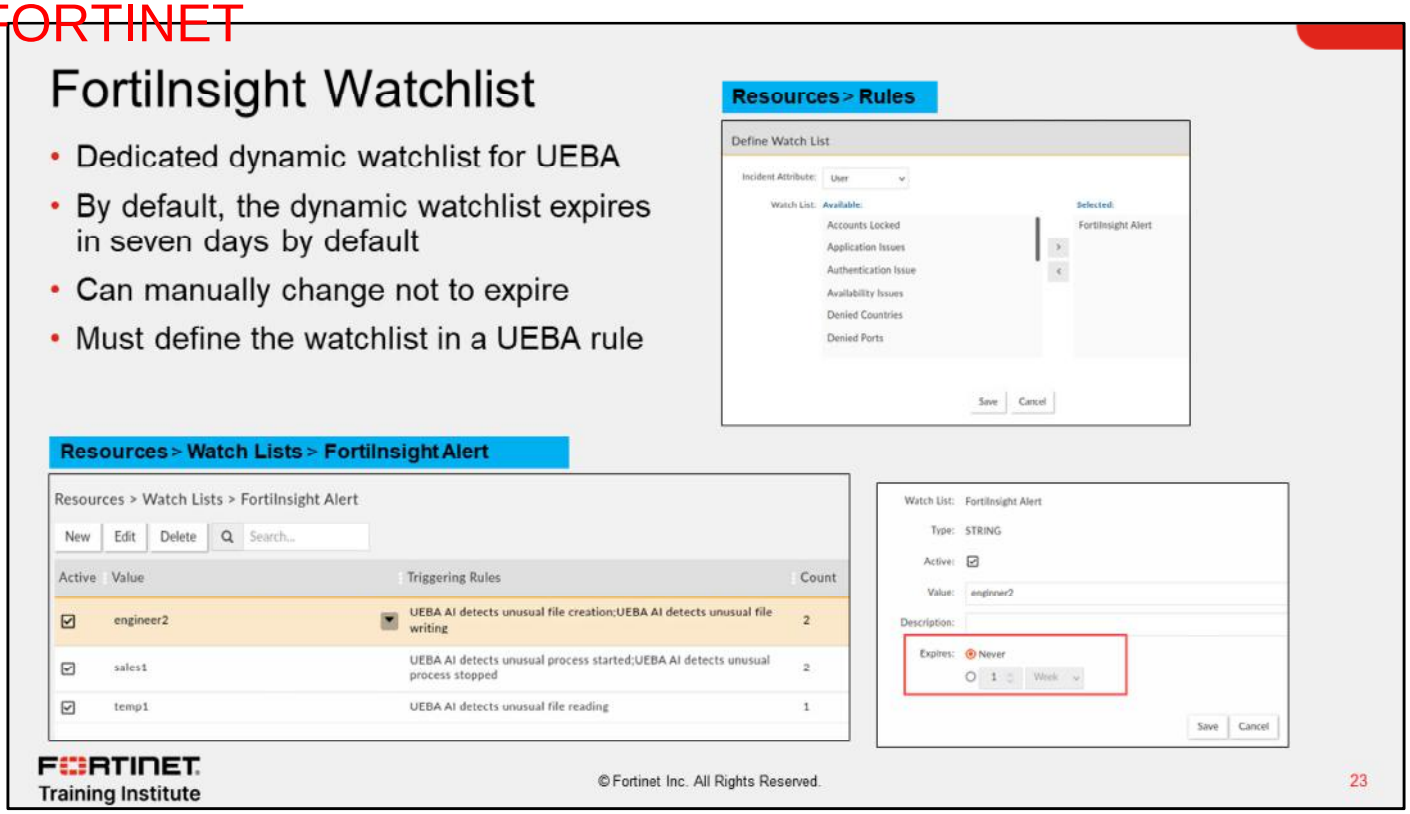
There are 38 built-in UEBA reports on FortiSIEM. You can clone any built-in report, and then edit it to customize the report according to your requirements. `FINS-Windows` and `FortiInsight-AiAlert` are two distinct event types that identify UEBA events.

DO NOT REPRINT
© FORTINET



This slide shows an example of a UEBA report of all incidents generated by the AI engine.

DO NOT REPRINT
© FORTINET



There is a dedicated watchlist for UEBA entities that can be populated dynamically as long as you define a **FortiInsight Alert** watchlist on the UEBA rule. The entries in the watchlist expire after seven days. If you want to permanently keep specific entries, you must set the expiry for those entries to **Never**, so that they are kept in the database indefinitely.

DO NOT REPRINT
© FORTINET

UEBA Rules Back End

GNU nano 2.9.8 FIN_UEBA_RULES.xml

```
<Rules>
  <Rule group="PH_SYS_RULE_UEBA" id="PH_Rule_UEBA_1" phIncidentCategory="Application" function="Security" subFunction="UEBA" technique="T1074.001">
    <Name>UEBA AI detects unusual file creation</Name>
    <active>true</active>
    <Description> FortiSIEM UEBA AI module detects unusual file Creation by a user </Description>
    <CustomerScope groupByEachCustomer="true"> <Include_all="true"/> <Exclude/> </CustomerScope>
    <IncidentDef eventType="PH_RULE_UEBA_AI_FILE_CREATED" severity="7">
      <ArgList>
        hostName = UEBA.hostName,
        domain = UEBA.domain,
        user = UEBA.user,
        tagName = UEBA.tagName,
        procName = UEBA.procName,
        fileName = UEBA.fileName,
        activityName = UEBA.activityName,
        resourceName = UEBA.resourceName
      </ArgList>
    </IncidentDef>
    <IncidentTitle>UEBA AI File creation anomaly from $user at $hostName file $fileName</IncidentTitle>
    <DynWatchListDef attr="user">
      <WatchList>PH_DYNLIST_UEBA</WatchList>
    </DynWatchListDef>
    <PatternClause window="300">
      <SubPattern displayName="UEBA" name="UEBA">
        <SingleEvtConstr> eventType = "FortiInsight-AiAlert-filecreated"</SingleEvtConstr>
        <GroupByAttr> hostName,domain,user,tagName,procName,fileName,activityName,resourceName </GroupByAttr>
        <GroupEvtConstr>COUNT(*) &gt;= 1</GroupEvtConstr>
      </SubPattern>
    </PatternClause>
    <TriggerEventDisplay>
      <AttrList> phRecvTime,hostName,domain,user,tagName,procName,fileName,activityName,resourceName</AttrList>
    </TriggerEventDisplay>
  </Rule>
```

Located in /opt/phoenix/data-definition/rules/FIN_UEBA_RULES.xml

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

24

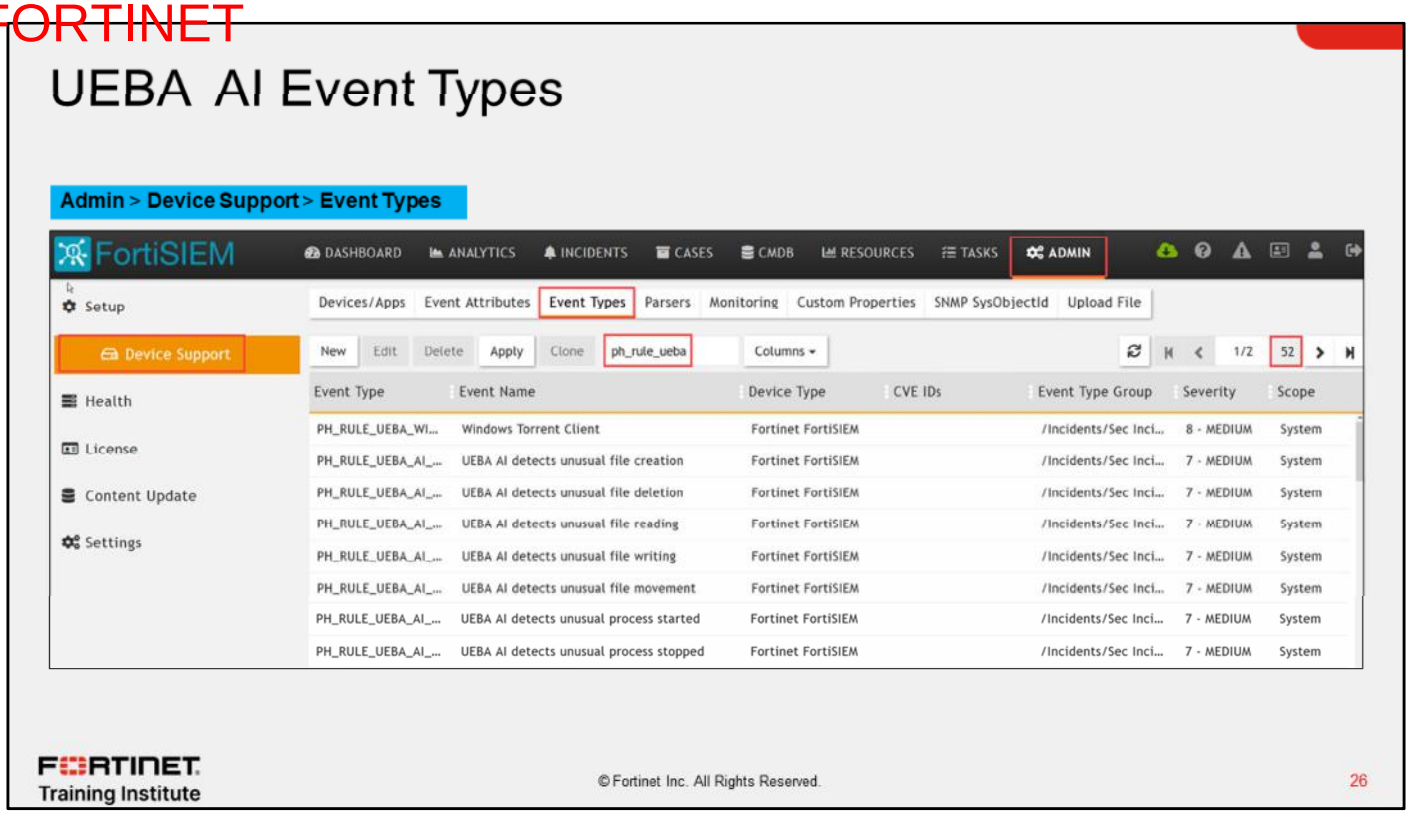
The UEBA rules are defined in the `/opt/phoenix/data-definition/rules` directory. The name of the XML file is `FIN_UEBA_RULES.xml`. All UEBA rules belong to the UEBA group called `PH_SYS_RULE_UEBA`.

DO NOT REPRINT
© FORTINET



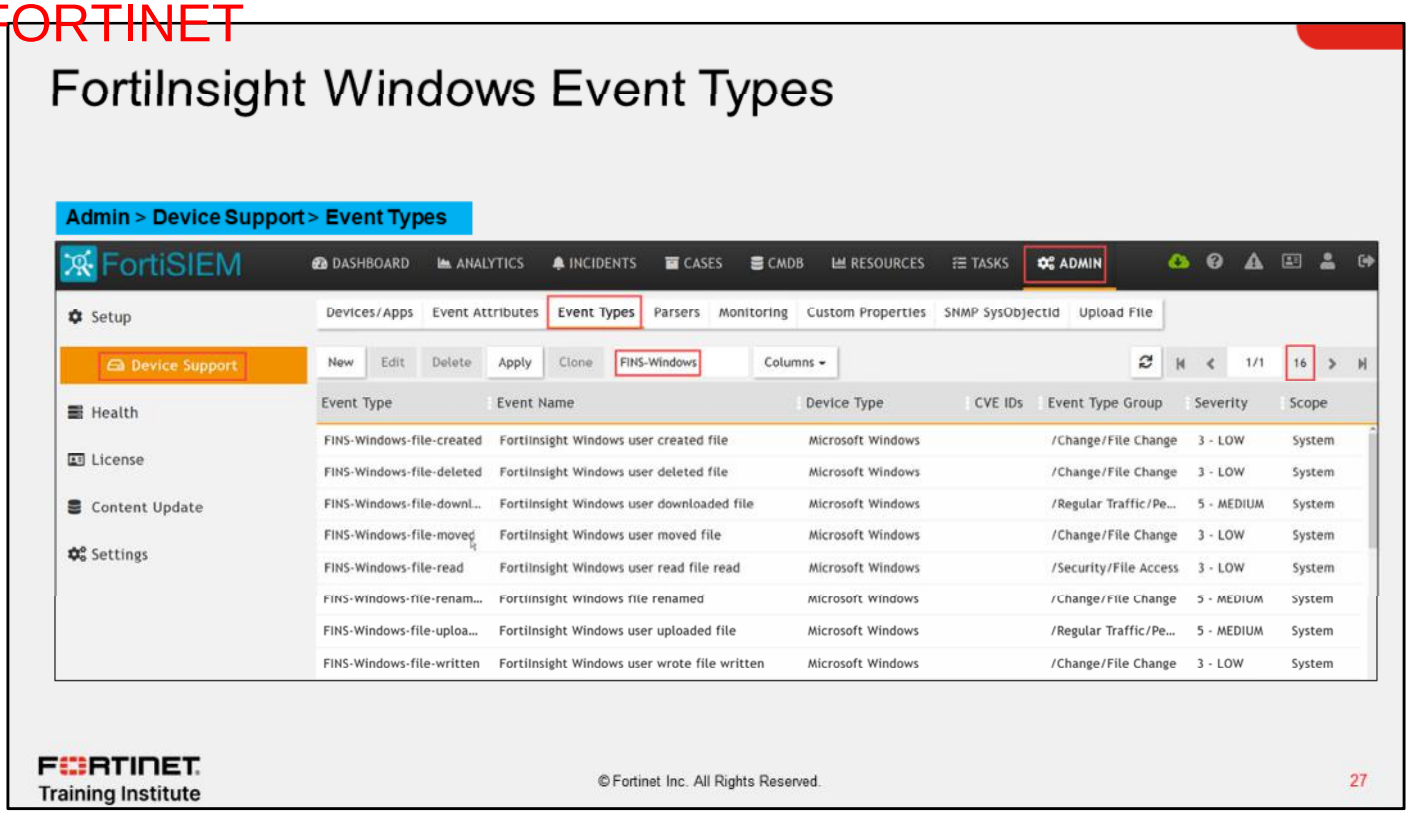
There are 51 out-of-box UEBA rules on FortiSIEM. A few of these rules are active by default, while others are inactive. You activate or deactivate a rule based on your requirements. You cannot delete the built-in UEBA rules, but you can clone and customize the rules.

DO NOT REPRINT
© FORTINET



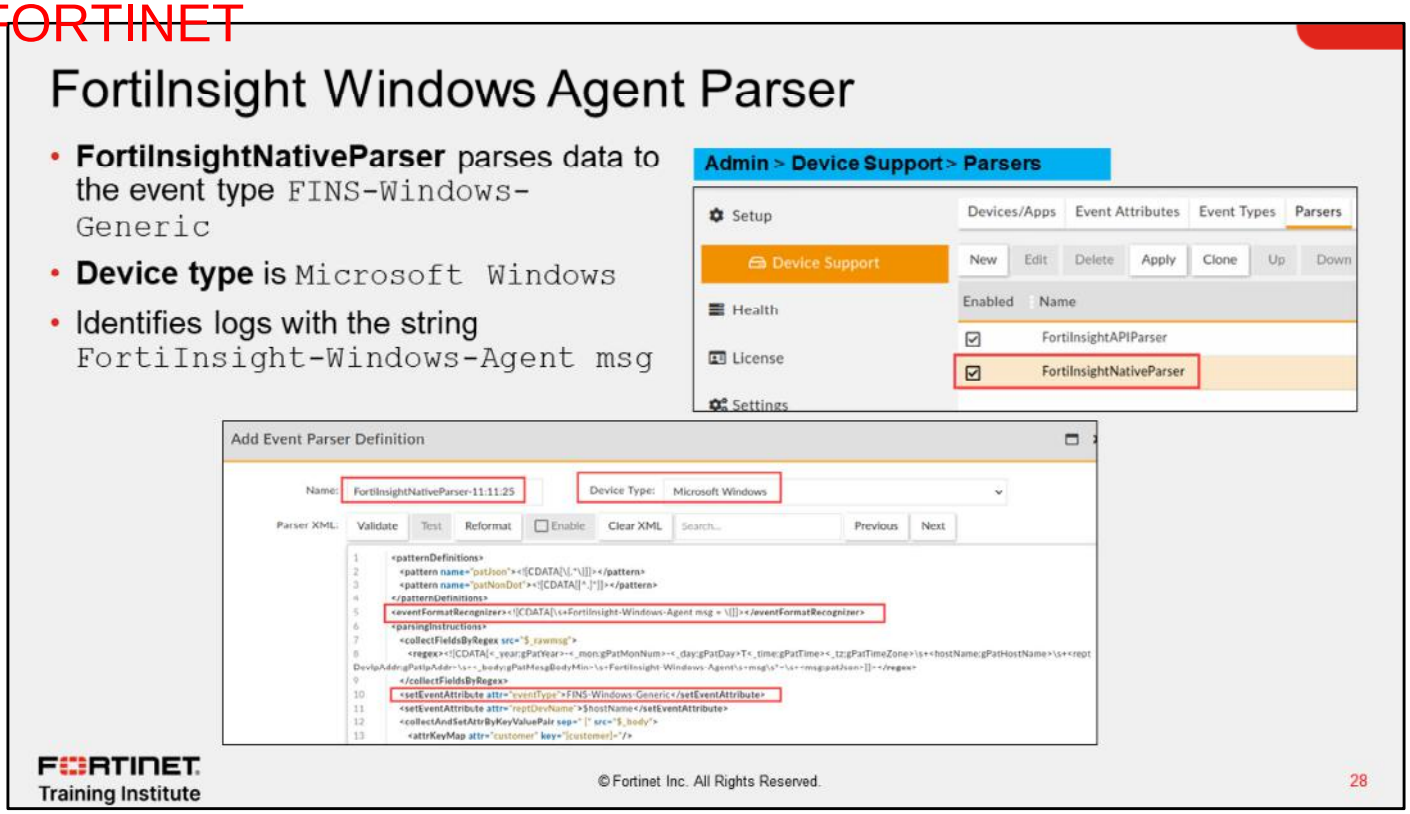
There are 52 built-in UEBA event types on FortiSIEM. You can clone and modify any of those built-in event types or you can create your own UEBA event types.

DO NOT REPRINT
© FORTINET



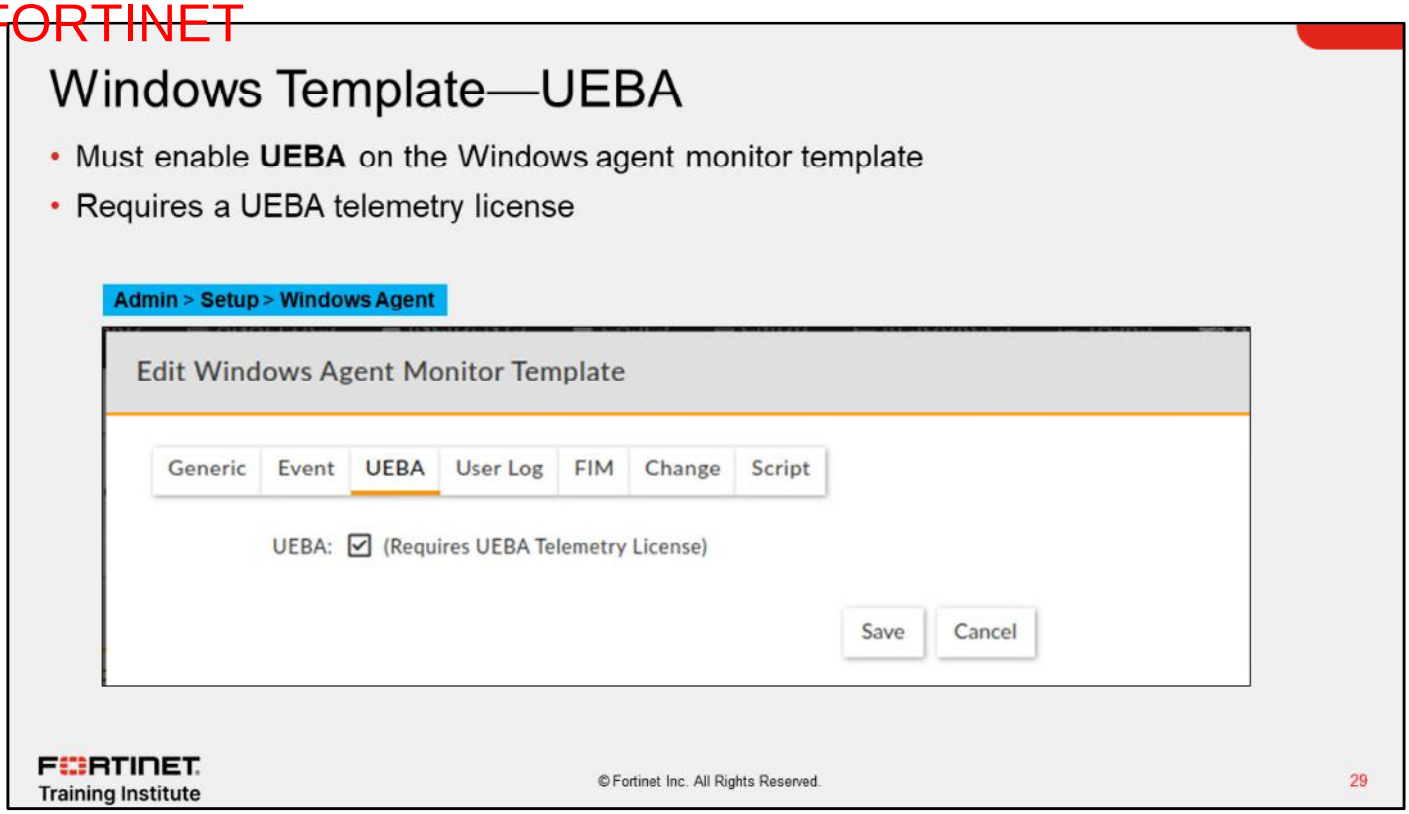
There are 16 built-in FortiInsight Windows event types on FortiSIEM. You can clone and modify any of these built-in event types or you can create your own UEBA event types.

DO NOT REPRINT
© FORTINET



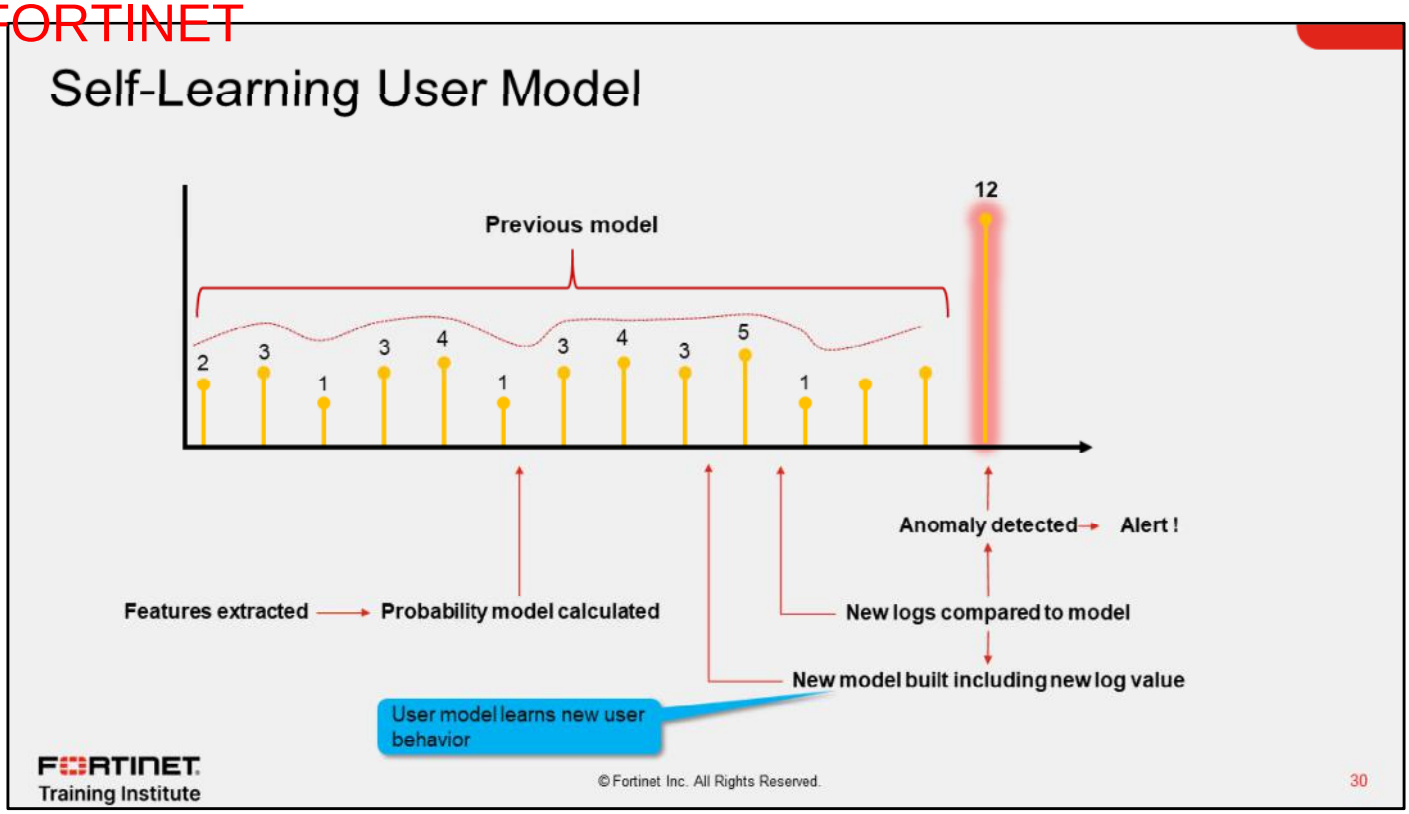
The **FortiInsightNativeParser** is a built-in parser that recognizes FortiInsight Windows agent logs. The logs sent by the UEBA agent that is installed on a Windows endpoint is parsed by this parser. The event type is mapped to `FINS-Windows-Generic`. The parser identifies logs that have the string `FortiInsight-Windows-Agent msg`.

DO NOT REPRINT
© FORTINET



For the Windows agent to collect UEBA events, you must enable UEBA on the Windows agent monitor template. You must also ensure that your FortiSIEM license supports UEBA.

DO NOT REPRINT
© FORTINET



The process of building the AI model is continuous. The model matures over time, and the more data that is fed into the AI model, the more accurate the model is at predicting anomalies. When new logs come in to the AI model, features from those logs are extracted and feed into the AI model. If an event appears to deviate from the AI model, an anomaly alert is triggered.

DO NOT REPRINT
© FORTINET

UEBA Higher Risk Entities

- **File Types**: Enter the type of file you want to monitor
- **File Paths**: Enter the path to the folder you want to monitor
- **User Accounts**: Enter the name of the Windows agent-side user account you want to monitor
- **Group Names**: Enter the name of the Windows agent-side group you want to monitor



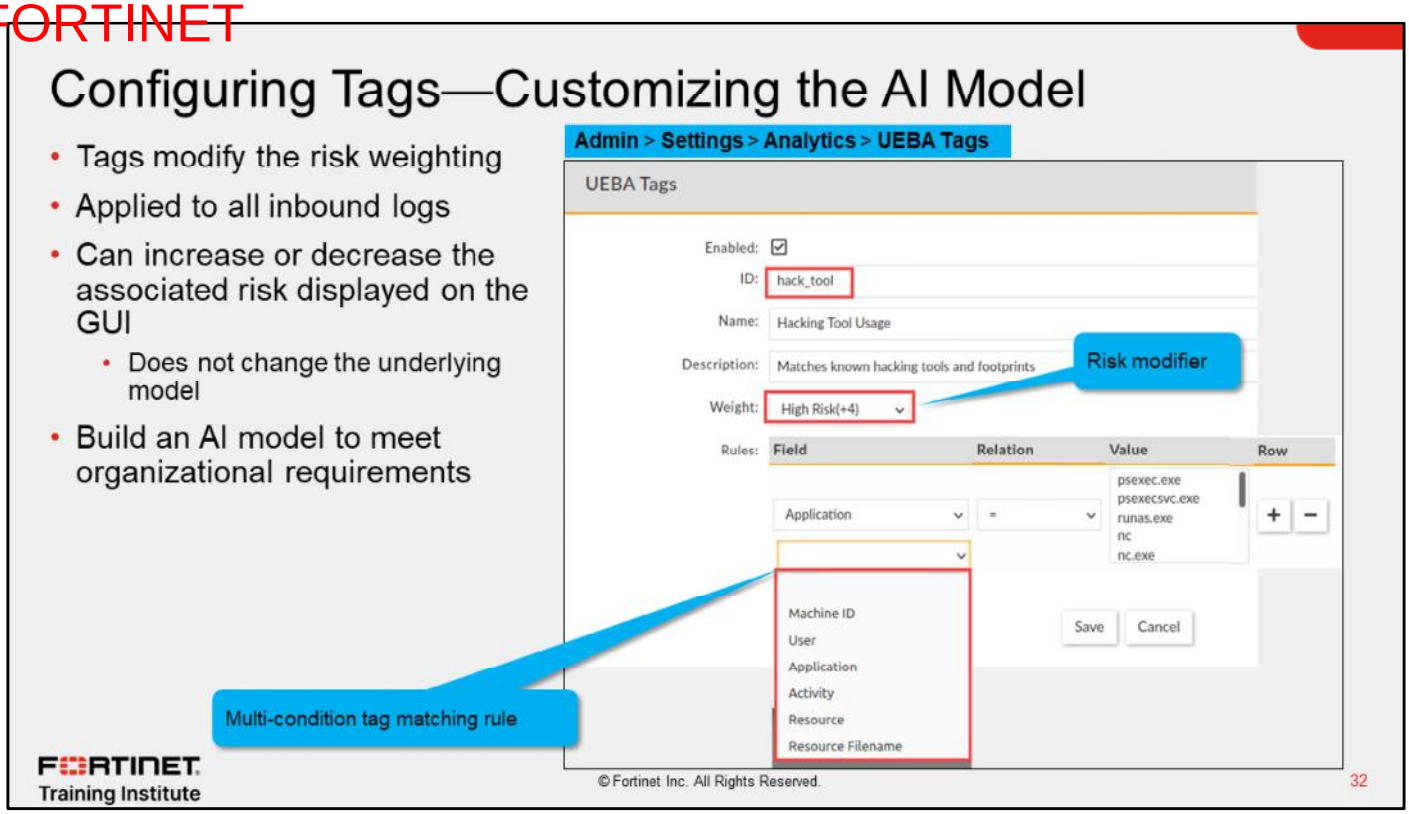
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

31

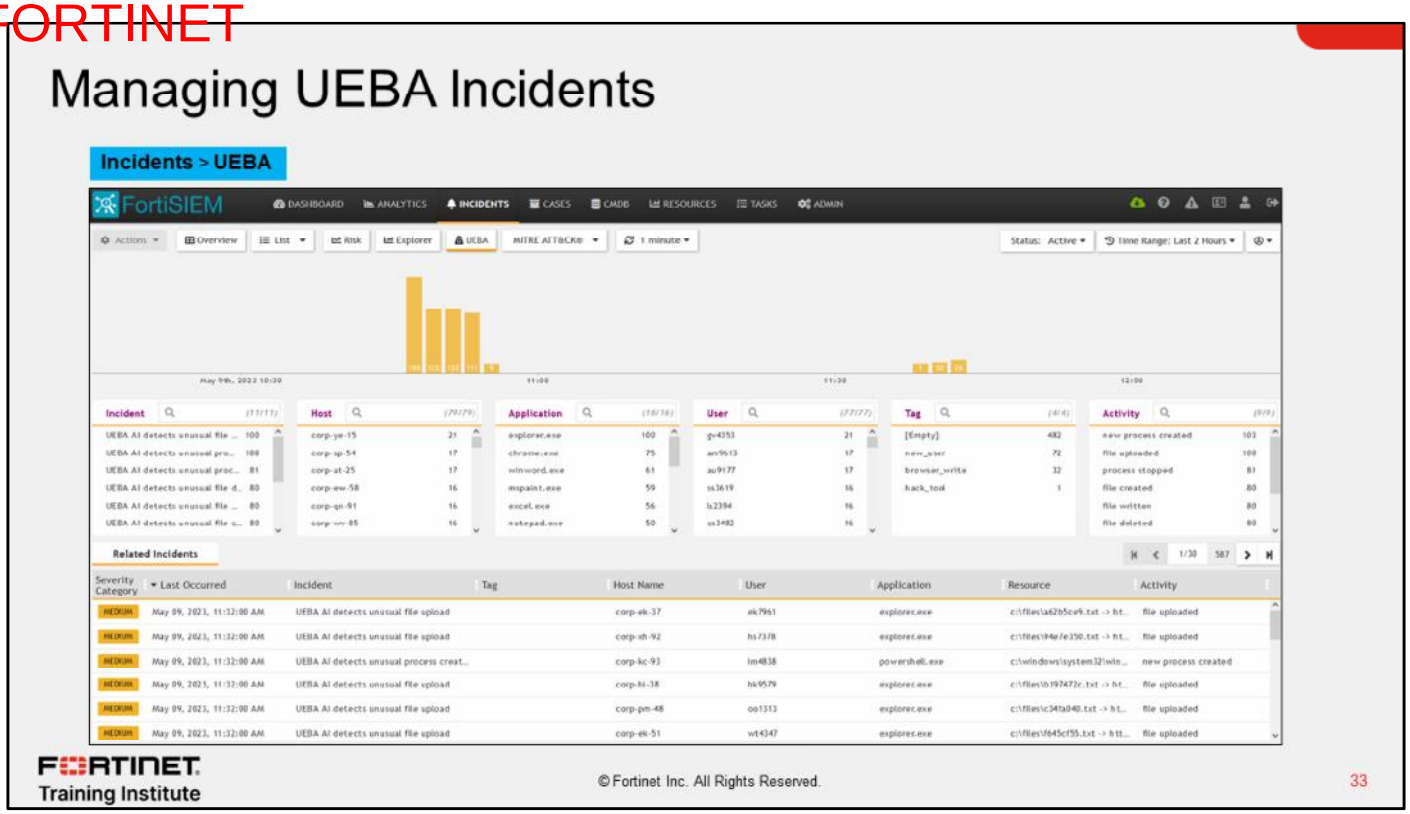
The **UEBA Higher Risk Entities** dialog box contains the following fields. All of the fields are optional. In each field, use the **+** and **-** buttons to add or remove entries.
File Types: Enter the type of file you want to monitor.
File Paths: Enter the path to the folder you want to monitor.
User Accounts: Enter the name of the Windows agent-side user account you want to monitor.
Group Names: Enter the name of the Windows agent-side group you want to monitor.

DO NOT REPRINT
© FORTINET



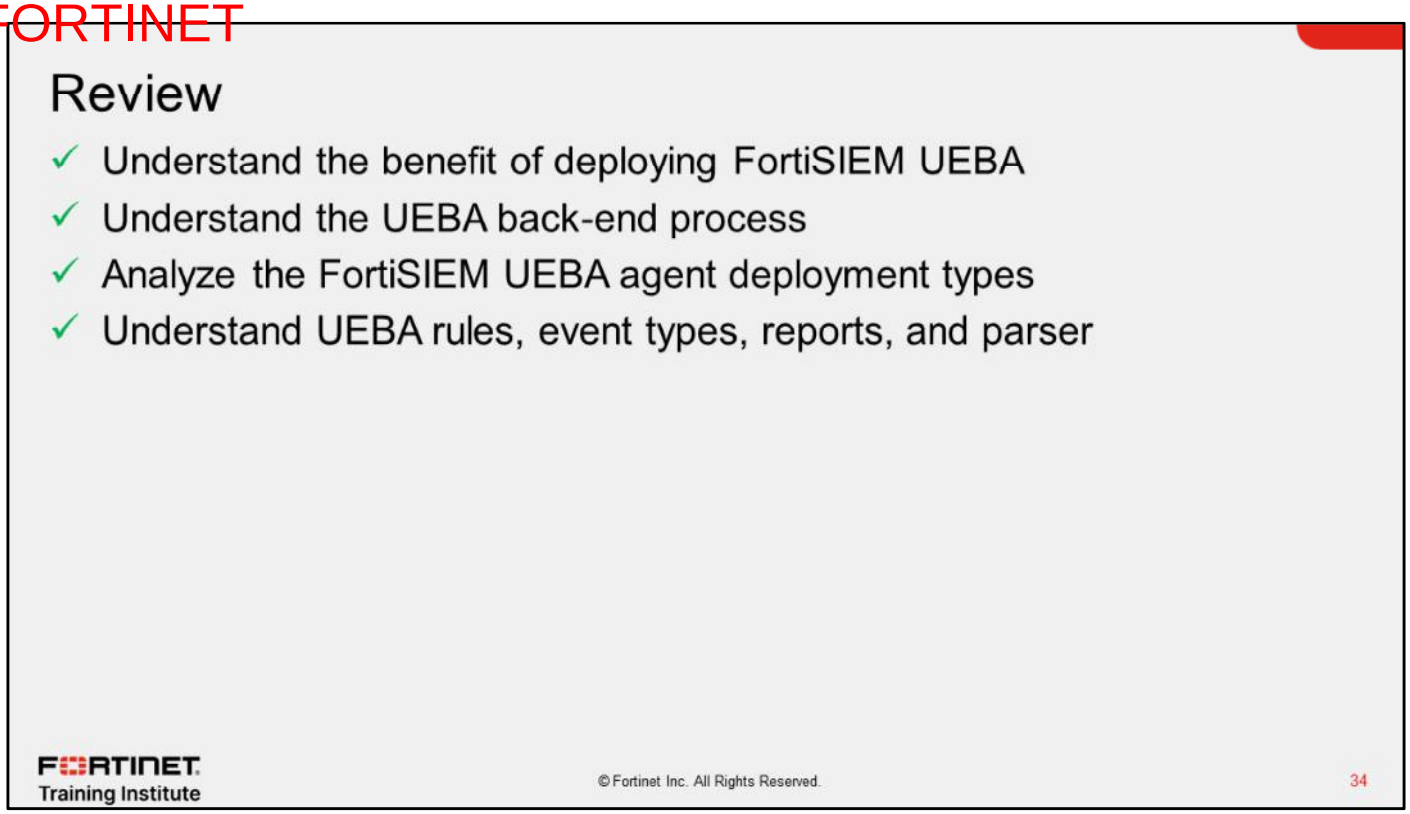
FortiInsight attempts to categorize anomalous events using tags. AI inspects the events for specific characteristics, as defined in the AI tag definitions, and applies the appropriate tags to events that match. Setting tags in FortiSIEM allows you to identify the FortiInsight tags that you want FortiSIEM to monitor. Tags modify risk weighting which is applied to all inbound logs. Tags can increase or decrease the associated risk displayed in the GUI, but they do not change the underlying AI model.

DO NOT REPRINT
© FORTINET



The UEBA incident dashboard shows the UEBA incidents that the AI module creates.
The attribute list table provides the following information about AI alerts received from FortiInsight:
Incident: the name of the incident that was detected.
Host: the host name or IP address where the alert originated.
Application: the name of the application that is the source of the incident.
User: the Windows agent user.
Tag: the tag used to categorize the alert.
Activity: the description of the activity that raised the alert.

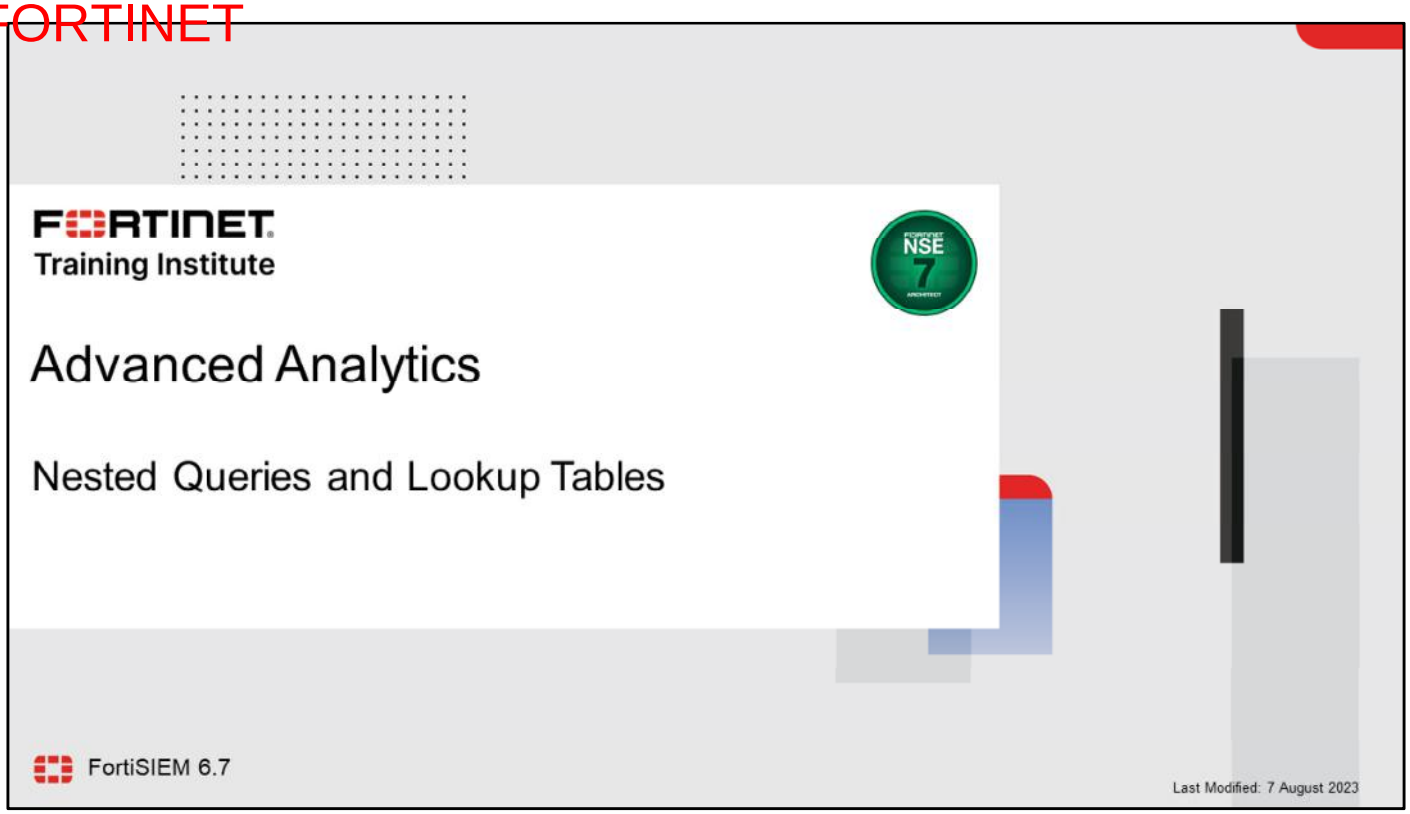
DO NOT REPRINT
© FORTINET



This slide shows the objectives that you covered in this lesson.

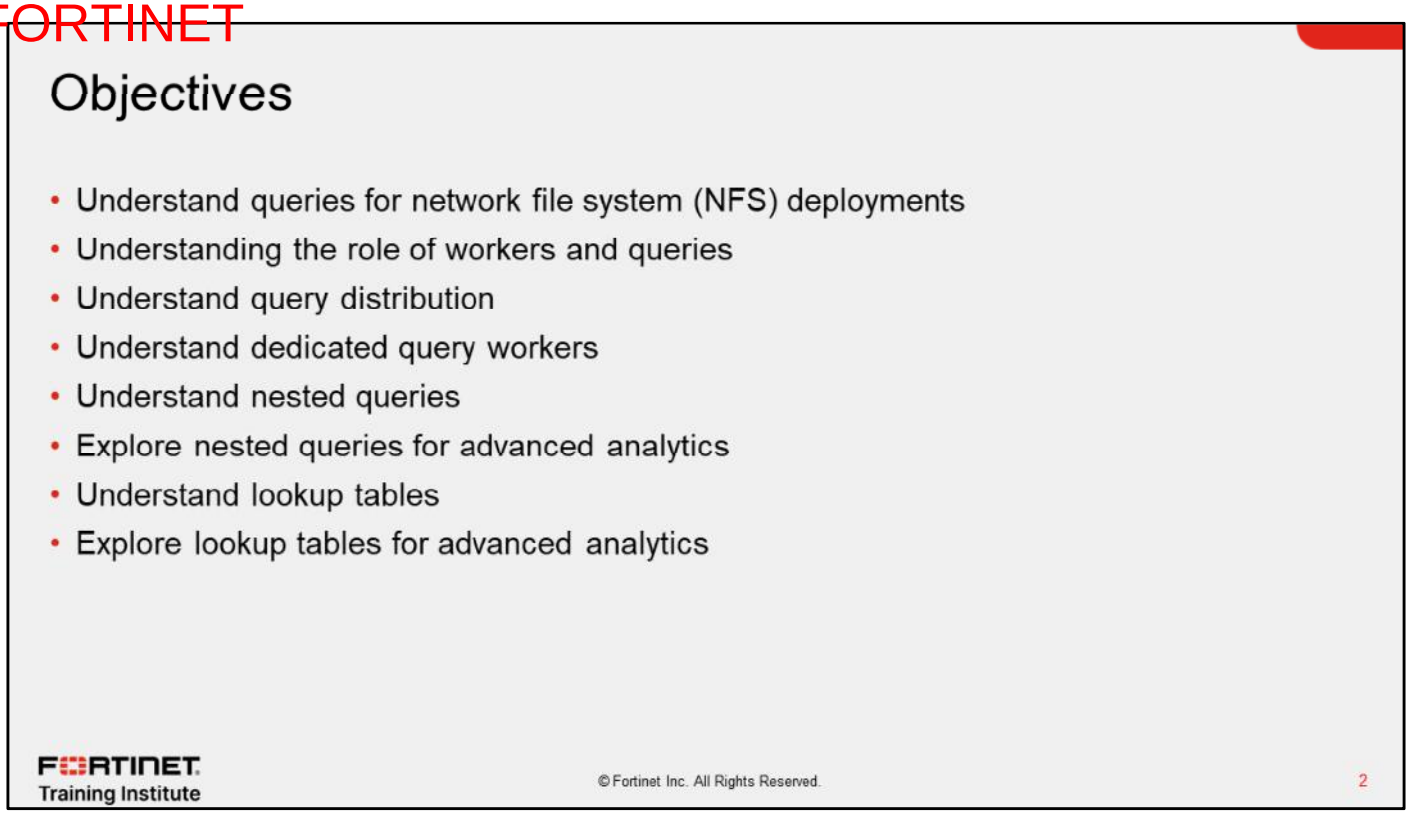
By mastering the objectives covered in this lesson, you learned how to deploy FortiSIEM UEBA in a multi-tenant environment and identify key implementation requirements. You also learned about the UEBA rules, event types, parser, reports, and watchlist, as well as the importance of UEBA tags in determining the risk level of UEBA incidents.

DO NOT REPRINT
© FORTINET



In this lesson, you will learn about the nested queries and lookup tables for advanced analytics on FortiSIEM.

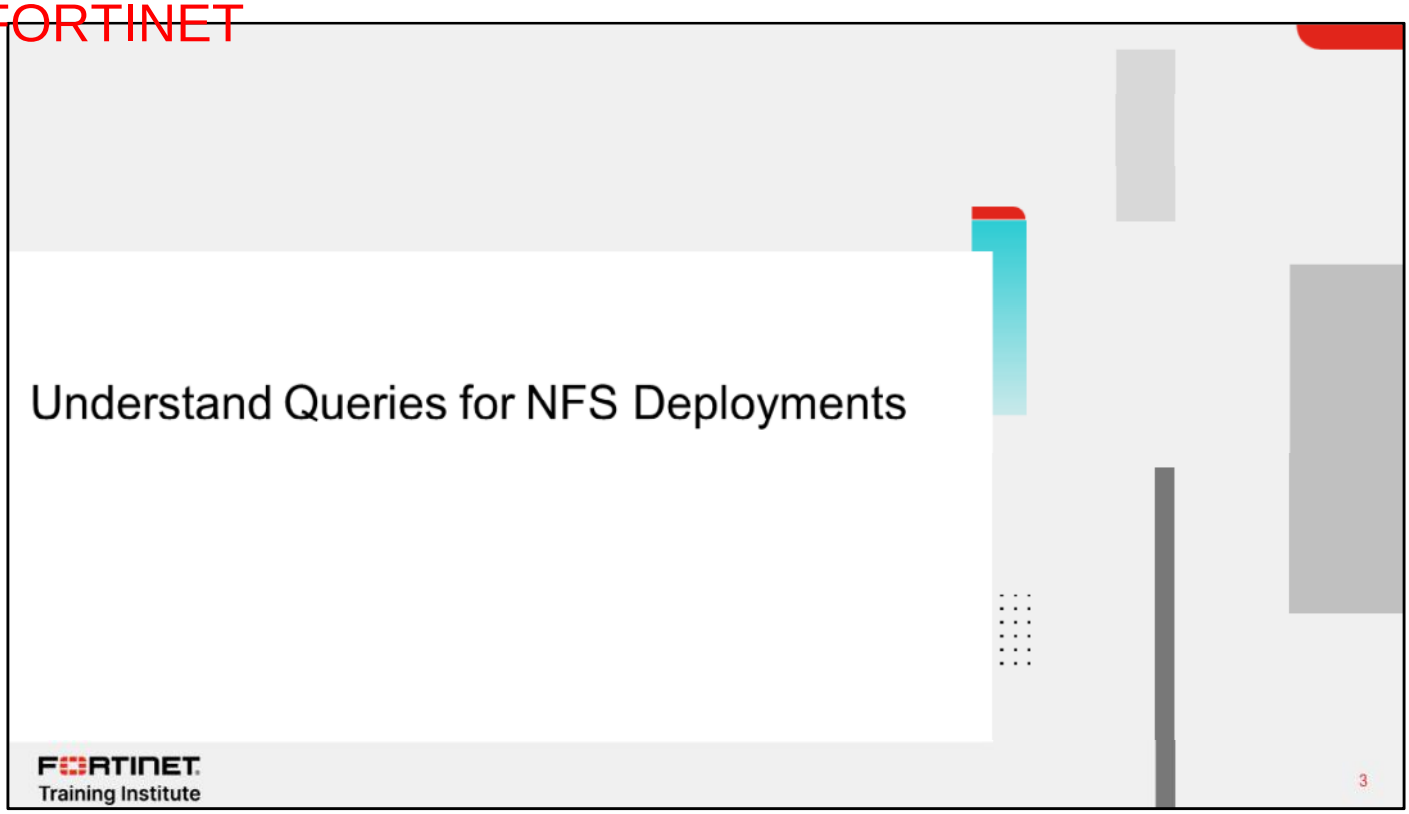
DO NOT REPRINT
© FORTINET



After completing this lesson, you should be able to achieve the objectives shown on this slide.

By demonstrating competent understanding of nested queries and lookup tables, you will be able to use the nested queries and lookup tables for advance analytics.

DO NOT REPRINT
© FORTINET

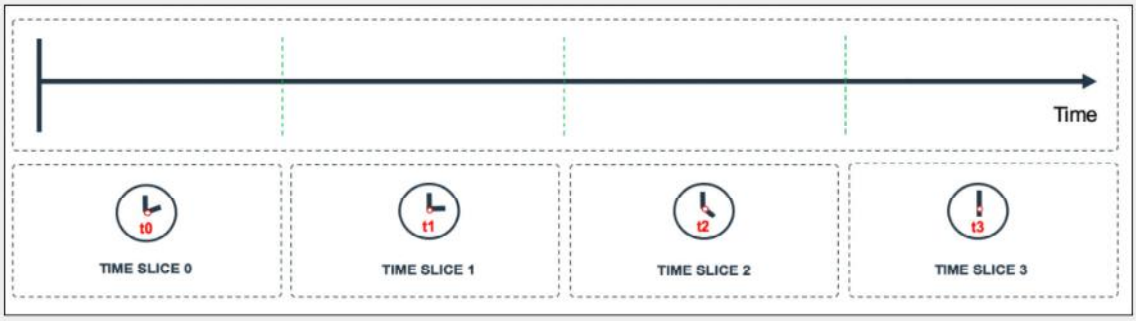


In this section, you will learn about queries for NFS deployments for advanced analytics.

DO NOT REPRINT
© FORTINET

Understanding the Role of Workers and Queries

- FortiSIEM runs two main processes that relate to running analytical queries, `phQueryMaster` and `phQueryWorker`
- `phQueryMaster`
  - Runs on the supervisor node only
- `phQueryWorker`
  - Runs on the supervisor, but also on any worker nodes in the cluster



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

4

FortiSIEM runs two main processes that relate to running analytical queries, `phQueryMaster` and `phQueryWorker`.

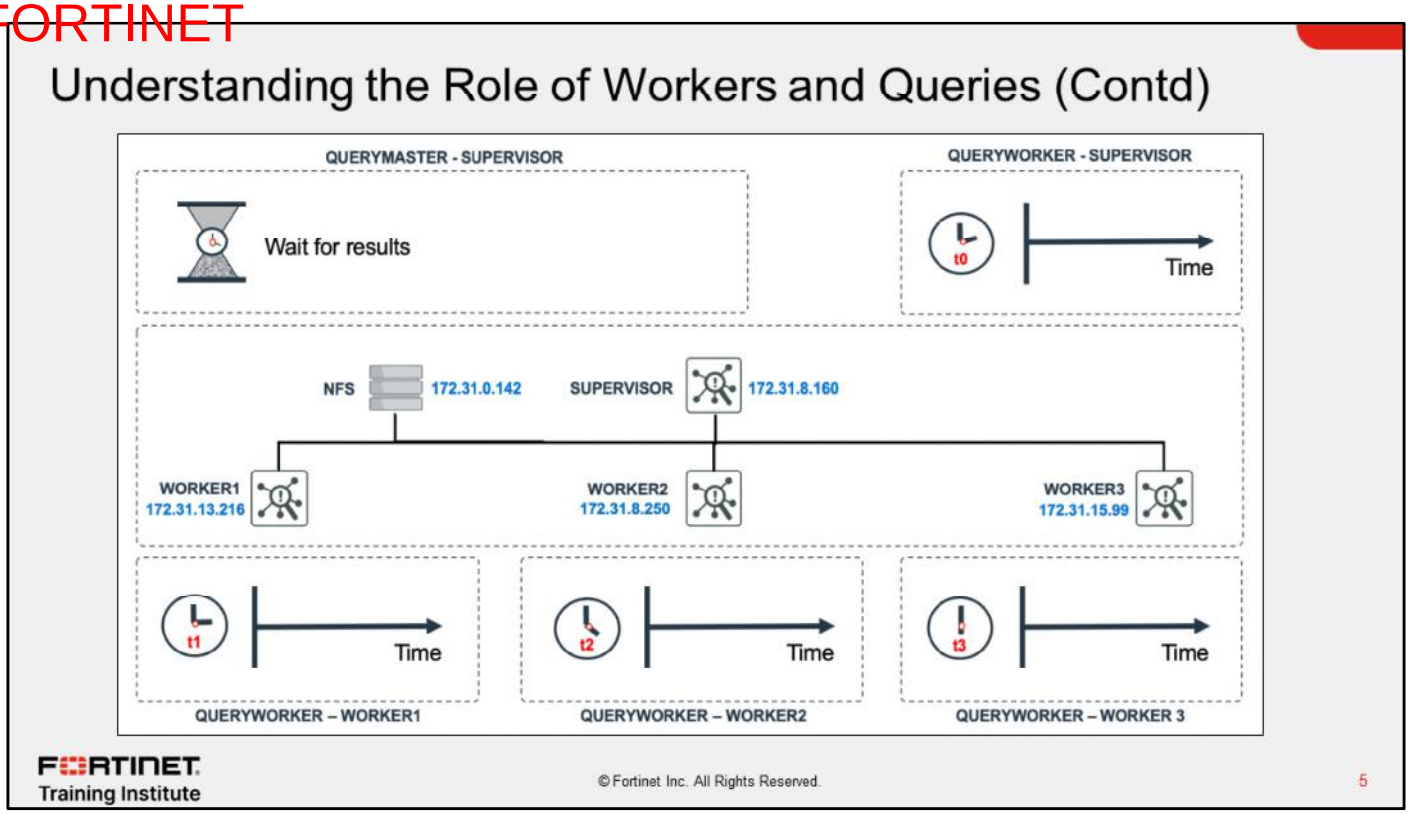
`phQueryMaster` runs on the supervisor node only. This process is responsible for spreading out the work associated with queries to another process named `phQueryWorker`. The process `phQueryWorker` runs on the supervisor, but also on any worker nodes in the cluster.

Consider an analytical search being run on the GUI over a specific time frame. The `phQueryMaster` process receives this query, splits the query into multiple time slices, and assigns these jobs as tasks. By default, the `phQueryMaster` then assigns these tasks to every `phQueryWorker` processes in the FortiSIEM cluster.

Note that you can change this behavior with the use of dedicated query workers.

.

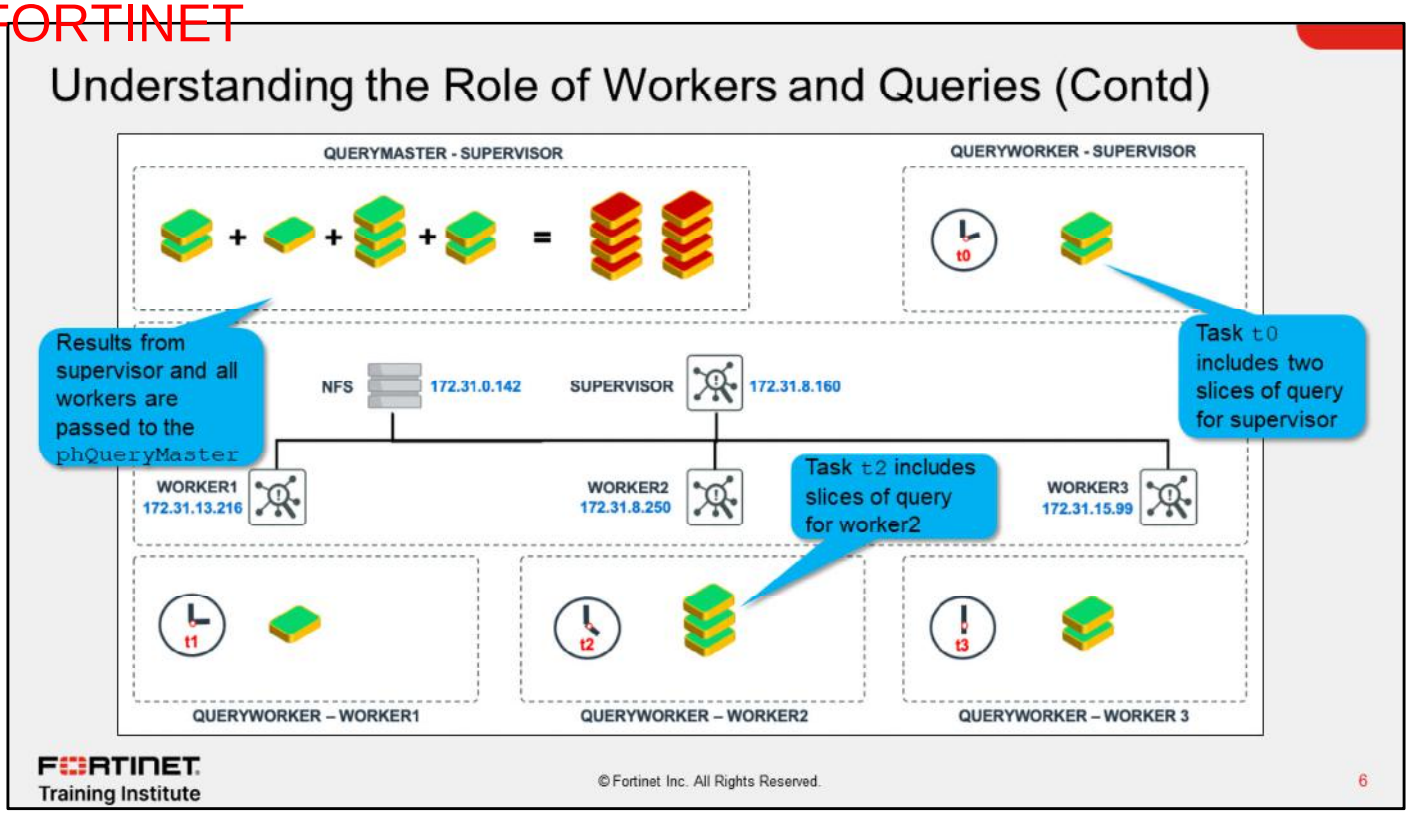
DO NOT REPRINT
© FORTINET



For example, consider a FortiSIEM system with a supervisor and three worker nodes. There are four phQueryWorker processes in total, meaning four tasks, which start at zero. Furthermore, the `phQueryMaster` will then instruct each `phQueryWorker` process (each worker and then itself on the supervisor node, by default) to query the event database for results over that period, which is why the load is spread across the supervisor and all workers.

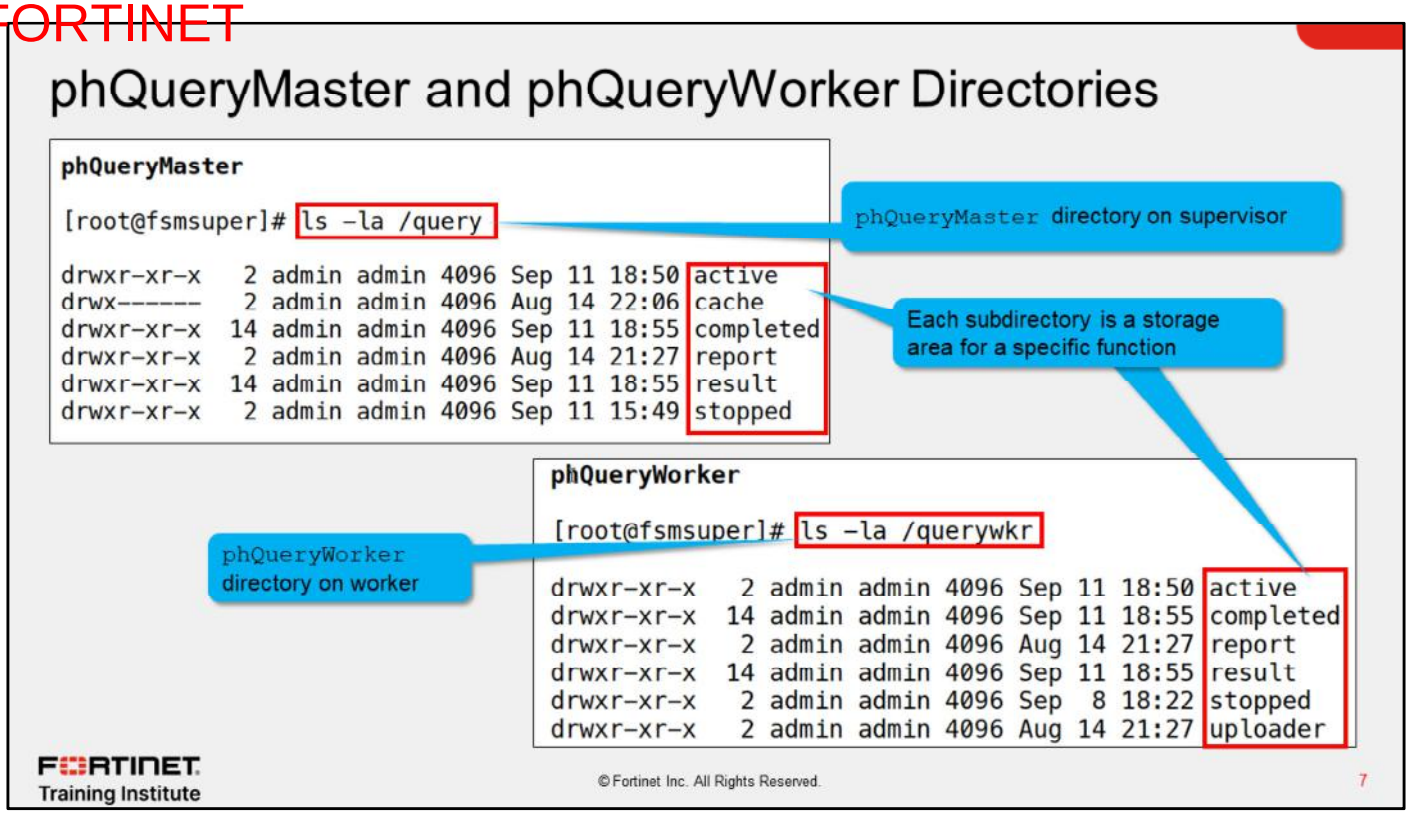
Note that the FortiSIEM version 5.x version behaviour was to query all workers, version 6.x onwards behaviour may not be all workers.

DO NOT REPRINT
© FORTINET



After the results are complete on each query worker, the results are passed to the `phQueryMaster`, which collates and sorts the results. This slide shows complete sequence.

DO NOT REPRINT
© FORTINET



The `phQueryMaster` and `phQueryWorker` directories are both present on the OS disk of the FortiSIEM supervisor and worker nodes, although you use the `/query` only on the supervisor whereas to list query folders on the worker, you need to check the directory `/querywkr`. Each subdirectory is a storage area for a specific function.

DO NOT REPRINT
© FORTINET

phQueryMaster and phQueryWorker Directories (Contd)

- `phQueryMaster` (`/query` on the supervisor)

Directory Name	Purpose
active	Currently running queries
cache	Not currently used.
completed	Completed queries
report	Inline reports
result	Result files for the whole completed query

- Symbolic link in FotiSIEM 6.x
  `/query -> /opt/phoenix/cache/query`

- `phQueryWorker` (`/querywkr` on supervisor or workers

Directory Name	Purpose
active	Currently running queries on this specific query worker.
completed	Completed queries on this query worker.
report	Inline reports
result	Result receipt, that task was completed
stopped	Queries stopped in the GUI, or `phQueryMaster` due to timeout values being met

- Symbolic link in FotiSIEM 6.x
  `/querywkr -> /opt/phoenix/cache/querywkr`

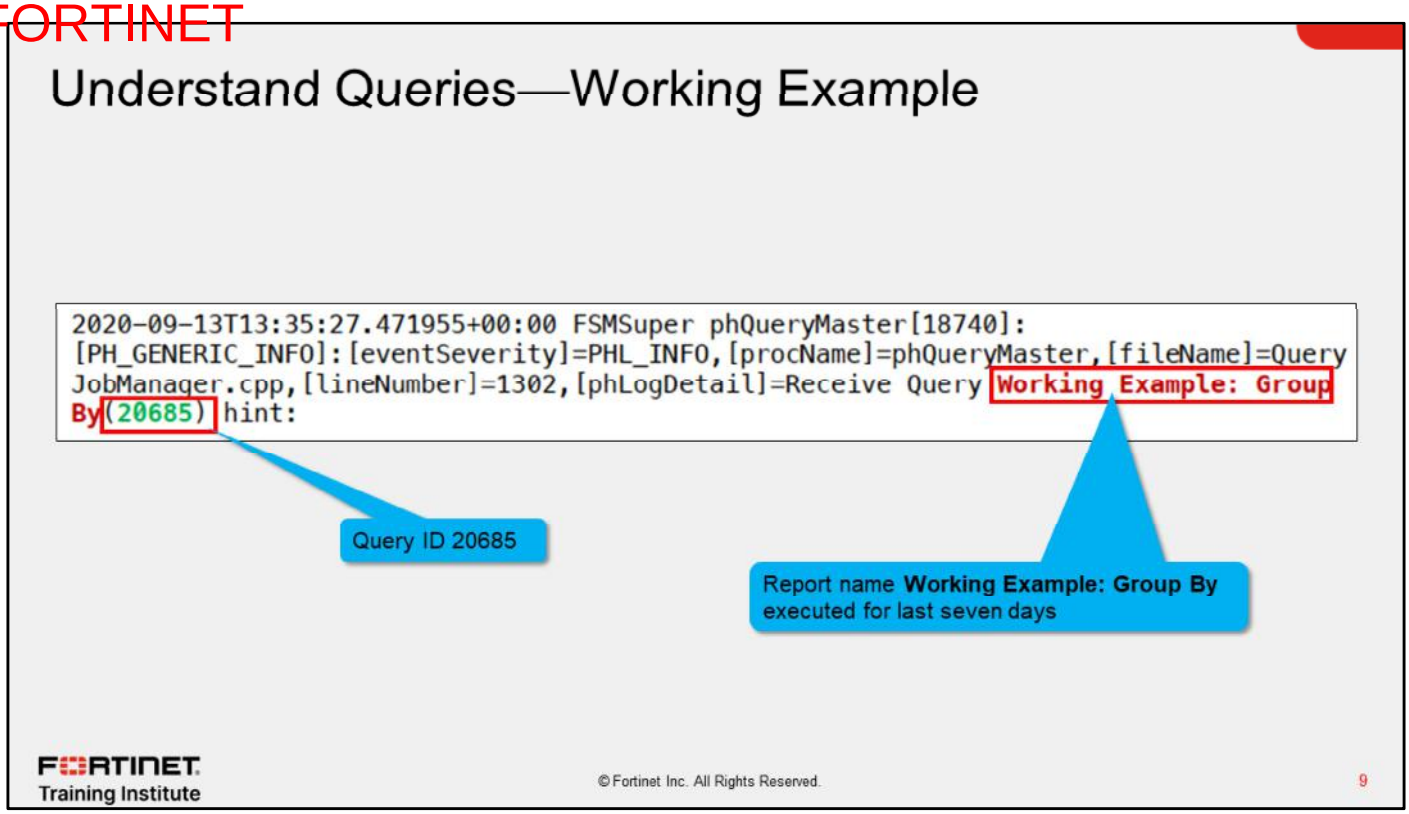
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

8

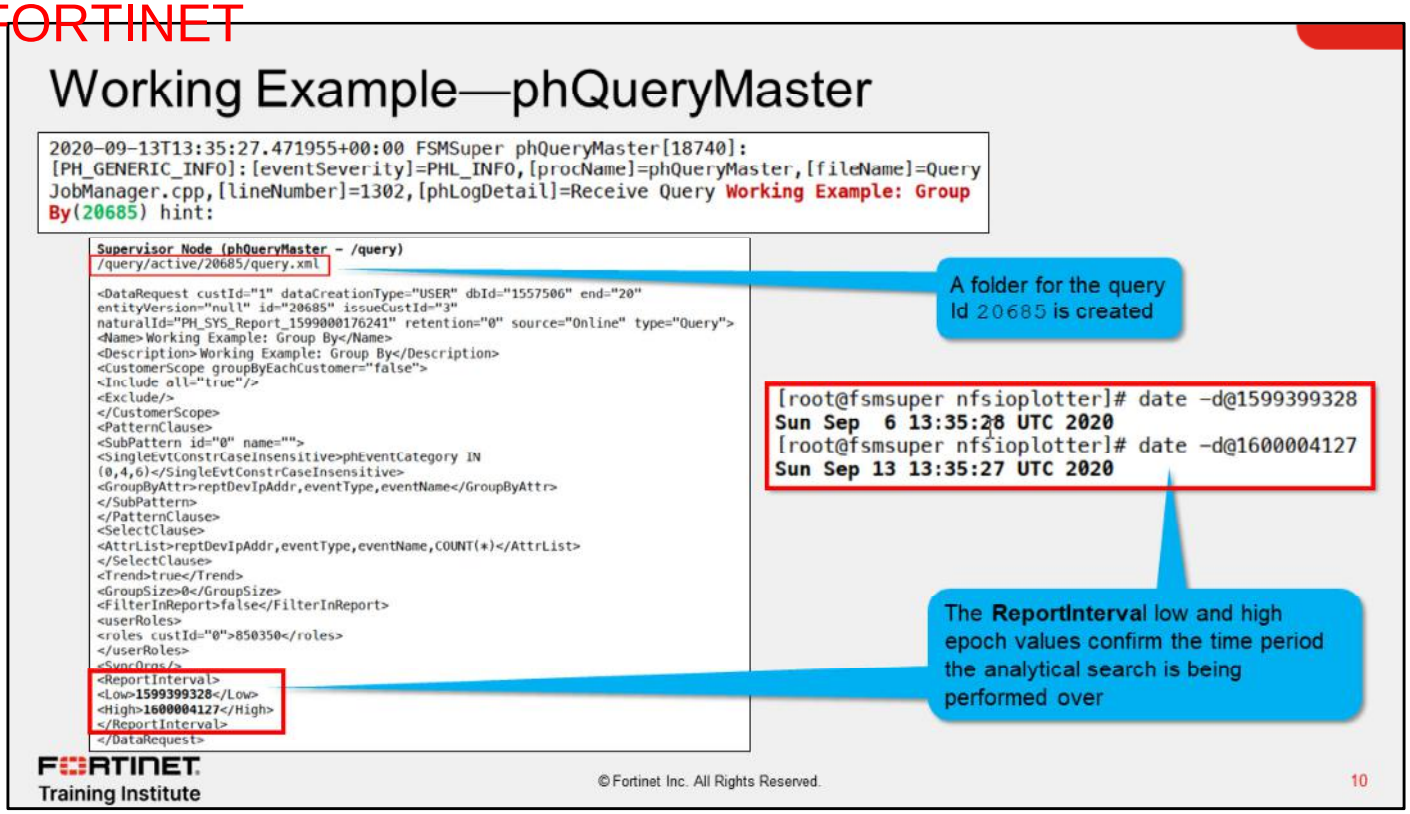
The tables on this slide list the directories for `phQueryMaster /query` and `phQueryWorker /querywkr`. These directories delete or clean up their content approximately every hour..

DO NOT REPRINT
© FORTINET



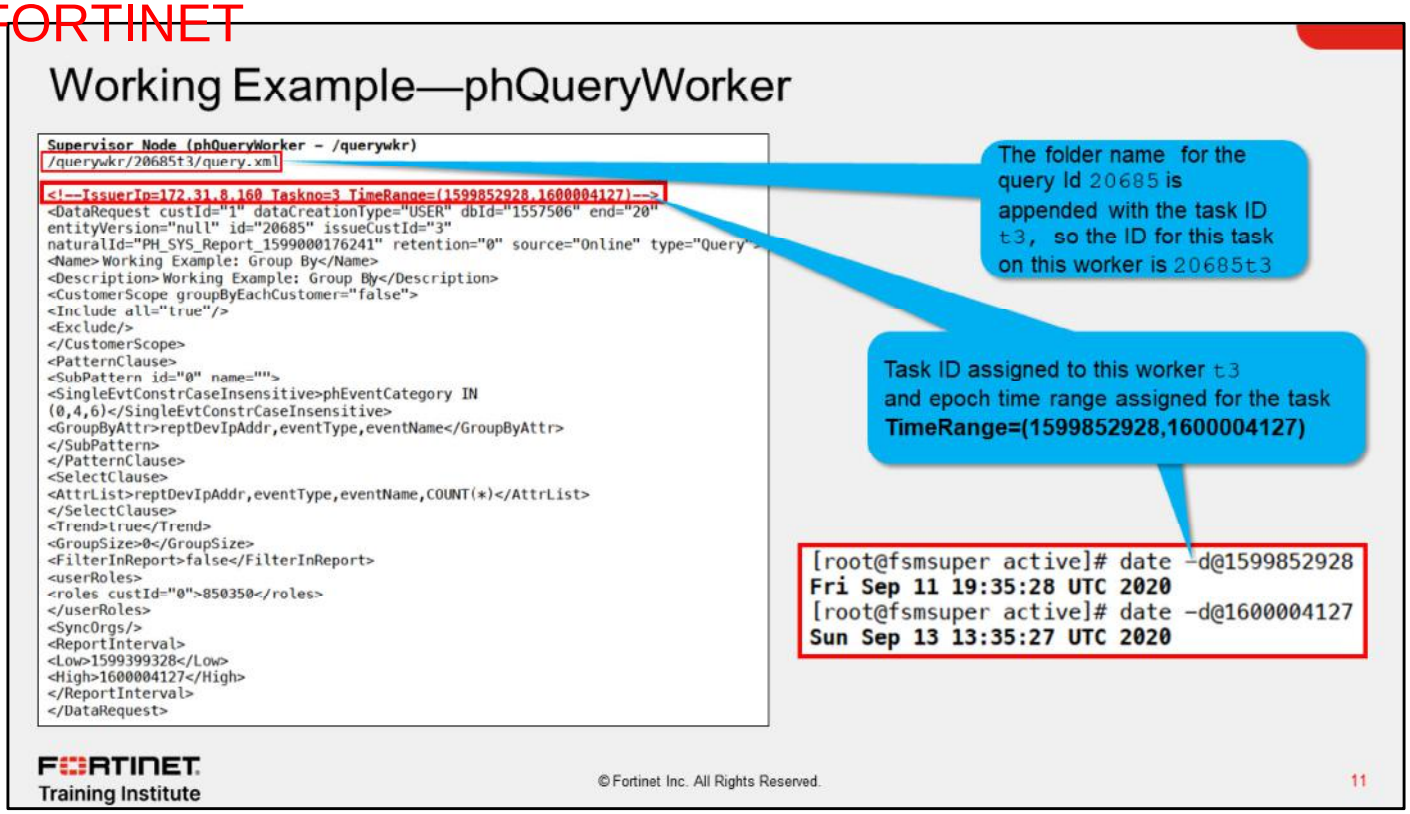
For example, , a report was created called **Working Example: Group By**, which was run over the last seven days. On the supervisor node in `phoenix.log`, the log recorded the entry as shown on the slide, which contained the report name **Working Example: Group By** and query ID `20685`.

DO NOT REPRINT
© FORTINET



On the supervisor under the `/query` directory, the system creates a folder for the query ID `20685`, creates and stores the query XML for the report being run, here. The **ReportInterval Low** and **High** epoch values confirm the time period the analytical search is being performed over.

DO NOT REPRINT
© FORTINET



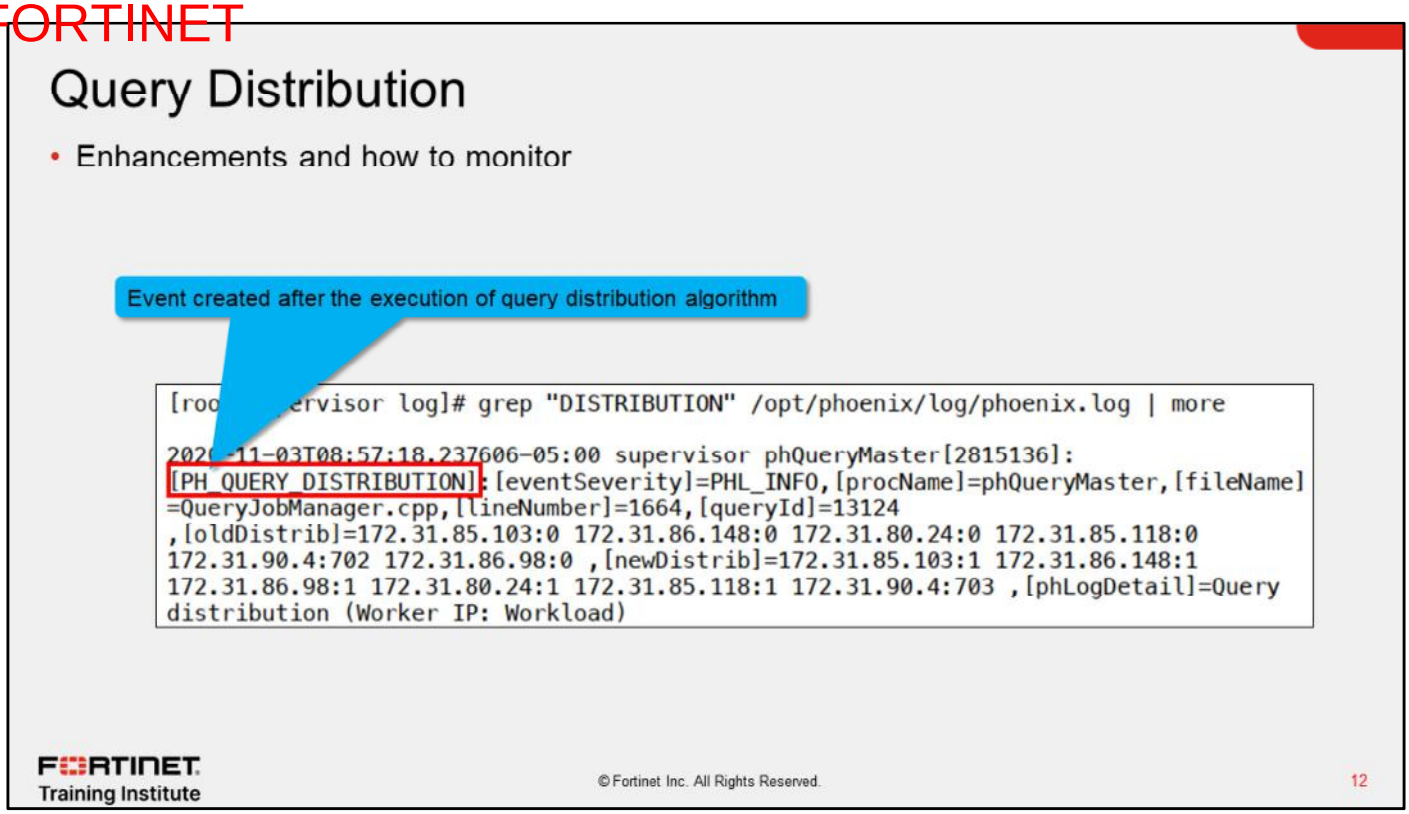
On each query worker node, the `query.xml` appears under the active directory, and query ID folder, but the folder name is appended with the task ID, for example `t3`. Note each query distributes the tasks randomly, so you may find different task ID folders for one query compared to another.

Note that, the `phQueryWorker` on the supervisor for this search has **Taskno=3**.

The first line of the query is appended, with the IP of the query master, the task number, and the time slice that this worker assigned to the query as an epoch timestamp. The full `query.xml` is shown on this slide, but it is the same as in the `phQueryMaster` folder, just with the addition of the time slice and so on.

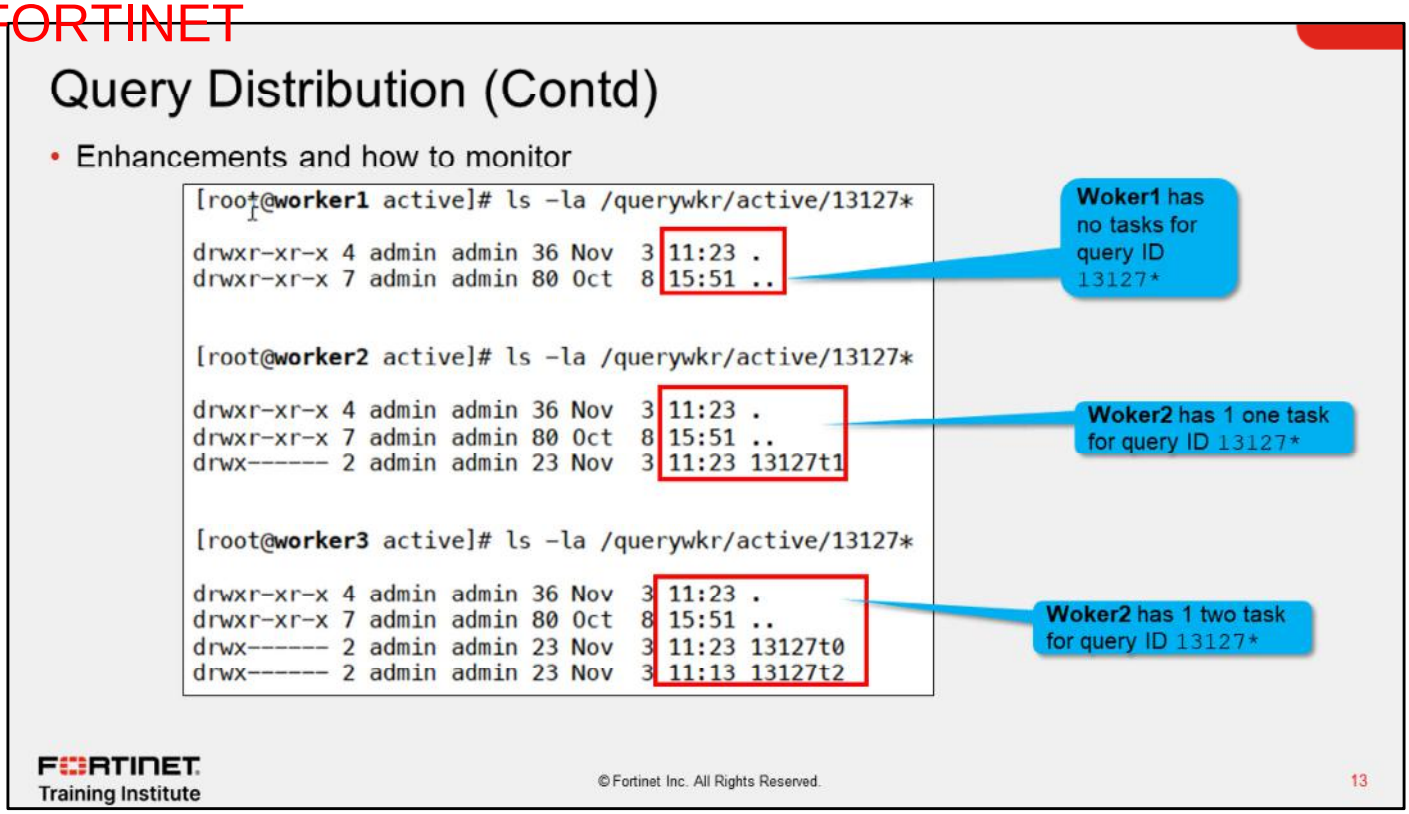
The time slice that `phQueryWorker` on the supervisor node is working on can be calculated through the statement **TimeRange=(1599852928,1600004127**).

DO NOT REPRINT
© FORTINET



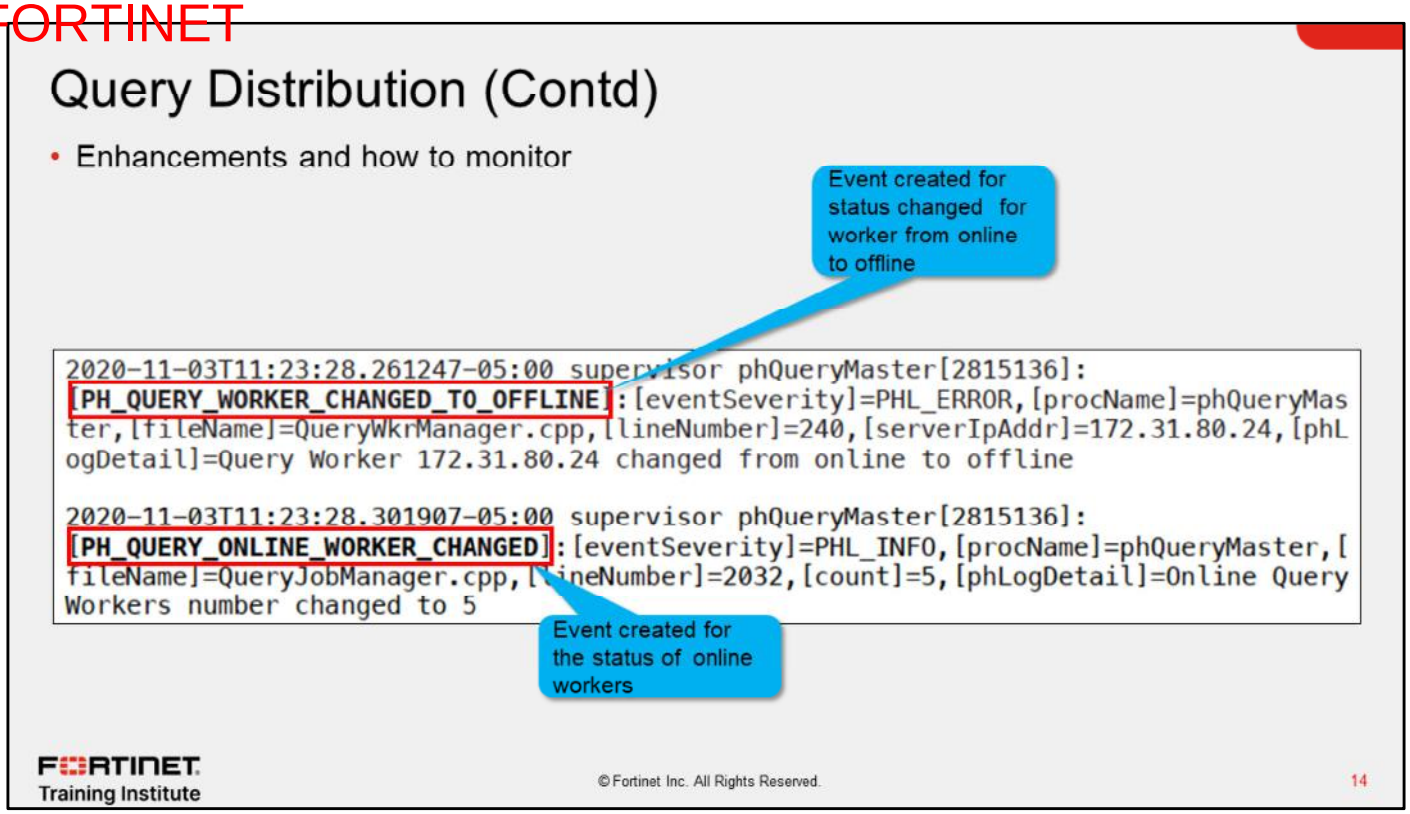
When a query is run, a distribution algorithm runs in the back end to identify which query workers will receive the query to run. It looks at the current active load on each and updates the distribution for the new query.

DO NOT REPRINT
© FORTINET



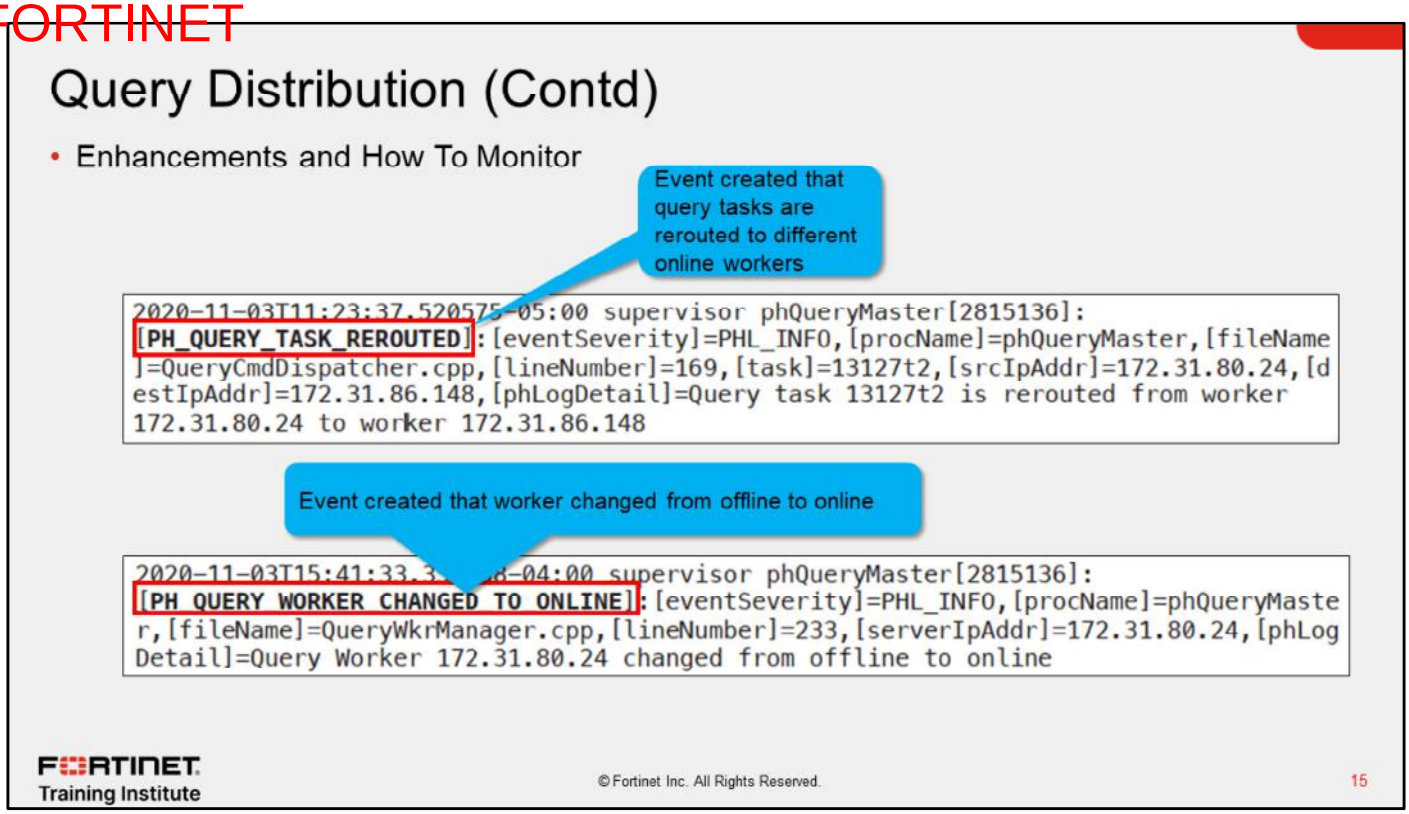
Some query workers may receive multiple query tasks by the distribution algorithm, while other query workers receive none. For example, for the query shown on this slide, narrowing the focus for query ID `13127`, the query worker on **Worker1** has no tasks for this query, whereas **Worker2** has one task, and **Worker3** has two tasks.

DO NOT REPRINT
© FORTINET



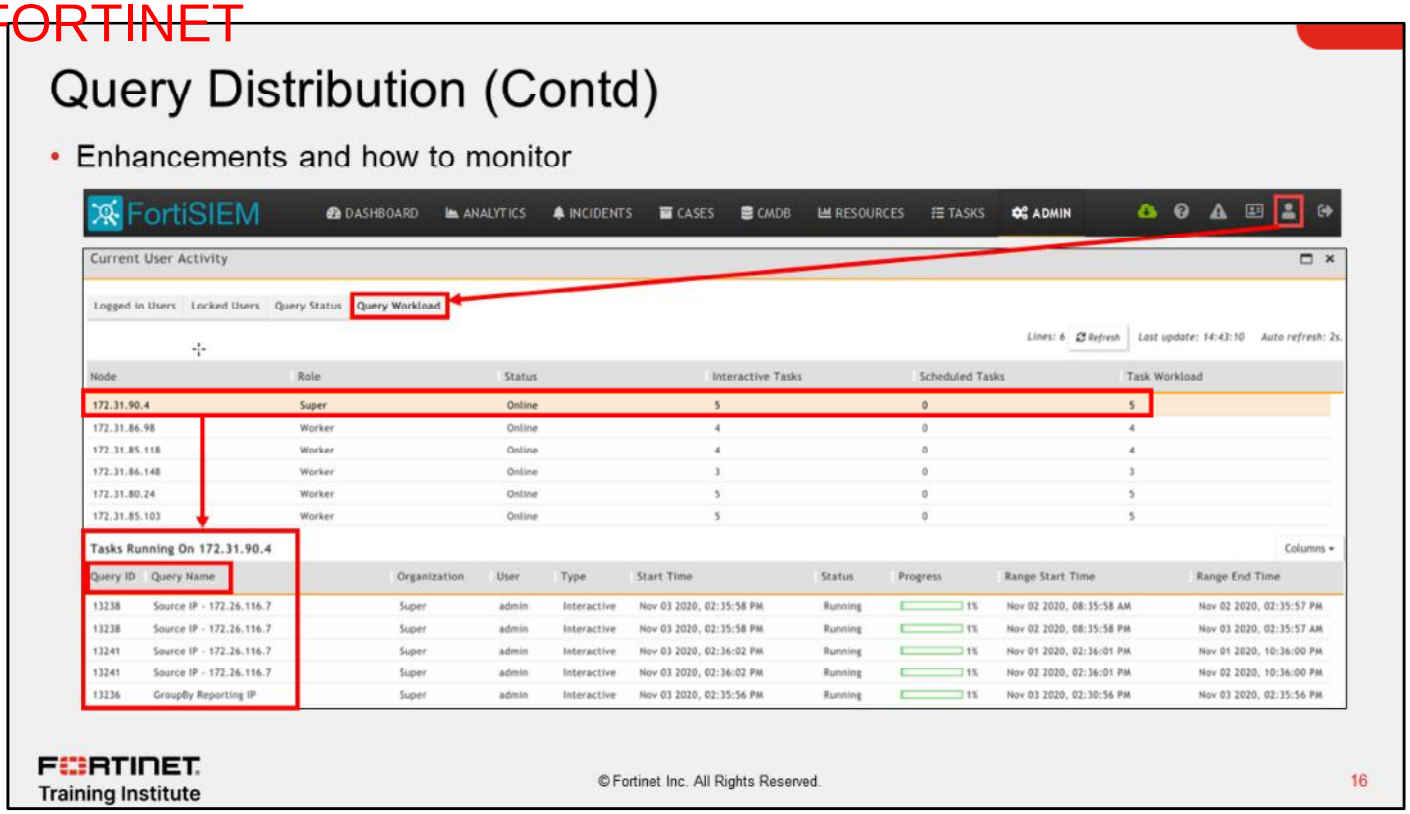
You may see multiple query tasks running on the same query worker, which when a query worker goes offline. This would produce back-end log entries in `/opt/phoenix/log/phoenix.log` as this slide shows.

DO NOT REPRINT
© FORTINET



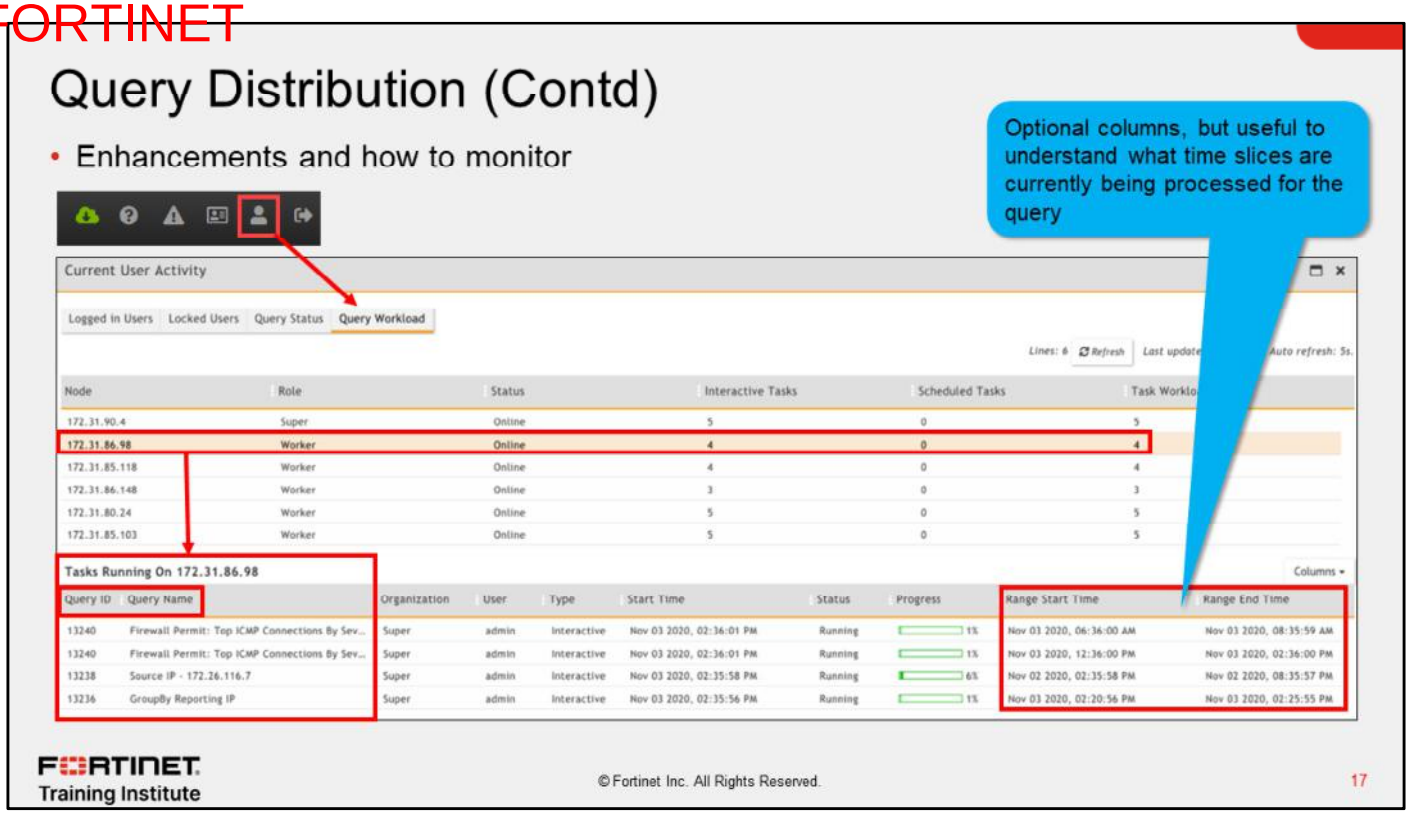
Any queries running on the query worker, which, unusually, has gone offline, would be automatically rerouted to another available node. This entry would then be logged on the back-end. Once the query worker is restored to normal operation, the log entries shown on this slide would be created.

DO NOT REPRINT
© FORTINET



Queries views and behaviors are available on FortiSIEM under the **Current User Activity** icon. For example, this slide shows that the query worker on the supervisor node is currently running five query tasks, where it has two tasks for query ID `13238` and `13241` and one for query ID `13236`.

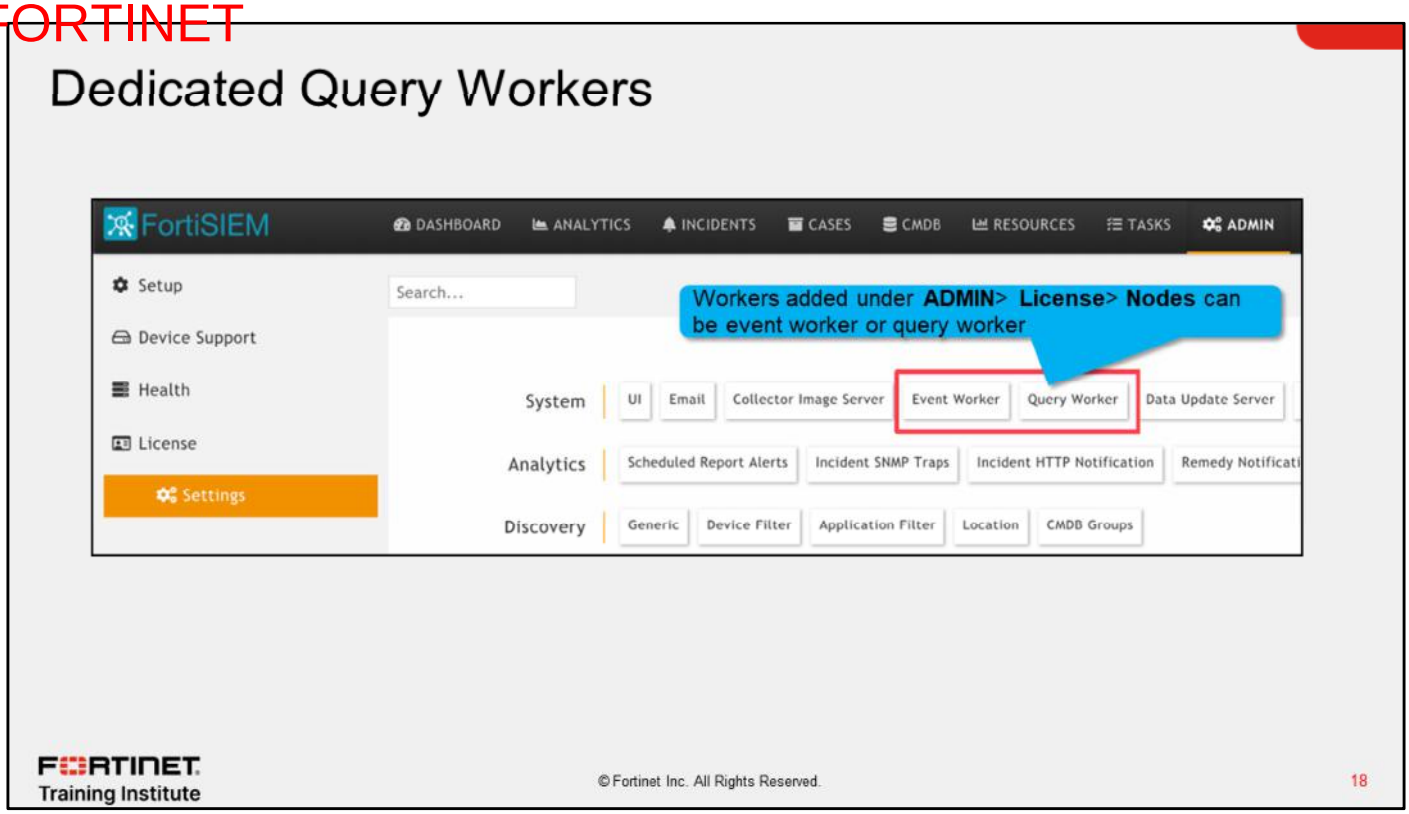
DO NOT REPRINT
© FORTINET



Whereas reviewing the workload on the first worker in the list shows two tasks for query ID `13240`, and one task for query IDs `13238` and `13236`, but no tasks for query ID `13238` and `13241` compared to the supervisor node tasks.

The **Range Start Time** and **Range End Time** are optional columns, but useful to understand what time slices are currently being processed for the query.

DO NOT REPRINT
© FORTINET

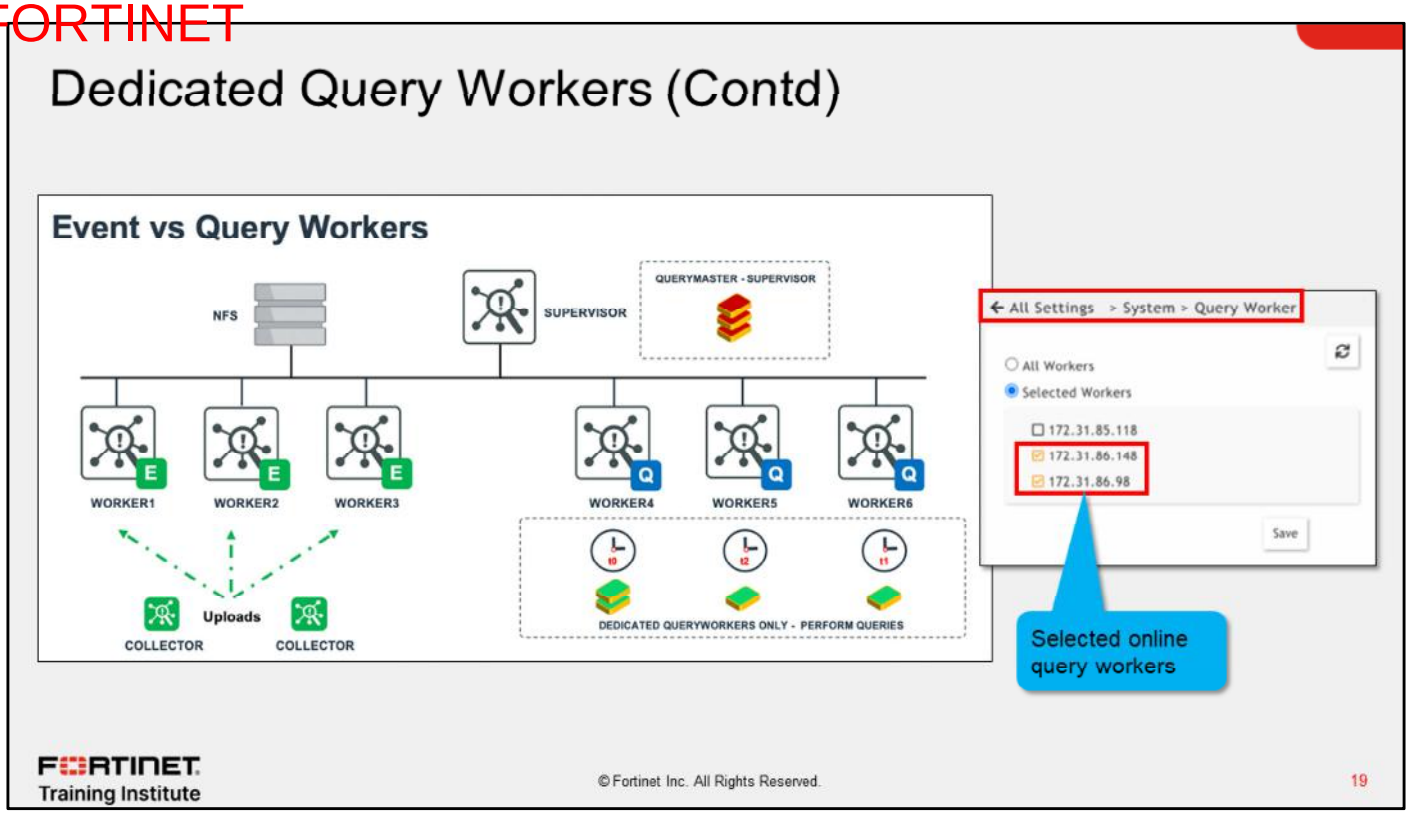


You can designate any workers added to FortiSIEM as an **Event Worker** or dedicated **Query Worker**.

By design, when you add workers, to the system under **ADMIN**> **License**> **Nodes**, they are assigned a license and become part of the cluster.

Any workers added to the system under **Event Workers** are assigned to any collectors for uploading and processing of customer data. Additionally, rule summarization is performed on that worker for incoming data.

DO NOT REPRINT
© FORTINET



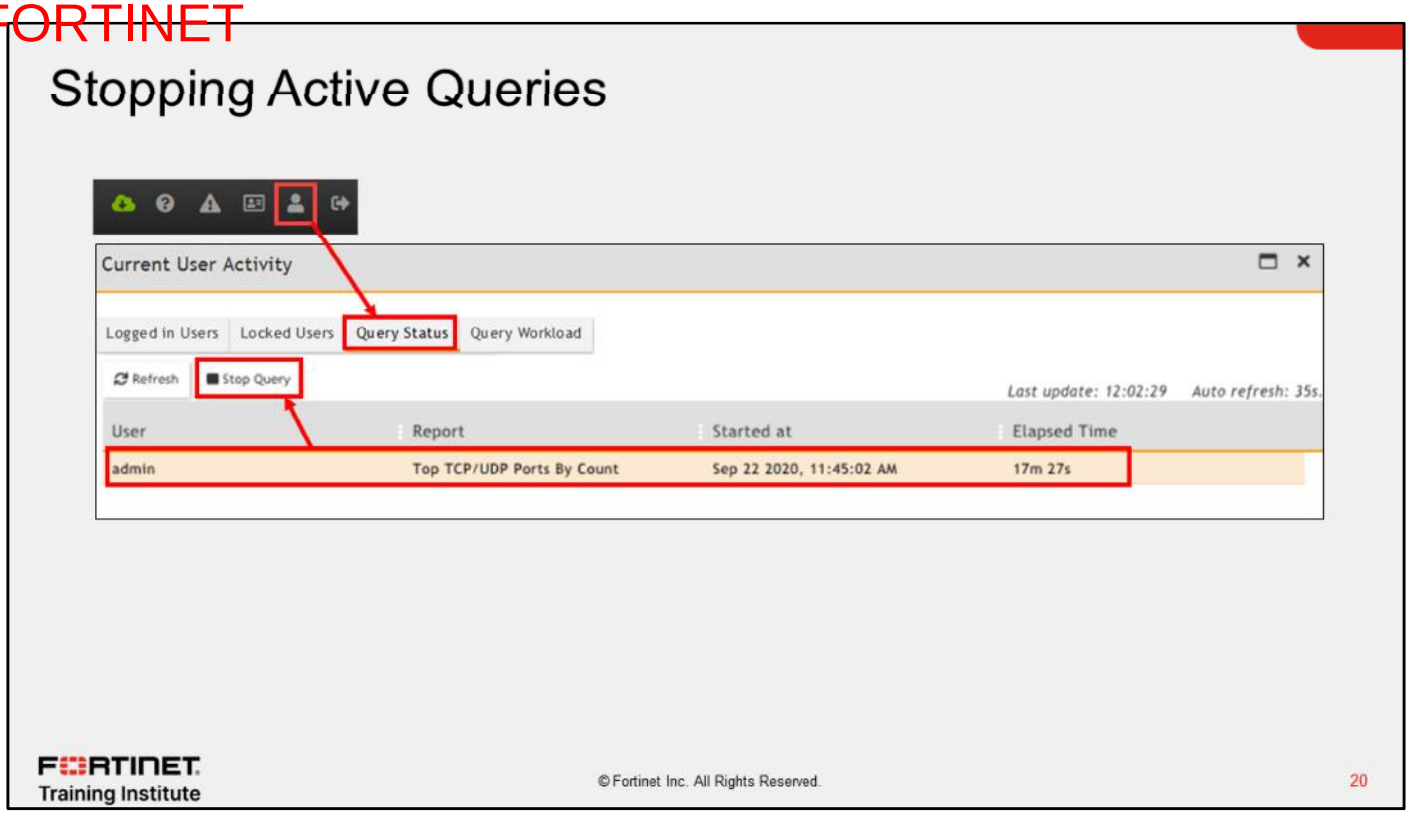
As previously discussed, when either the system or a user performs a query or runs a report, whether unscheduled or scheduled, every worker is involved by default.

If workers are under heavy load processing incoming events from collectors, it may be desirable for those workers not be involved in running queries or add new workers with the sole task of performing queries only. To this aim, the **Query Worker** feature was added in FortiSIEM.

You can choose either all workers by default, or specific workers, to be query workers, meaning only a specific list of workers will performs queries, reports, and scheduled reports.

In this context, worker here means any worker that is listed under **ADMIN > License > Nodes.** Note that a worker can be in the **Query Worker** list and be an **Event Worker** at the same time. Any change you make under worker settings is not instant, It takes a couple of minutes to take effect.

DO NOT REPRINT
© FORTINET



To stop any active user query or report, click the **Current User Activity** icon, and then the> **Query Status** tab.

DO NOT REPRINT
© FORTINET

Listing Active Queries From the Back End

- Listing running queries with name on the supervisor

```
[root@supervisor ~]# cat /query/active/*/query.xml | grep "Name>" | cat -n
     1          <Name>Scale Test B: GroupBy Reporting IP</Name>
     2          <Name>Scale Test B: GroupBy Reporting IP</Name>
     3          <Name>Scale Test API3 GroupBy Reporting IP - 2 Hours</Name>
     4          <Name>Firewall Permit: Top Internal Destinations, Ports By Connections,
Bytes</Name>
```

- Listing running queries on any query worker with name and query ID

```
[root@supervisor ~]# grep -H "Name>" /querywkr/active/*/query.xml | cat -n
     1          /query/active/24994/query.xml:<Name>Scale Test B: GroupBy Reporting IP</Name>
     2          /query/active/25059/query.xml:<Name>Scale Test B: GroupBy Reporting IP</Name>
     3          /query/active/25503/query.xml:<Name>Scale Test API3 GroupBy Reporting IP - 2
Hours</Name>
     4          /query/active/25512/query.xml:<Name>GroupBy Reporting IP</Name>
     5          /query/active/25516/query.xml:<Name>Specific Event Type, Reporting IP and Source
IP - 172.16.20.146</Name>
     6          /query/active/25518/query.xml:<Name>Connection ID</Name>
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

21

Similar to what this slide shows, for reference, it is possible to list running queries from the back end on the supervisor and on any query worker.

DO NOT REPRINT
© FORTINET

Hitting FortiSIEM Limits on Active Queries

Reporting IP	Reporting Device	Event Type	Event Name	Raw Event Log
172.31.15.99	Worker3	PH_GENERIC_INFO	PH system generic info	<14>Sep 23 21:10:28 Worker3 phQueryWorker[31820]: [PH_GENERIC_INFO]:[eventSeverity]=PHL_INFO, [procName]=phQueryWorker,[fileName]=EventSearcher.cpp, [lineNumber]=89,[phLogDetail]=Too many event search threads, will not spawn further
172.31.7.193	worker4	PH_GENERIC_INFO	PH system generic info	<14>Sep 23 21:10:28 worker4 phQueryWorker[16723]: [PH_GENERIC_INFO]:[eventSeverity]=PHL_INFO, [procName]=phQueryWorker,[fileName]=EventSearcher.cpp, [lineNumber]=89,[phLogDetail]=Too many event search threads, will not spawn further

```
2020-09-23T20:13:49.824178+00:00 FSMSuper phQueryWorker[18701]:
[PH_GENERIC_INFO]:[eventSeverity]=PHL_INFO,[procName]=phQueryWorker,[fileName]=Event
Searcher.cpp,[lineNumber]=89,[phLogDetail]=Too many event search threads, will not
spawn further
```

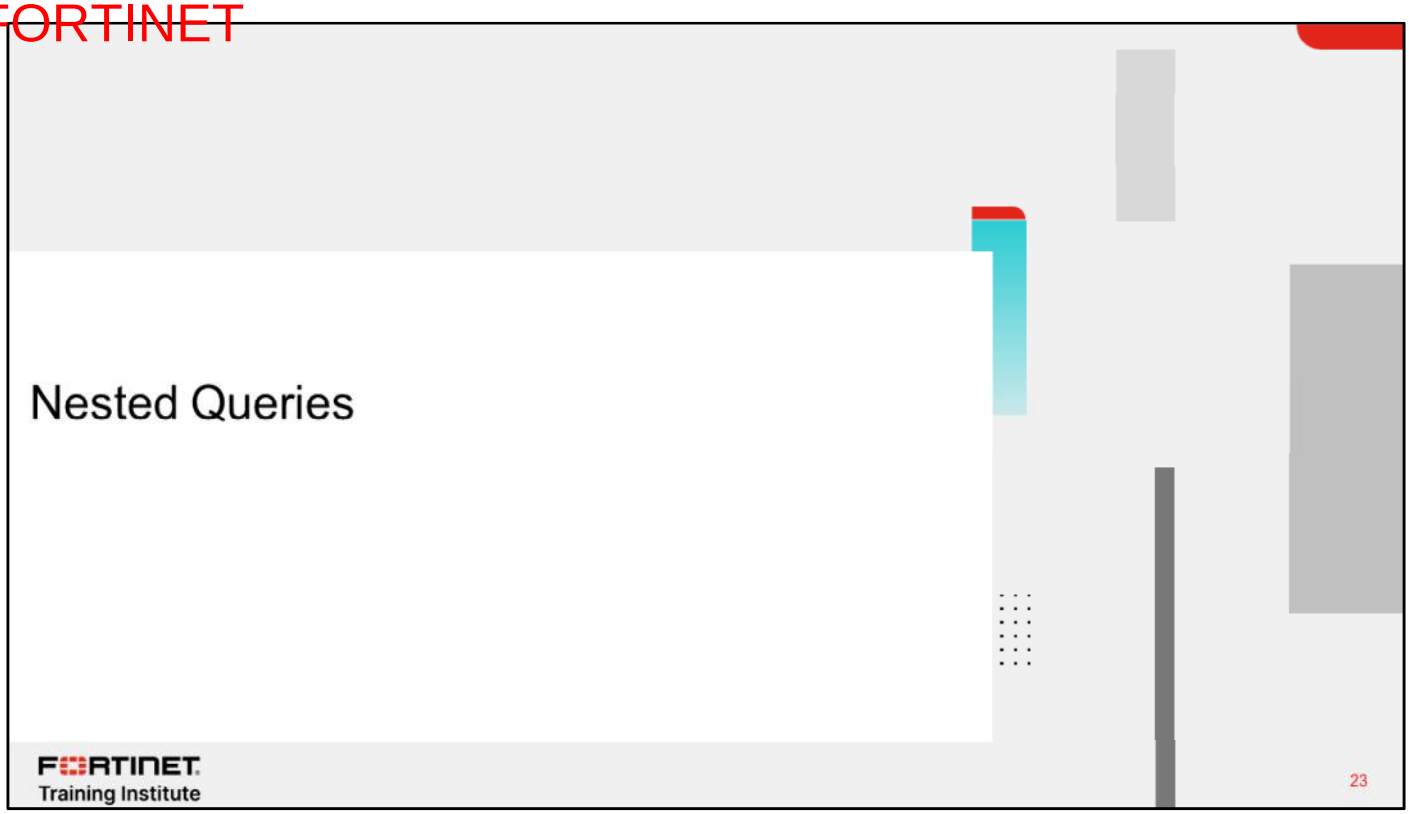
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

22

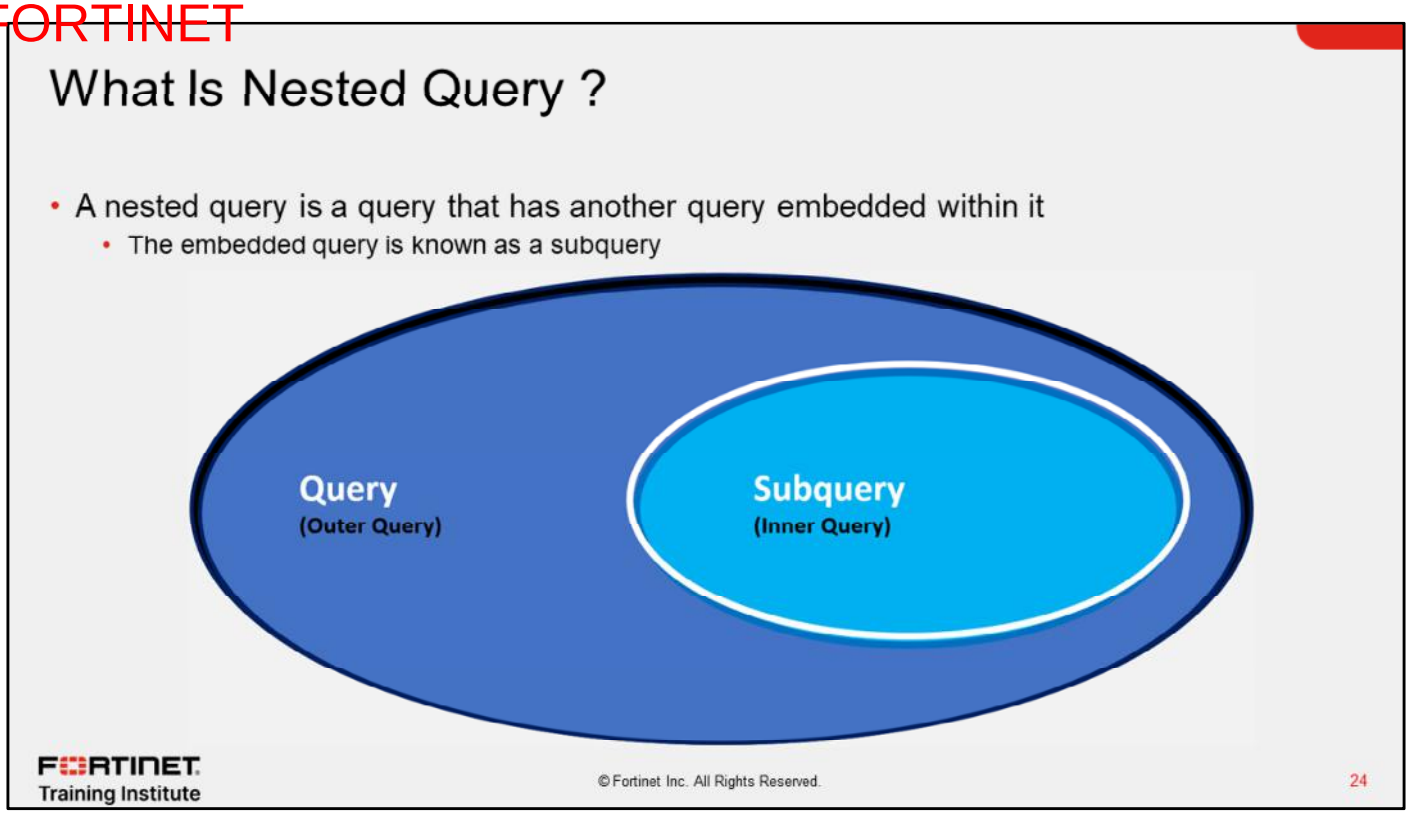
Each query worker has a limit on the number of active queries that it can perform. Watch for this error message in a large scale environment.

DO NOT REPRINT
© FORTINET



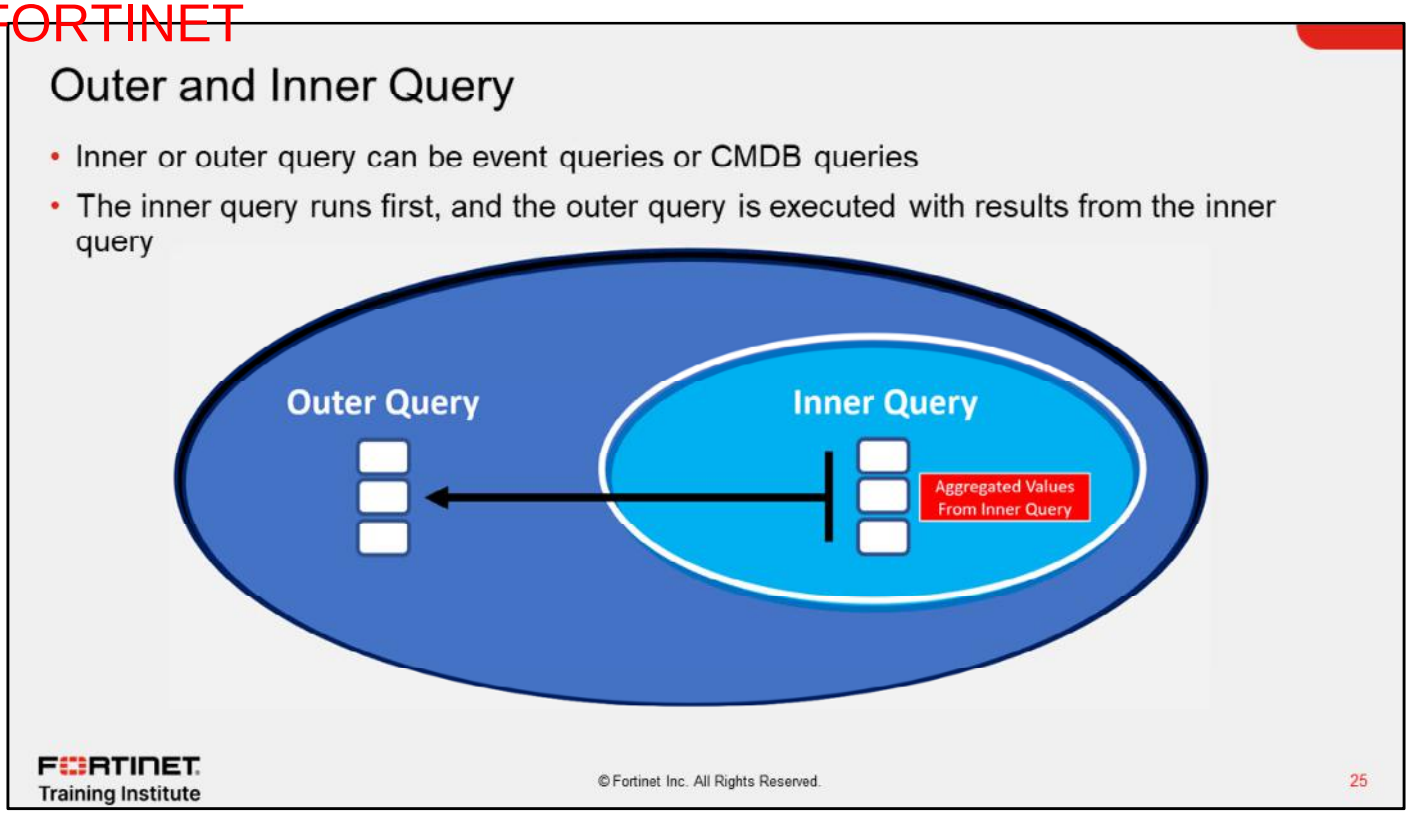
In this section, you will learn about nested queries for advanced analytics.

DO NOT REPRINT
© FORTINET



A nested query is a query that has another query embedded within it. The embedded query is known as a subquery. When a query is included inside another query, the outer query is known as the main query, and the inner query is known as a subquery.

DO NOT REPRINT
© FORTINET



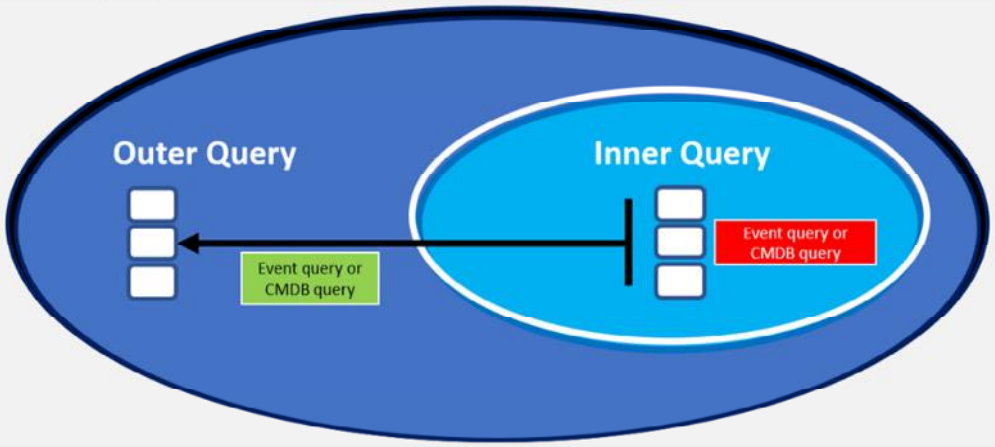
The inner query runs first, and the outer query is executed with results from the inner query. Both the inner and outer queries can be event queries or CMDB queries. It is very important to understand that if the inner query is an event query, then the query results need to be aggregated results, for example, single row with a COUNT, rather than a list of raw events. The outer query references a single column of the inner query results, and both the queries can be run over different time intervals.

It sounds complicated, but it is actually very easy once you understand the basics.

DO NOT REPRINT
© FORTINET

Nested Queries—Supported Scenarios

- Outer CMDB query, inner event query
- Outer event query, inner event query
- Outer event query, inner CMDB query



FORTINET Training Institute

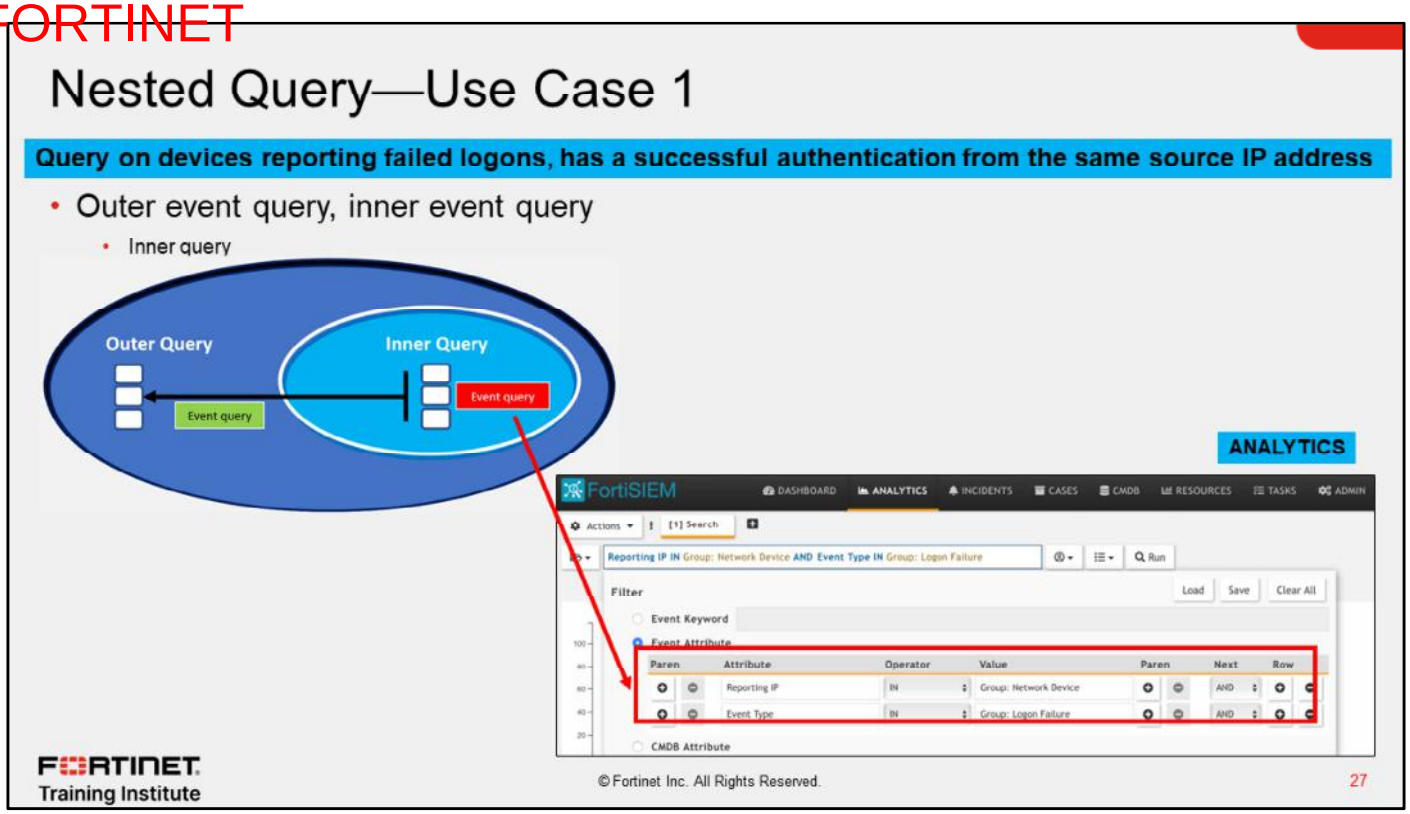
© Fortinet Inc. All Rights Reserved.

26

This slide shows the supported combinations of inner and outer queries, which are:

- Outer CMDB query, inner event query
- Outer event query, inner event query
- Outer event query, inner CMDB query

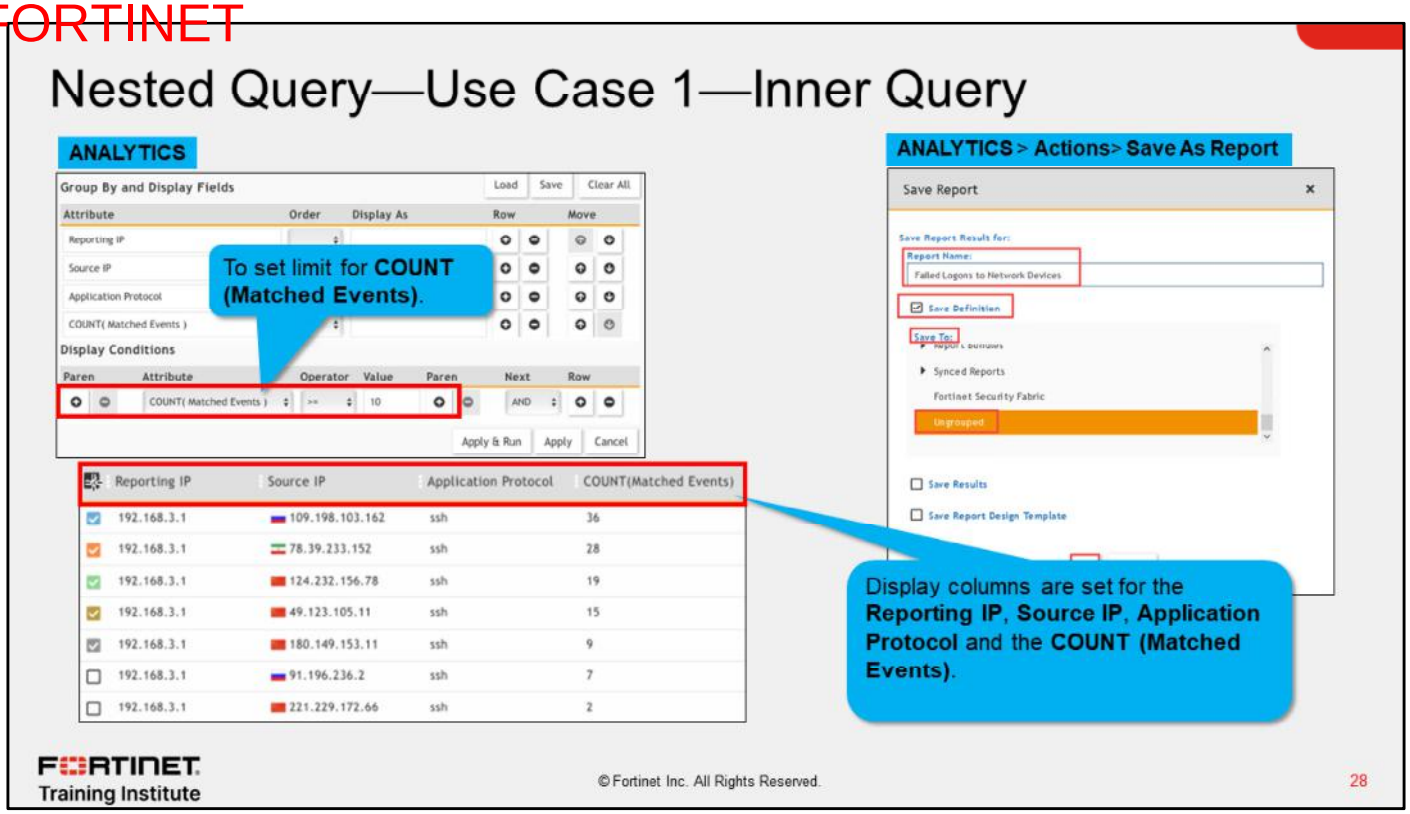
DO NOT REPRINT
© FORTINET



Consider the use case where a customer has some network devices out in the public network that always have failed authentications. The customer would like to query if any of those devices reporting failed logons, has a successful authentication from the same **Source IP** address within a specified time period.

In this case, the inner query is an aggregated set of results of the failed logons reported by the network devices, saved as the **Report definition**: **Failed Logons to Network Devices**.

DO NOT REPRINT
© FORTINET

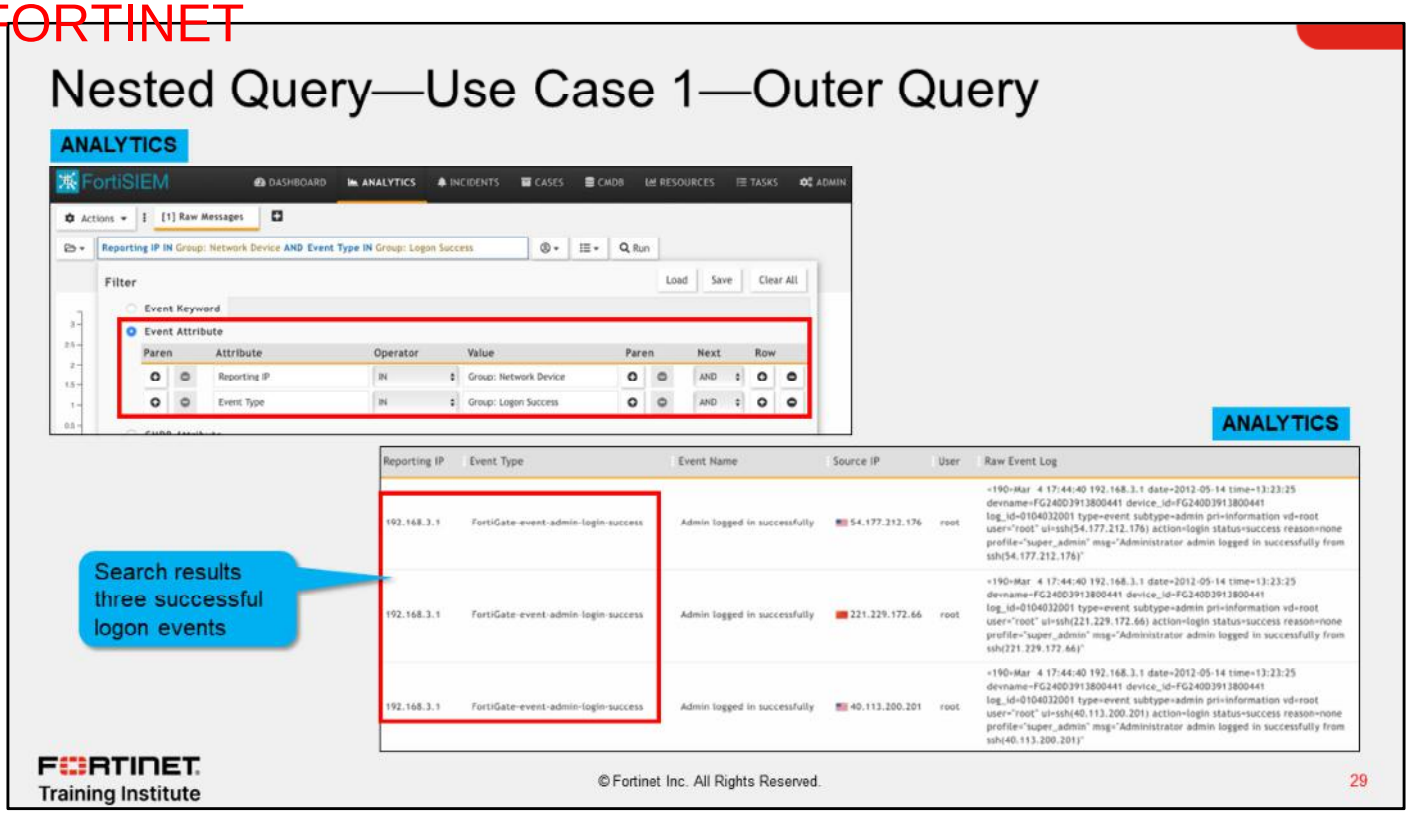


Running the report in the example dataset produces the set of aggregated results over the last 24 hours, as this slide shows.

Notice display columns are for the **Reporting IP**, **Source IP**, **Application Protocol**, and the **COUNT (Matched Events)**.

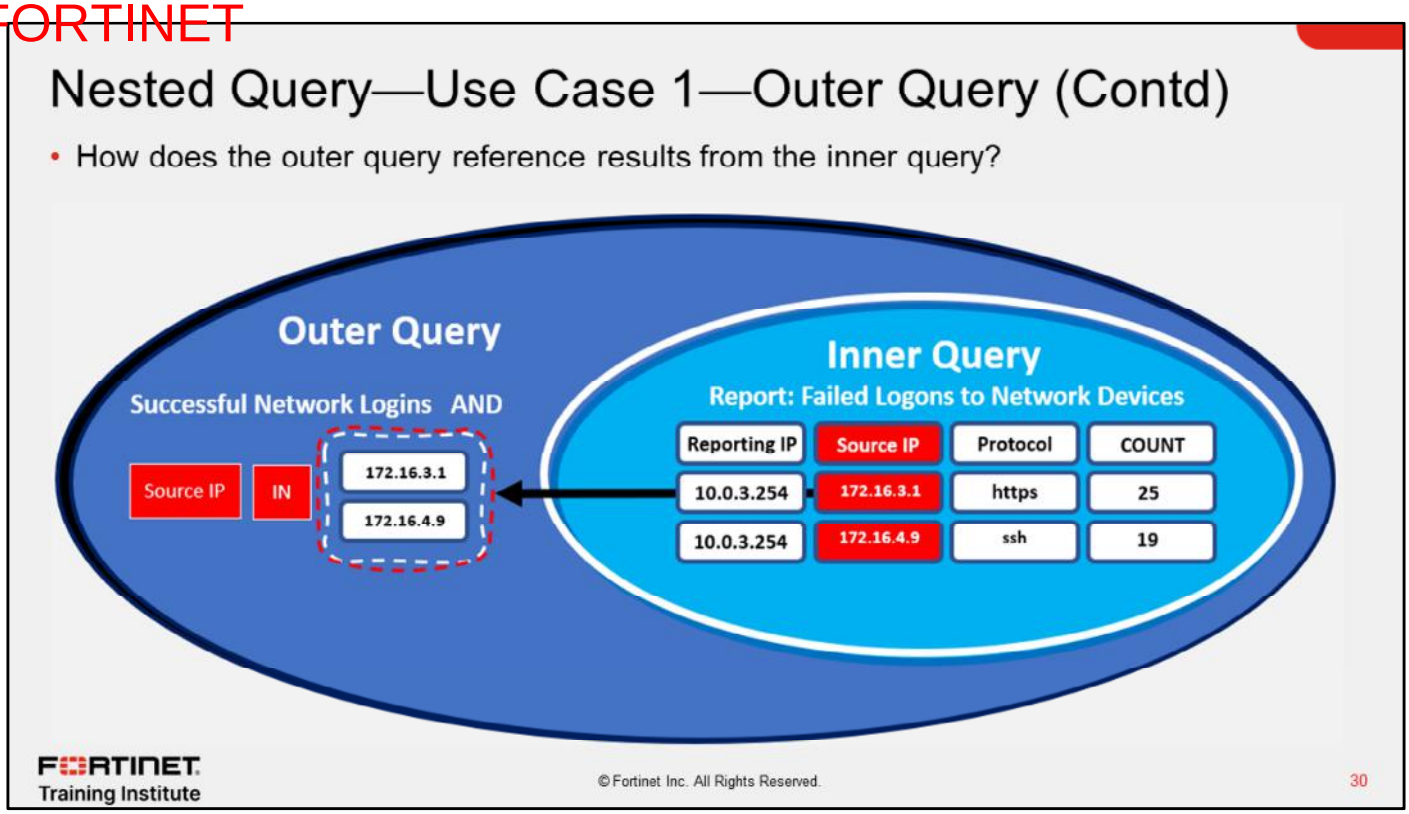
Note that if the use case was to include **Source IP** addresses that had 10 or more logon failures, then this aggregation restriction could have been made under the **Display Conditions**.

DO NOT REPRINT
© FORTINET



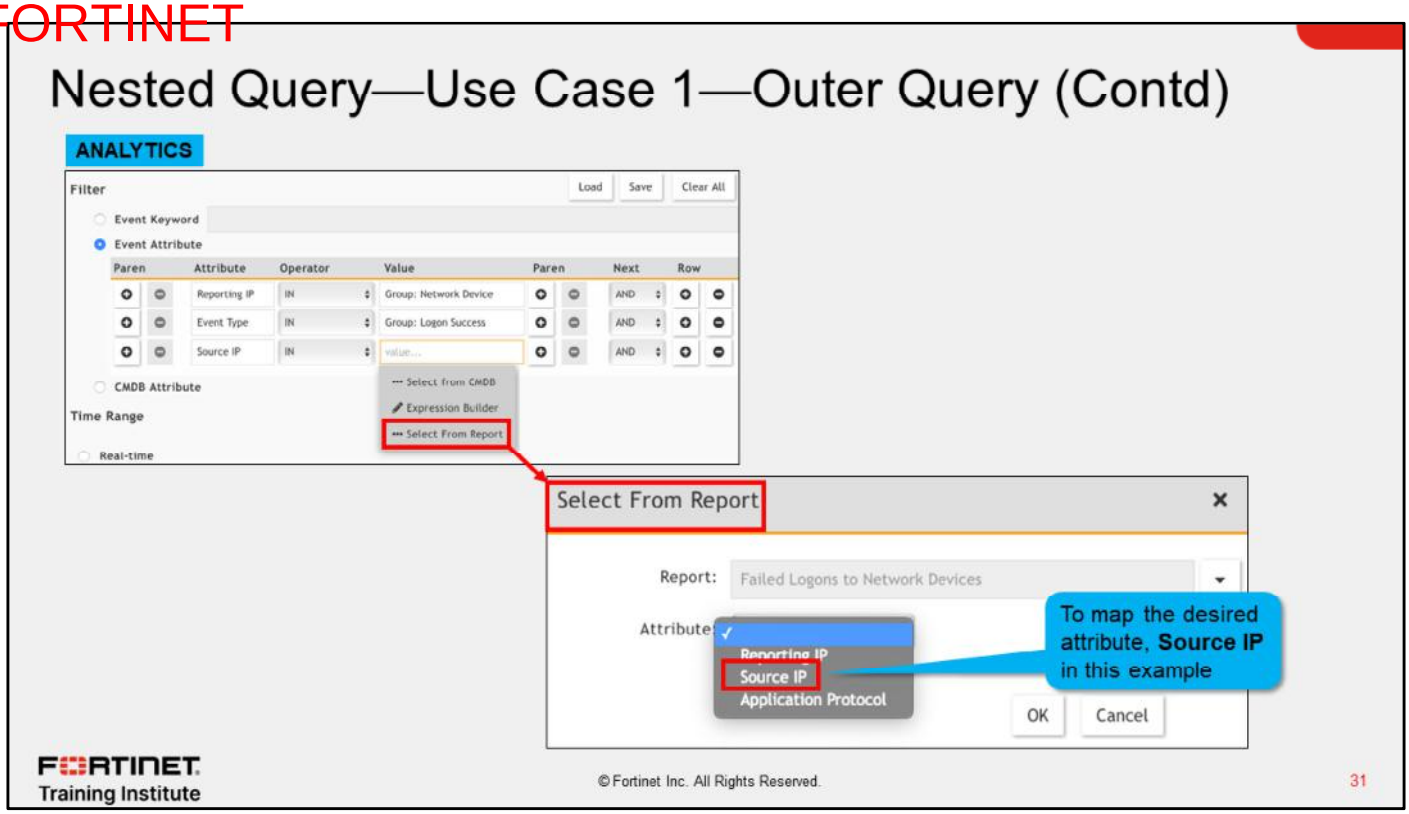
The outer query looks at the successful authentications, partially constructing the search condition for events from the group **Network Device**, which are of **Event Type Logon Success**.

DO NOT REPRINT
© FORTINET



You must decide which **Display Column** in the inner query report contains the data you need, and which event attribute you going to query for it. In this use case the inner query **Source IP** display column contains the data of the source of the failed logins, and you want to query if the **Source IP** in the **Successful Logon** events matches any of those results in that dataset.

DO NOT REPRINT
© FORTINET

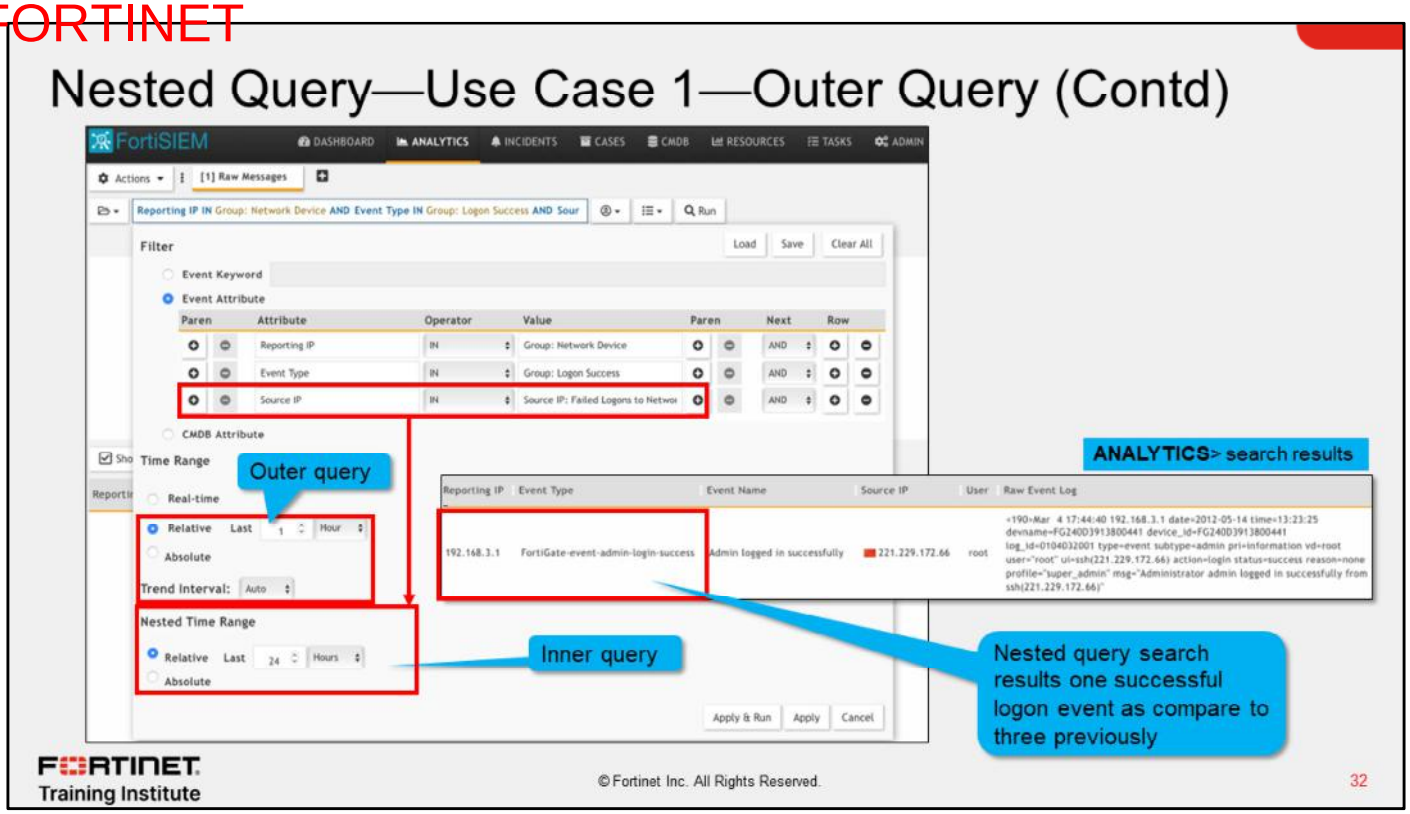


You can construct this by adding an additional row to the outer query, by using **Source IP** as the **Attribute**, using the **Operator IN** and using the **Select From Report** option for the **Value**.

When the **Select From Report** option is used, you must find the event or CMDB report definition you saved as the inner query and then select **Display Column** to map to the **Attribute**.

Note that the **Select From Report** option shows only the non-aggregation columns.

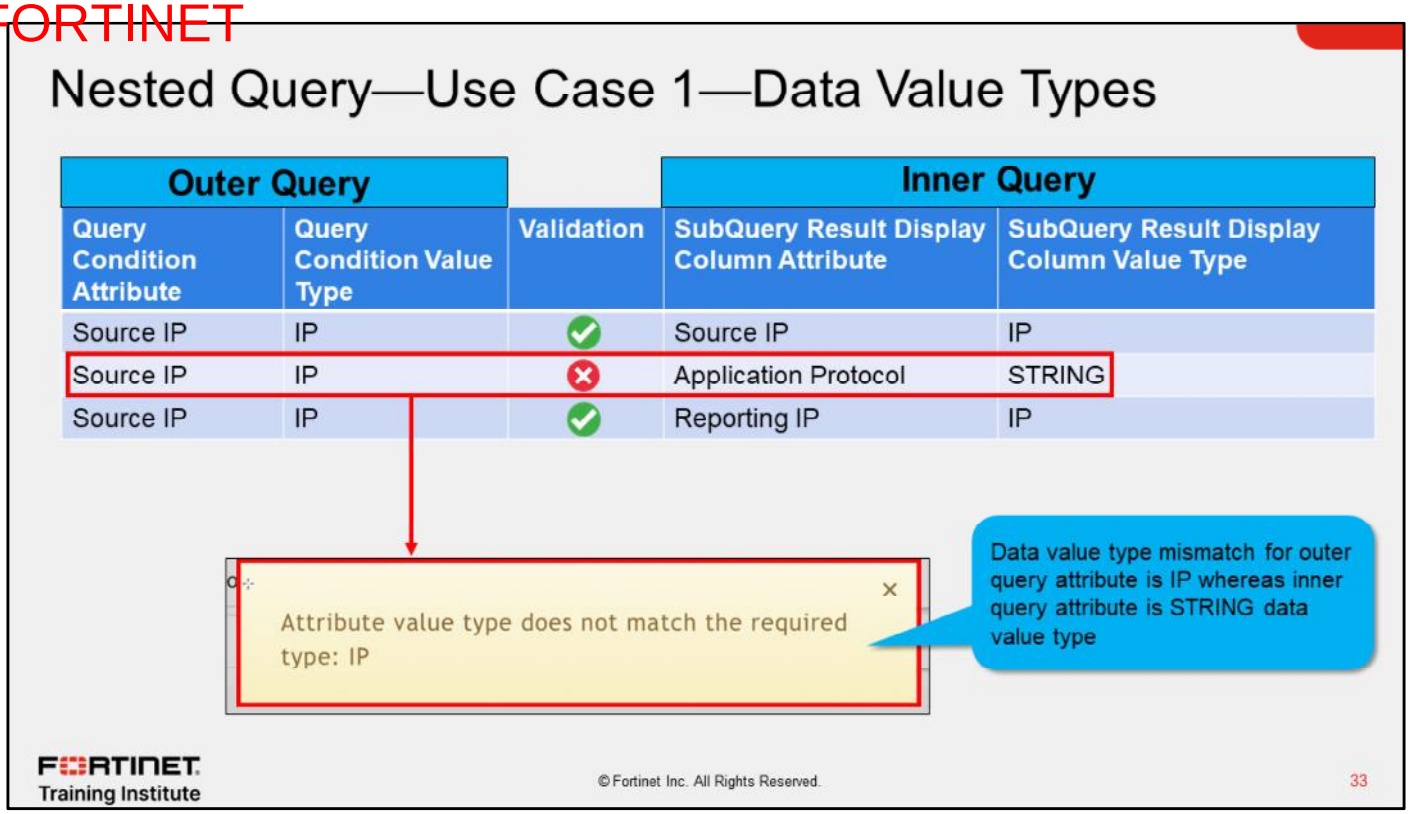
DO NOT REPRINT
© FORTINET



This slide shows how the final query looks. Notice the different time ranges possible for the inner query **Nested Time Range**, and outer query **Time Range**, when both queries are event queries.

Running the search now against the sample dataset returns only one entry rather than three as before because this was the only event where the **Source IP** was also seen to have **Logon Failures** against the network devices.

DO NOT REPRINT
© FORTINET



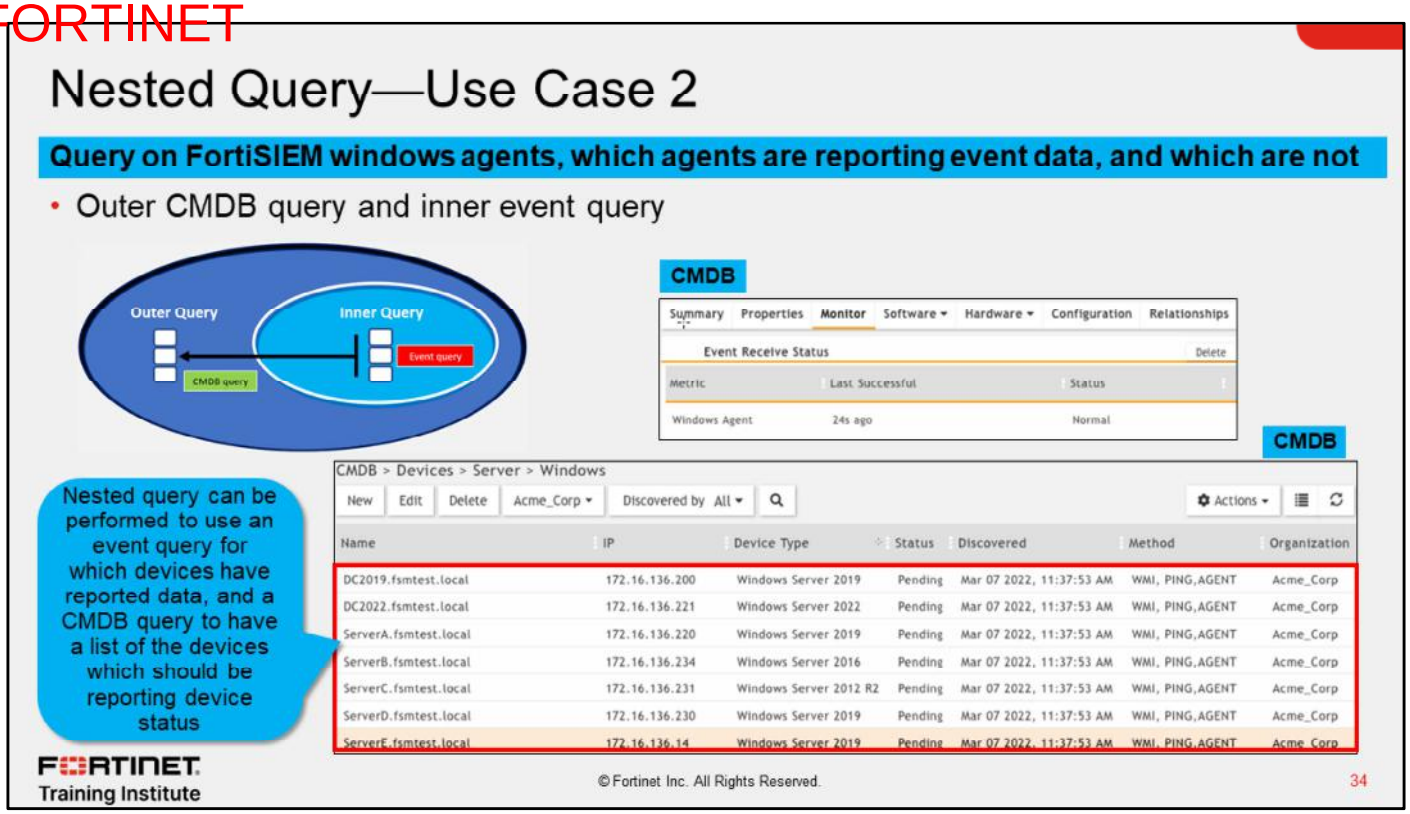
Another important aspect to understand is that the data value types must match. In this case, the inner query, had three display columns that the outer query could reference. Each of those display column attributes has a data value type such as IP, String, INT16, and so on. You can verify this under **ADMIN** > **Device Support** > **Event Attributes**. The attribute you use in the outer query to query the inner query must match the data value types.

In this use case, the first row used **Source IP** in both cases for the outer and inner attribute, and the data value type of IP matches.

In the strange event that the **Application Protocol** was selected as the inner query attribute, which has a data value type of **STRING**, then attempting to use **Source IP** in the outer query would not work because the data value type of **Source IP** is **IP**. This would result in a pop-up error window, as this slide shows.

The third row above would also work, even though the attributes used in the inner and outer query are different, the data value type of both is IP, which satisfies the logic.

DO NOT REPRINT
© FORTINET

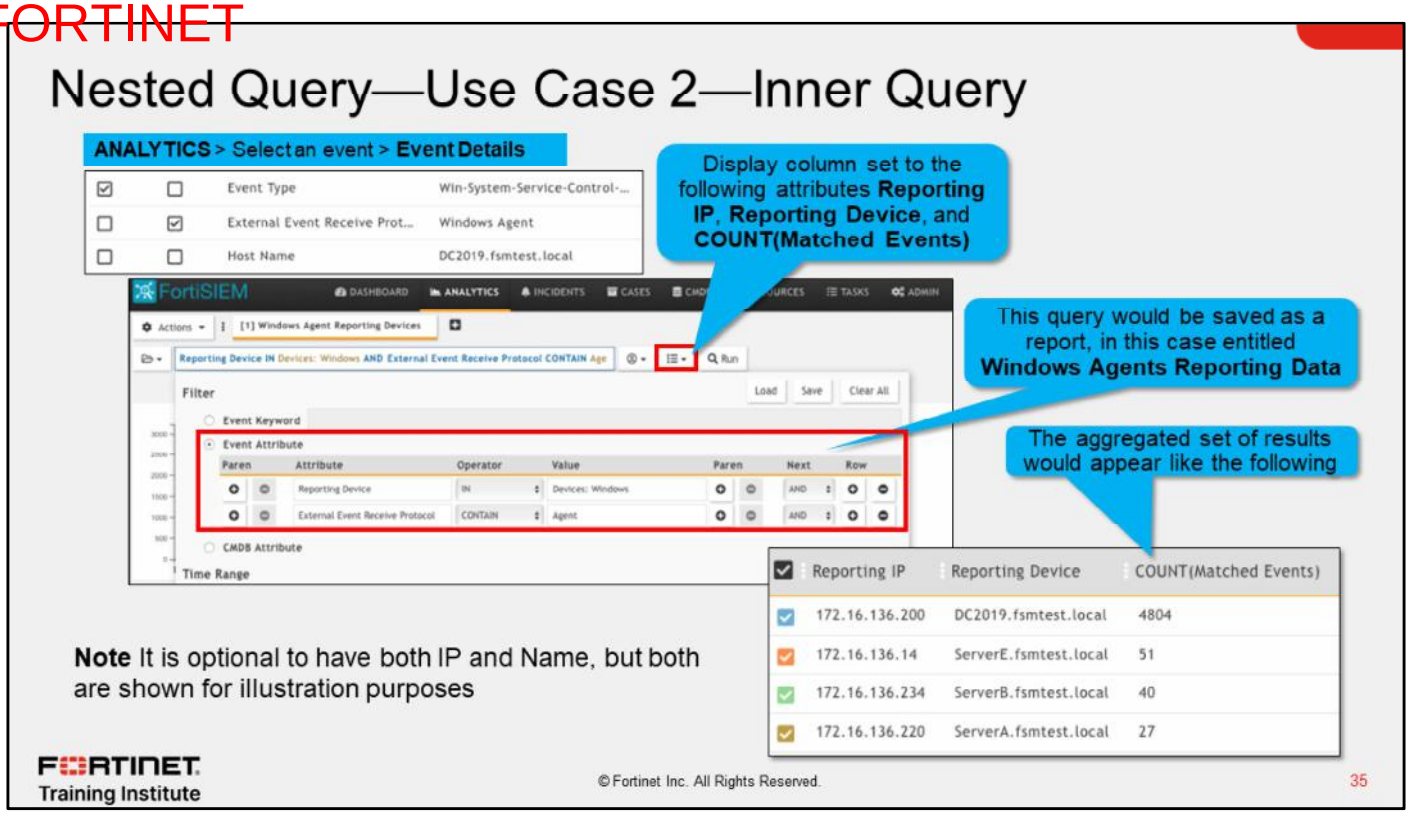


Consider the use case where you have a large estate of FortiSIEM windows agents but because of changes in personnel and network issues, you do not know which are reporting event data and which are not.

On an individual device level in the CMDB, the **Monitor** tab records the last event time stamp for the **Metric Windows Agent**, but you cannot report on these values globally.

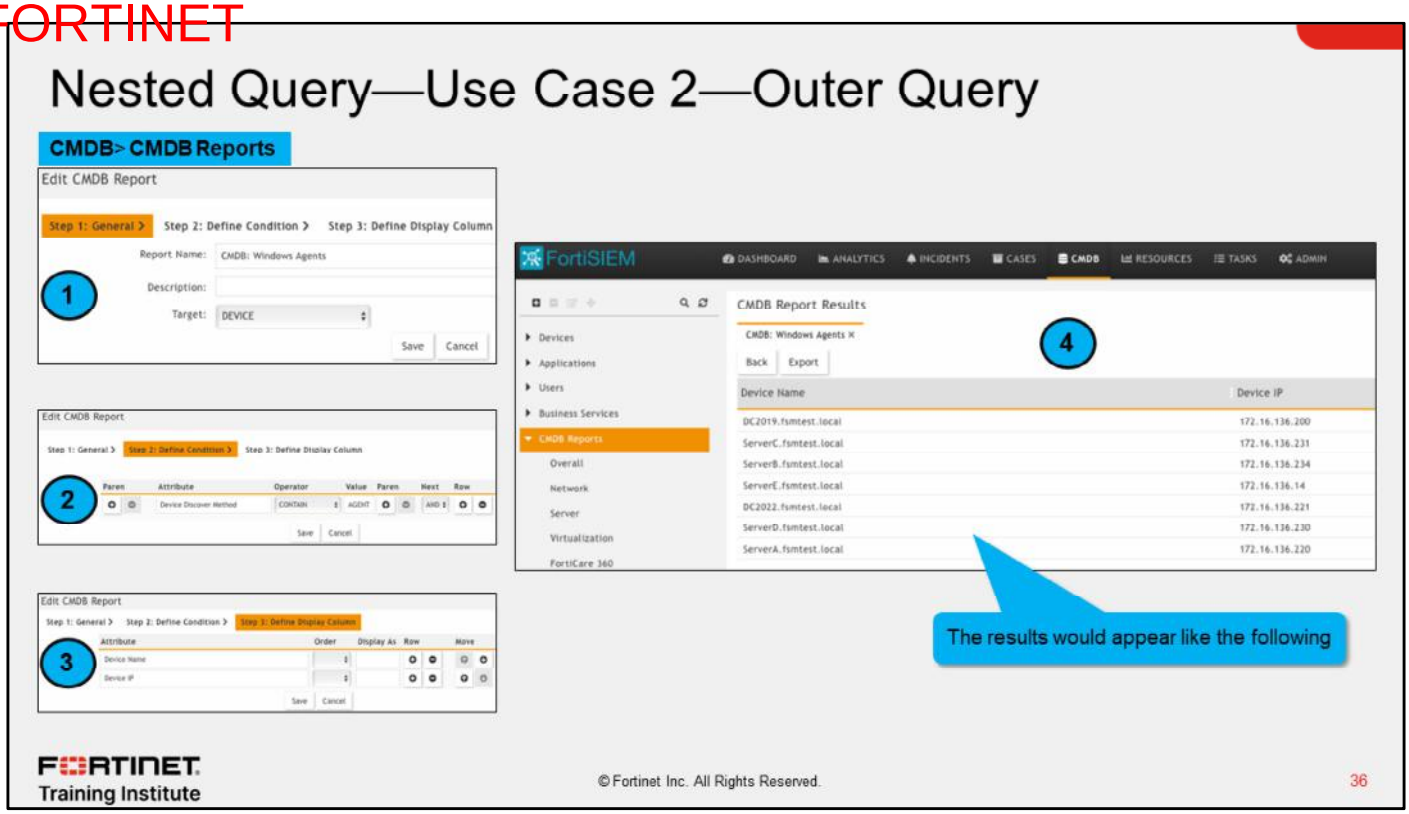
Consider an environment with few windows agents installed, as the image on the slide shows. This is a good example of how you can perform a nested query to use an event query for which devices have reported data, and a CMDB query to have a list of the devices which should be reporting.

DO NOT REPRINT
© FORTINET



In this case, the inner query is an event query of the windows agents that are reporting data. It is simply looking for any windows device in the CMDB where the events produced contain the metadata field **External Event Receive Protocol** containing the word **Agent**, as all windows and Linux agents do and grouping the results by the **Reporting IP** and **Reporting Device**. This slide shows the aggregated set of results are shown on slide. This query would be saved as a report, in this case entitled **Windows Agents Reporting Data**.

DO NOT REPRINT
© FORTINET

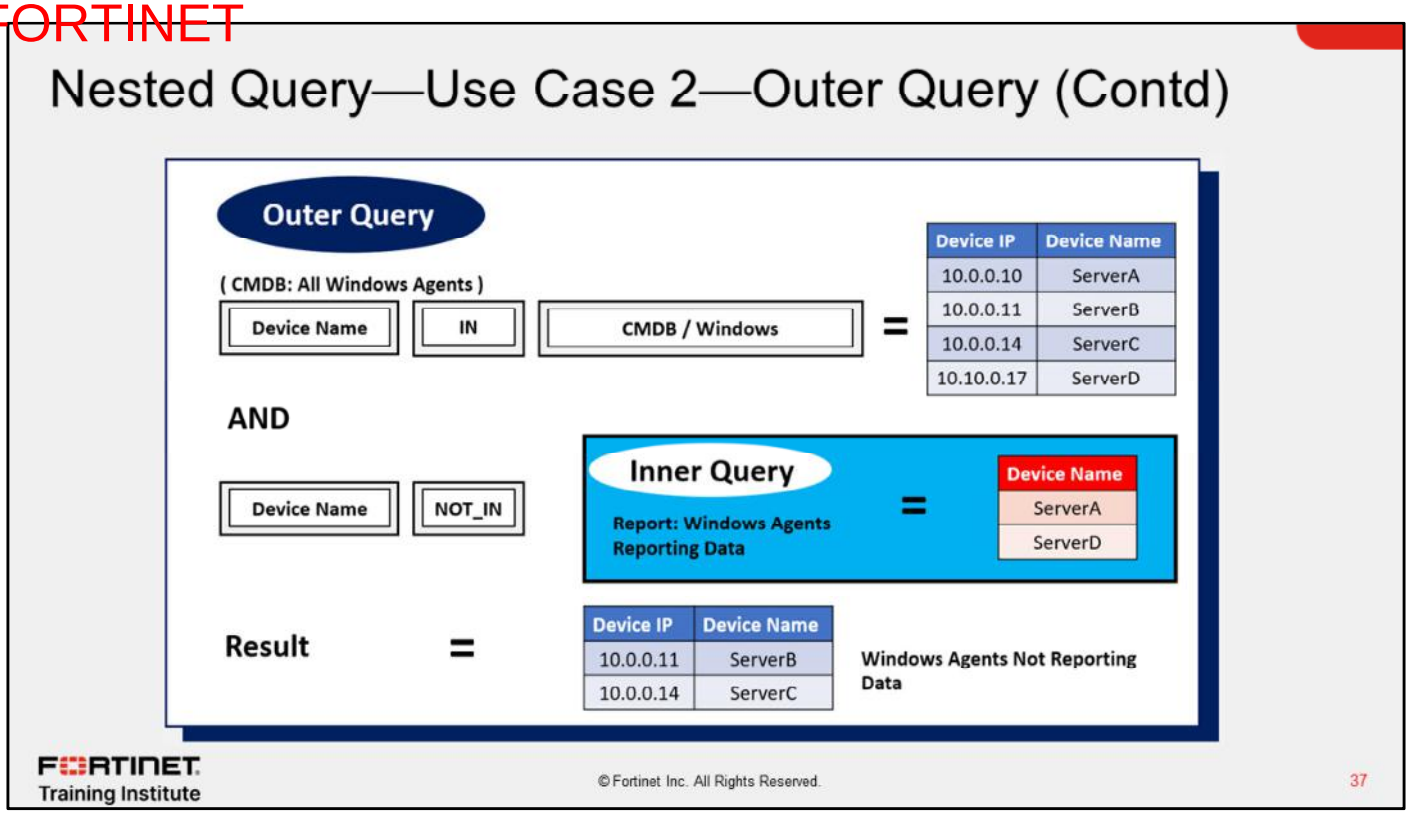


The outer query is a CMDB query. The CMDB reporting feature in FortiSIEM allows reporting on the current content of the CMDB, for example, not time-based event data.

If the report was constructed under **CMDB** > **CMDB Report** of **Target DEVICE**, and the report was saved and run, then FortiSIEM would query the postgres CMDB database and return the rows and data related to device entries themselves in the database, rather than anything to do with received event logs.

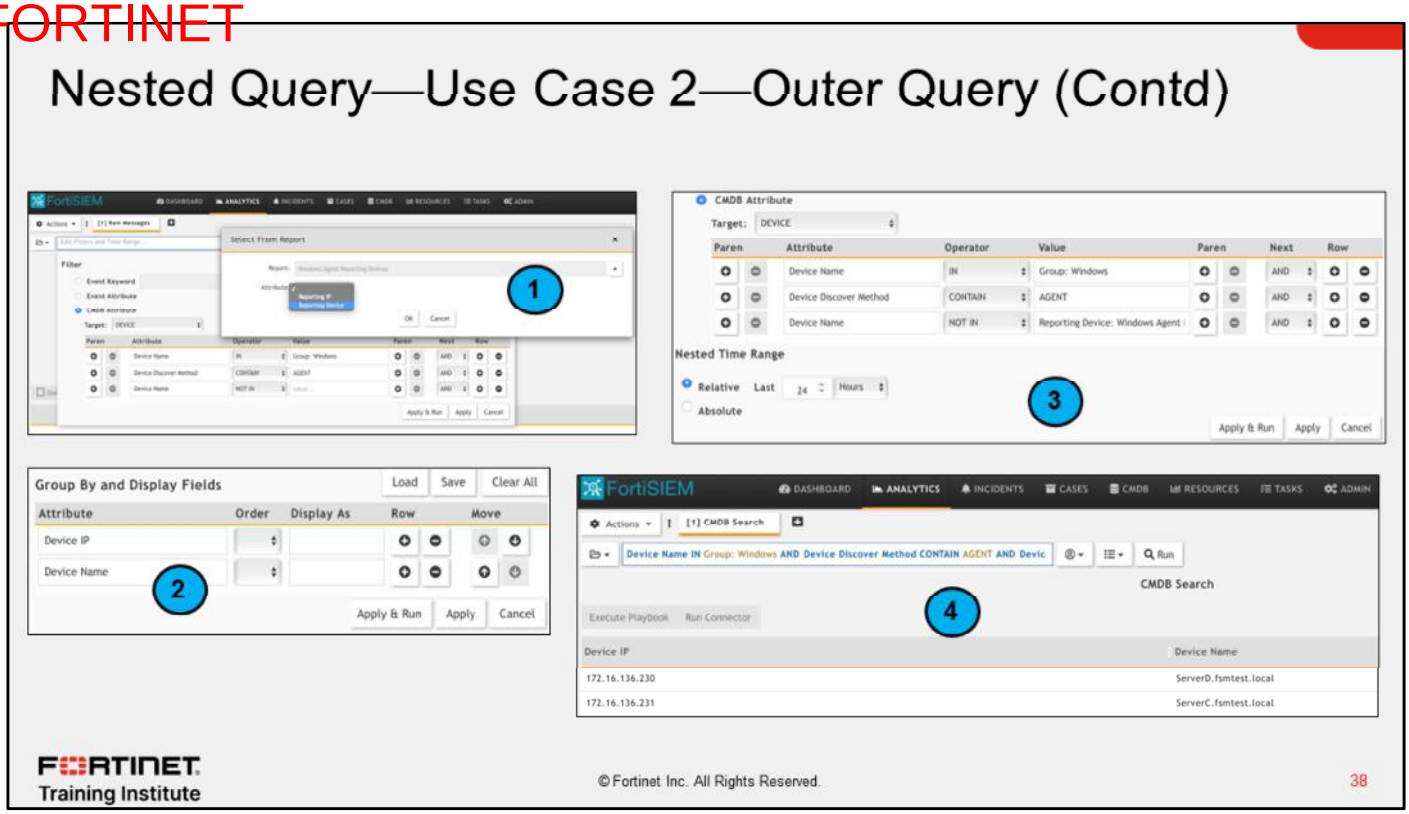
Notice **Device Name** and **Device IP** in the **CMDB** and compare them with the **Reporting IP** and **Reporting Device** in event logs covered on the previous slide.

DO NOT REPRINT
© FORTINET



So, putting this all together, you must construct an outer query to produce a list of all the windows agents in the CMDB, and then subtract from that list through an inner query the list of agents that have reported events within a defined time range. The result is then the list of agents that have not reported data.

DO NOT REPRINT
© FORTINET

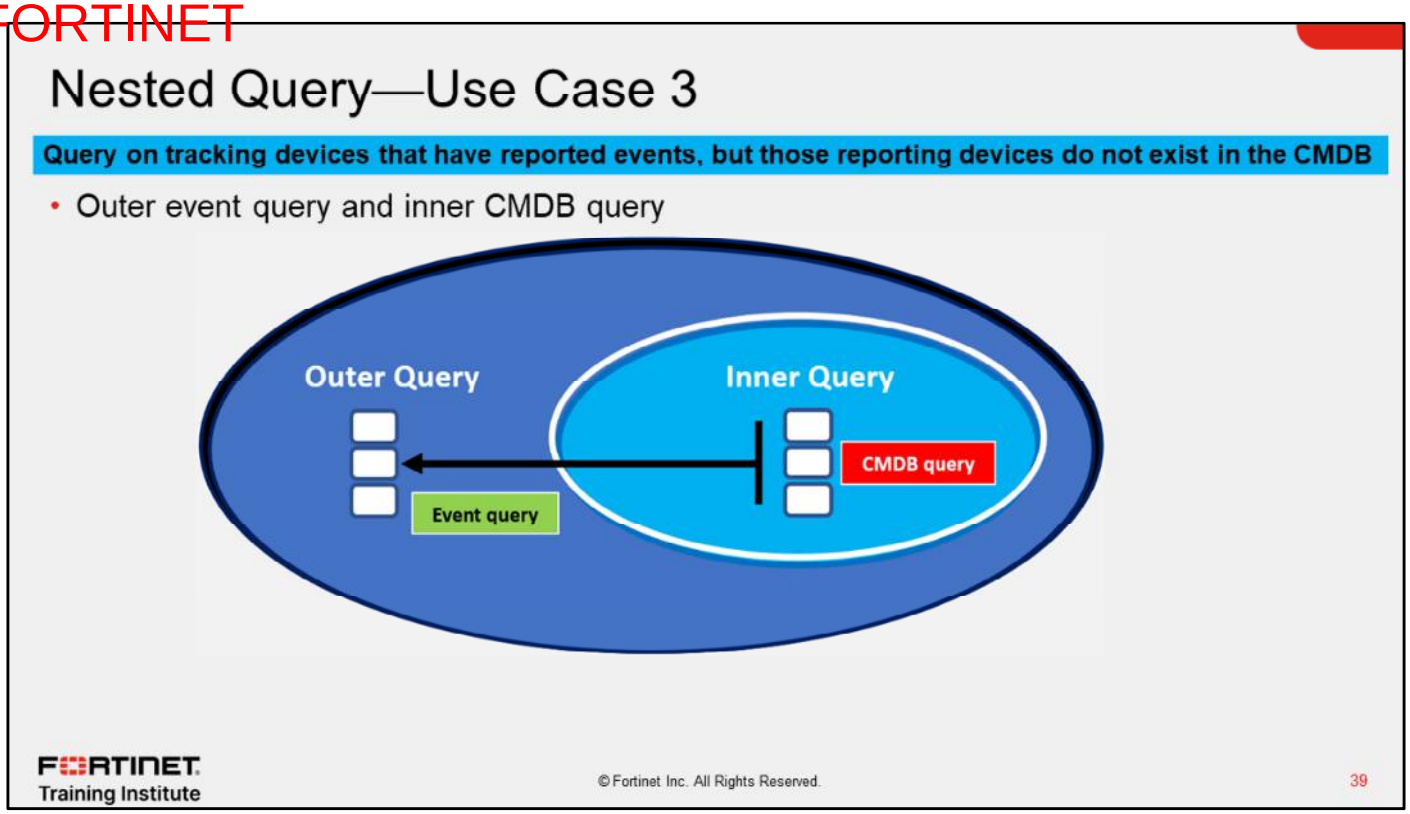


On the **ANALYTICS** tab, construct the outer query using the **CMDB Attribute** option because the outer query is a CMDB in this use case. Add an extra row to ignore entries that have reported agent data through the inner query by selecting **Select From Report** option and mapping the **Device Name** to the **Reporting Device.**

And finally, add **Display Columns** to show from the CMDB the **Device IP** and **Device Name** results from the query. Running the query for the last `24` hours produces the result shown on the slide for devices that did not report any data.

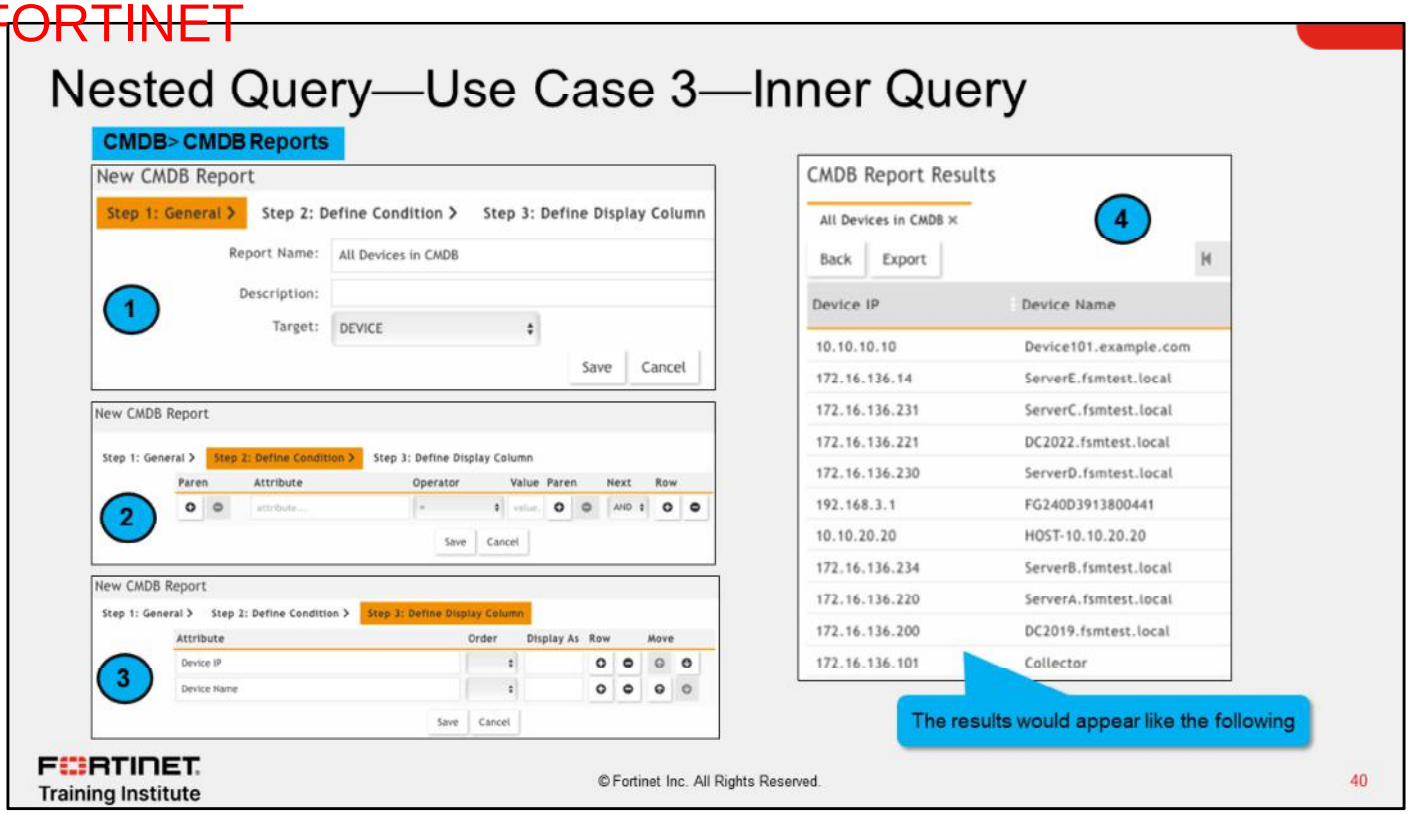
Remember the time range is for a nested query, which is an inner query. The outer query is an CMDB query and therefore the time range does not apply to CMDB queries.

DO NOT REPRINT
© FORTINET



Consider another scenario where the outer query is an event query and the inner query is a CMDB query. The example for this use case would be tracking devices that have reported events, but those reporting devices do not exist in the CMDB.

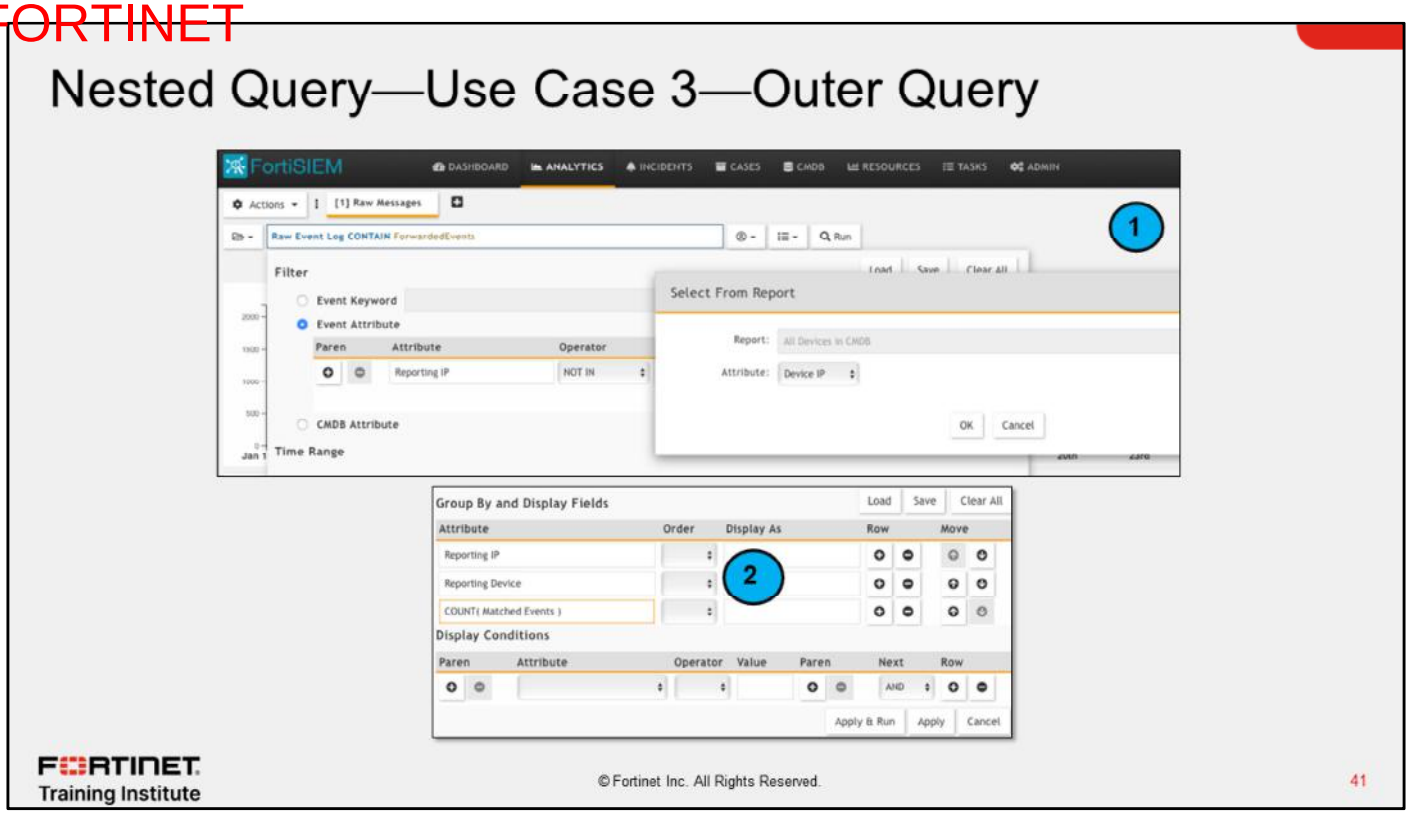
DO NOT REPRINT
© FORTINET



In this scenario, the inner query is a CMDB query. In **Step 1: Target** is set to **Device**, and in **Step 2:** there is no there is no specific condition. Leave it blank to report on all devices. Finally, in **Step 3**: for the **Display Column** add **Device IP** and the **Device Name**.

The CMDB report produces a list of all the devices with an example data set in the **All Devices in CMDB** report to list all devices in CMDB, as this slide shows.

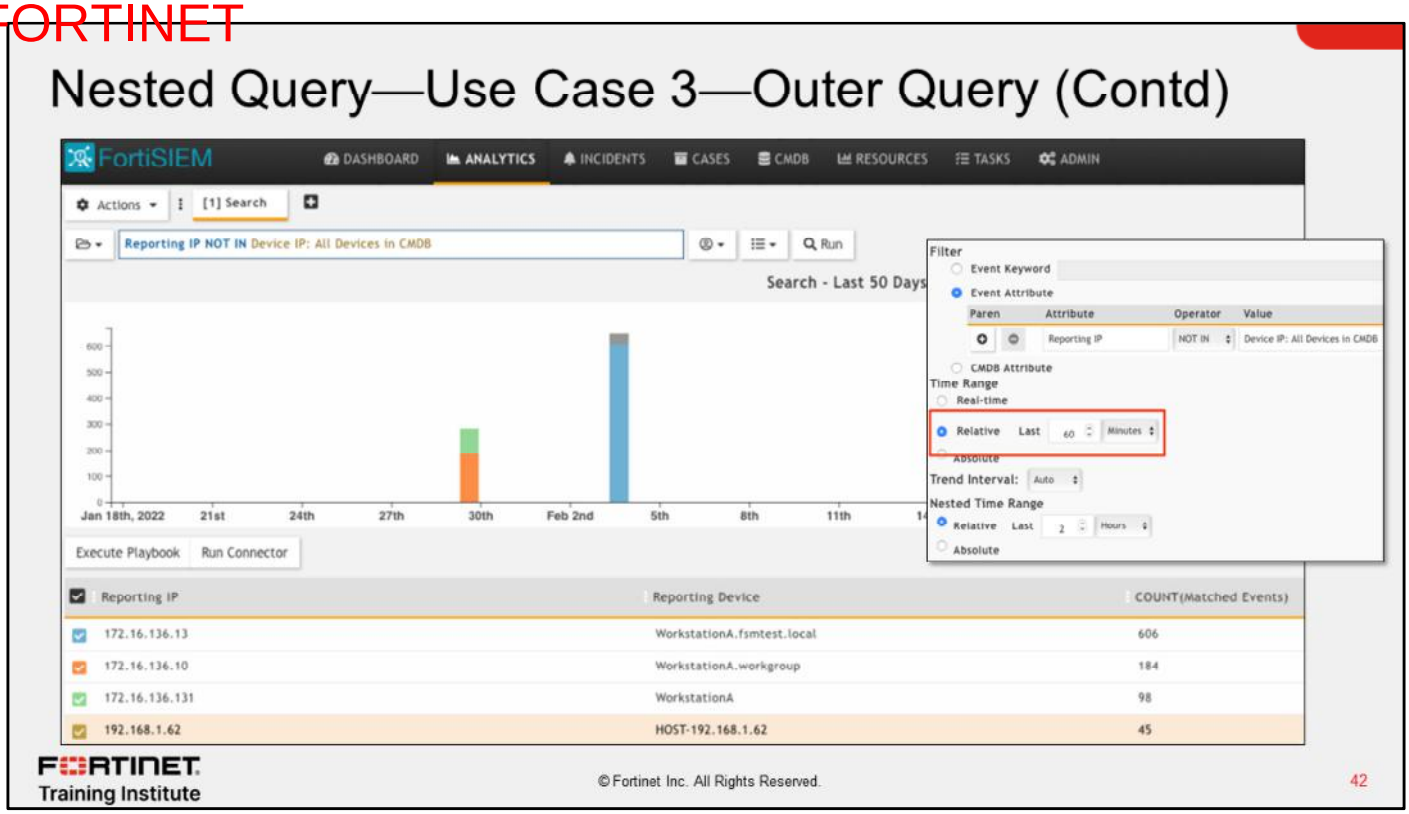
DO NOT REPRINT
© FORTINET



The outer query is an event query. This time, because you are looking at data being reported from all devices, there is no specific condition applied except to query the contents of the inner query.

This slide shows that, the query is looking for any **Reporting IP** that is **NOT IN** in **Select From Report** and references the inner query CMDB report map of the **Device IP**.

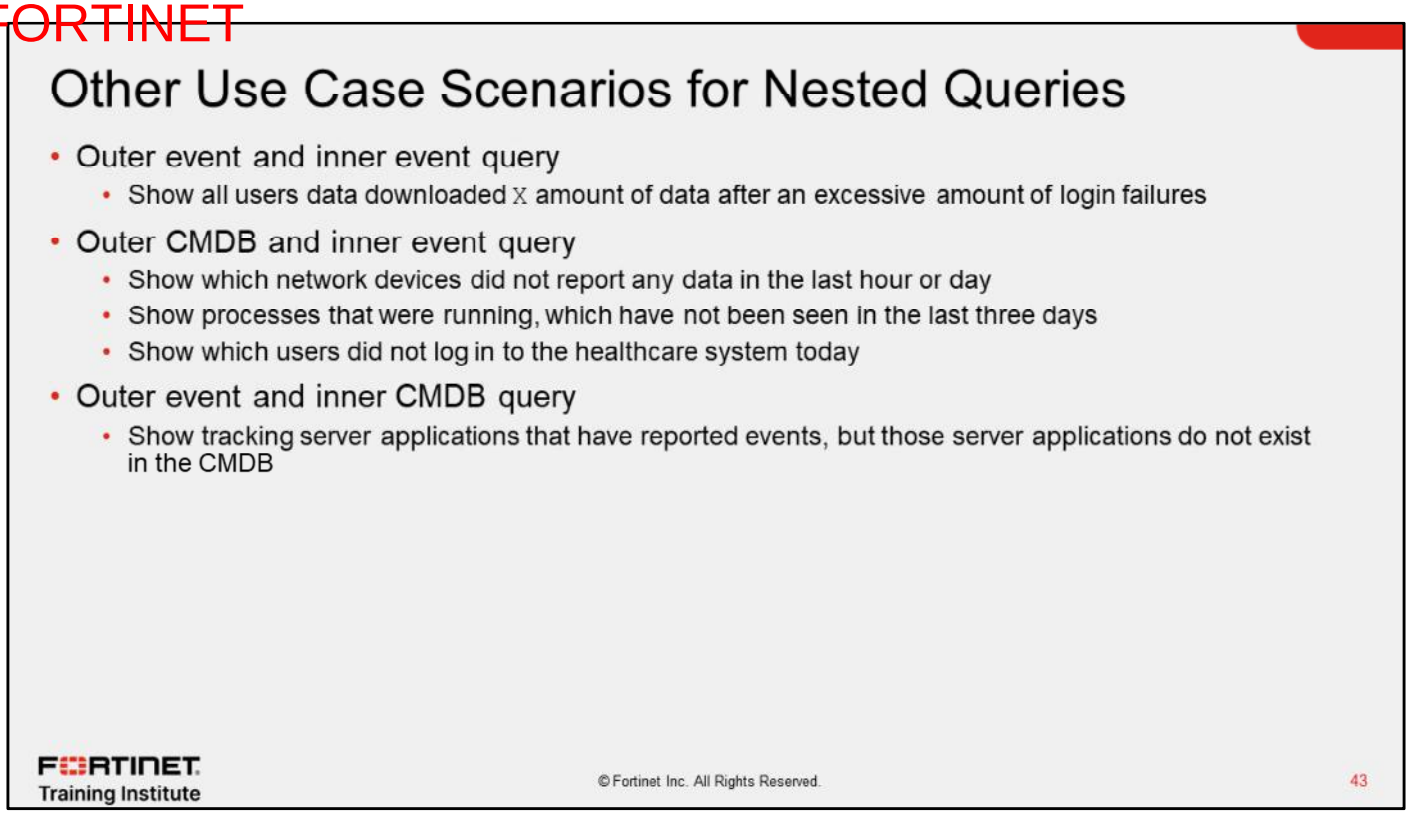
DO NOT REPRINT
© FORTINET



This slide shows the results of a nested query, where these devices had reported data but are not currently present in the CMDB.

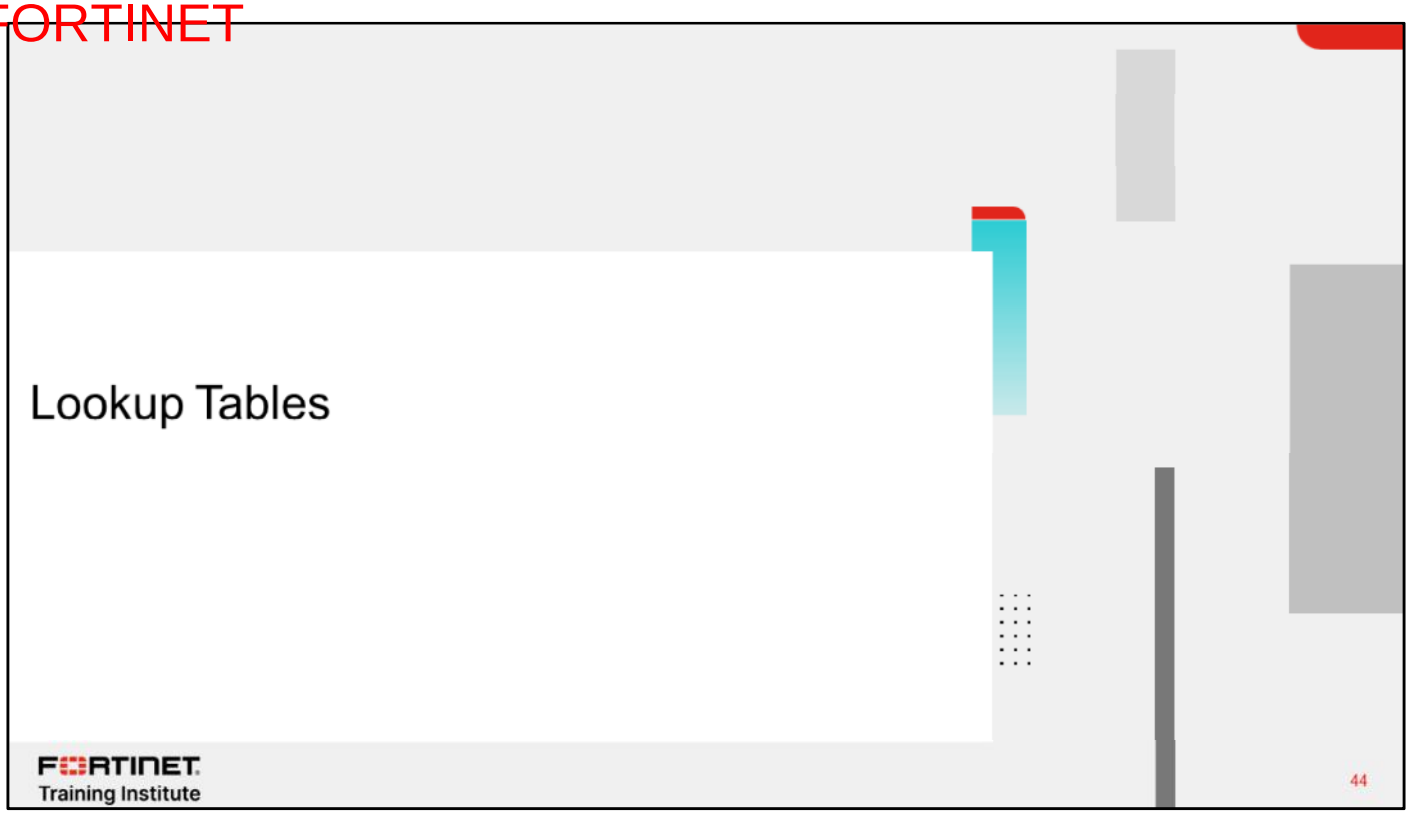
Note that in this case, since the inner query was a CMDB query, the **Nested Time Range** is not relevant because it is static data so the standard **Time Range** selector would determine how far back the query was to be run.

DO NOT REPRINT
© FORTINET



This slide shows examples of a few more use cases for nested queries.

DO NOT REPRINT
© FORTINET

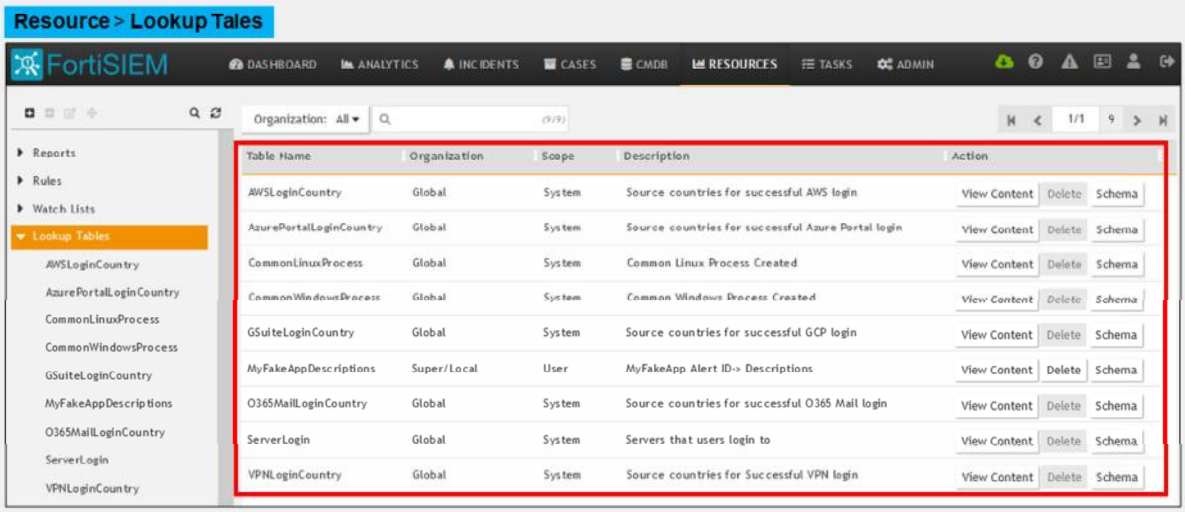


In this section, you will learn about a lookup tables on FortiSIEM for advance analytics.

DO NOT REPRINT
© FORTINET

What Are Lookup Tables?

- Lookup tables are data tables that contain one or more columns, outside of event log data, which can be used in analytical queries, reports, and rules



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

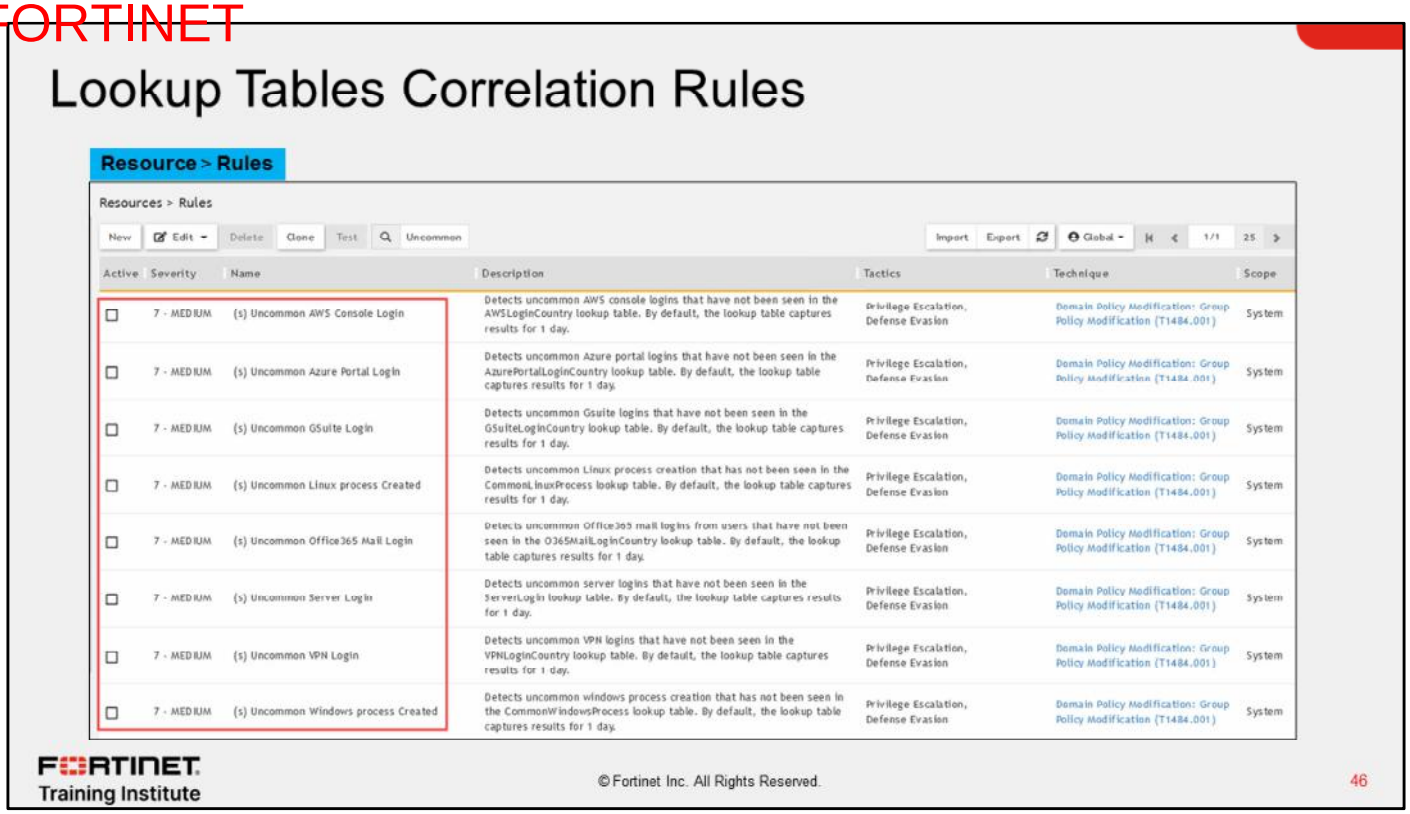
45

Lookup tables are data dictionaries that contain mappings of unique keys to values. They allow you to reference any key or value to enrich or filter the data in your analytics and rules conditions. A single lookup table can contain up to five keys and as many values required.

FortiSIEM has the ability to enrich analytics with information from lookup tables to provide insight into the data being retrieved from logs. It can also create correlation rules based on contents and conditions in one or more lookup tables.

This slide shows nine out-of-the-box lookup tables. You can populate lookup manually, by running an on-demand or scheduled report, by uploading a CSV file, or using the new lookup table API.

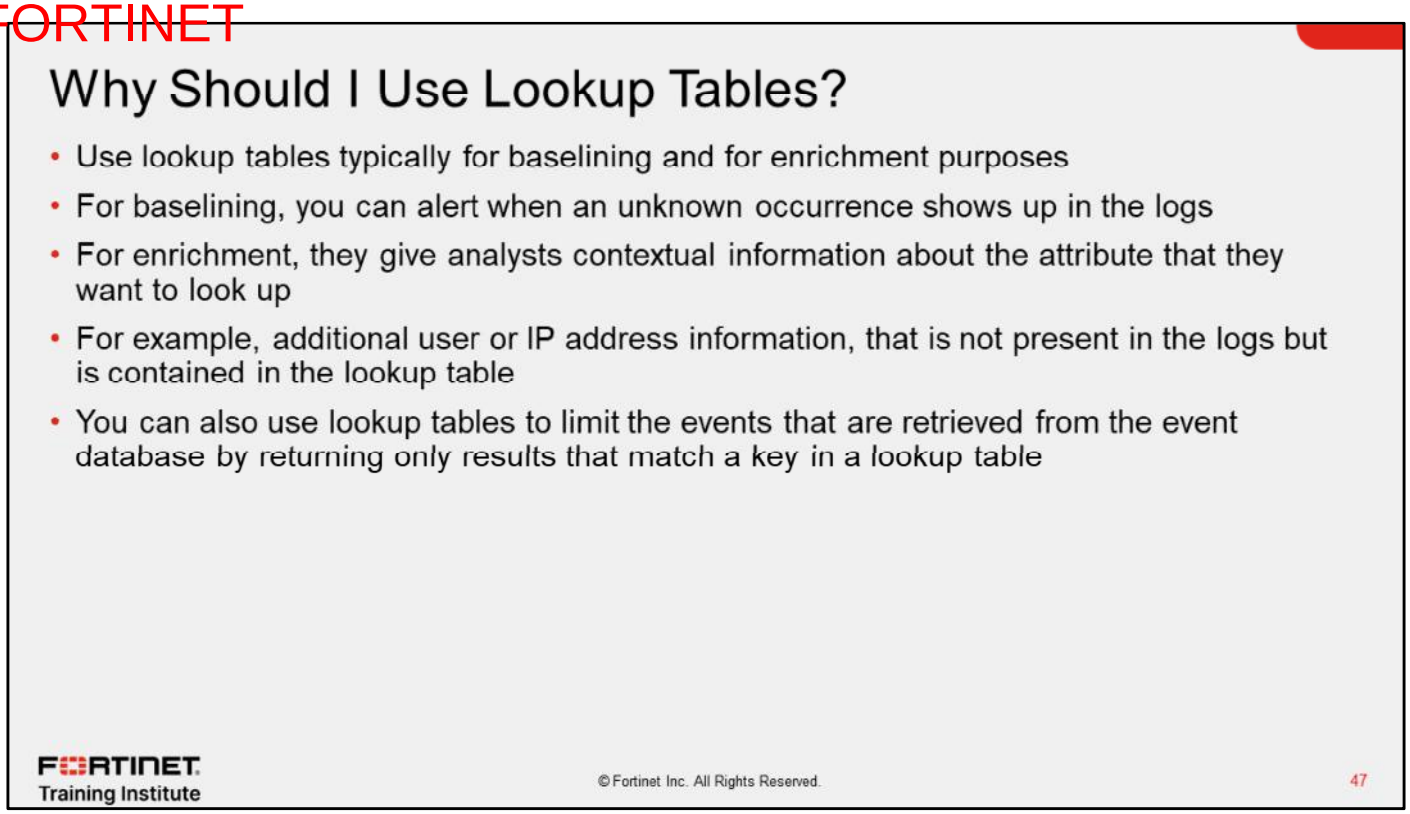
DO NOT REPRINT
© FORTINET



Lookup tables are accompanied by new correlation rules that trigger in real time when FortiSIEM receives an event that does not match the existing baseline.

The new rules associated with lookup tables are:
- Uncommon AWS Console Login
- Uncommon Azure Portal Login
- Uncommon GSuite Login
- Uncommon Linux process Created
- Uncommon Office365 Mail Login
- Uncommon Server Login
- Uncommon VPN Login
- Uncommon Windows process Created

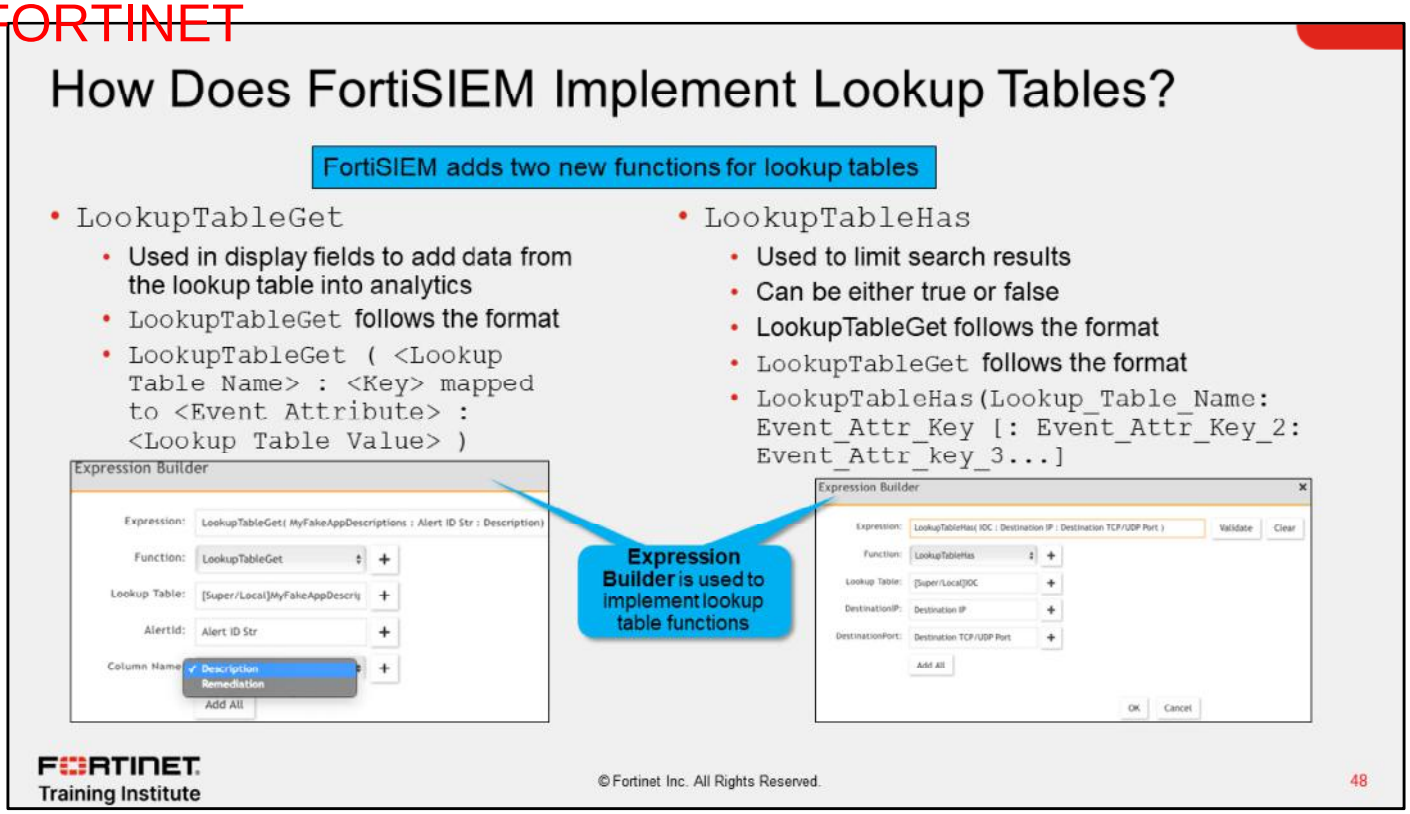
DO NOT REPRINT
© FORTINET



You typically use lookup tables for baselining and for enrichment purposes. By using them for baselining, you can be notified when an unknown occurrence shows up in the logs.

When used for enrichment, they give analysts contextual information about the attribute that they want to look up. For example, additional user or IP address information that is not present in the logs but is contained in the lookup table. You can also use lookup tables to limit the events that are retrieved from the event database by returning only results that match a key in a lookup table.

DO NOT REPRINT
© FORTINET



FortiSIEM adds two new functions for lookup tables:

- `LookupTableGet`: Used for data enrichment and filtering purposes
- `LookupTableHas`: used for data filtering purposes

`LookupTableHas` can be either true or false. If set to true then it will filter event retrieval and only return data that is present in the lookup table. If set to false then it will retrieve all data when the attributes are not present in the lookup table.

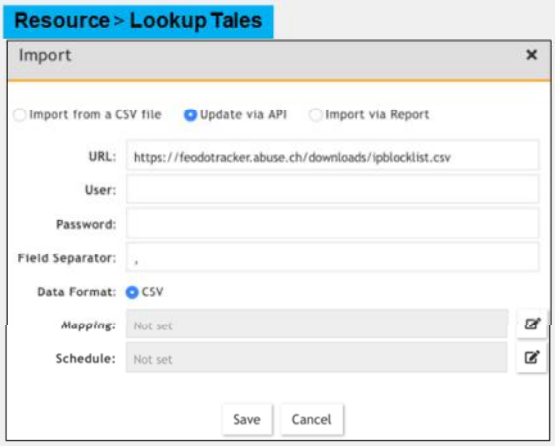
`LookupTableGet` retrieves a value from a lookup table and can be used as a filter or within display fields.

FortiSIEM provides a REST API for lookup table management, as well as manual creation using the GUI, and automatic population through running reports or uploading CSVs.

DO NOT REPRINT
© FORTINET

Configuring Lookup Tables

- Manually
- Using CSV file
- Using analytic report
- Using URL (API in the GUI)
  - Lookup tables can import data using a URL
  - *This is the slightly misleading through API option in the GUI*
- Using API



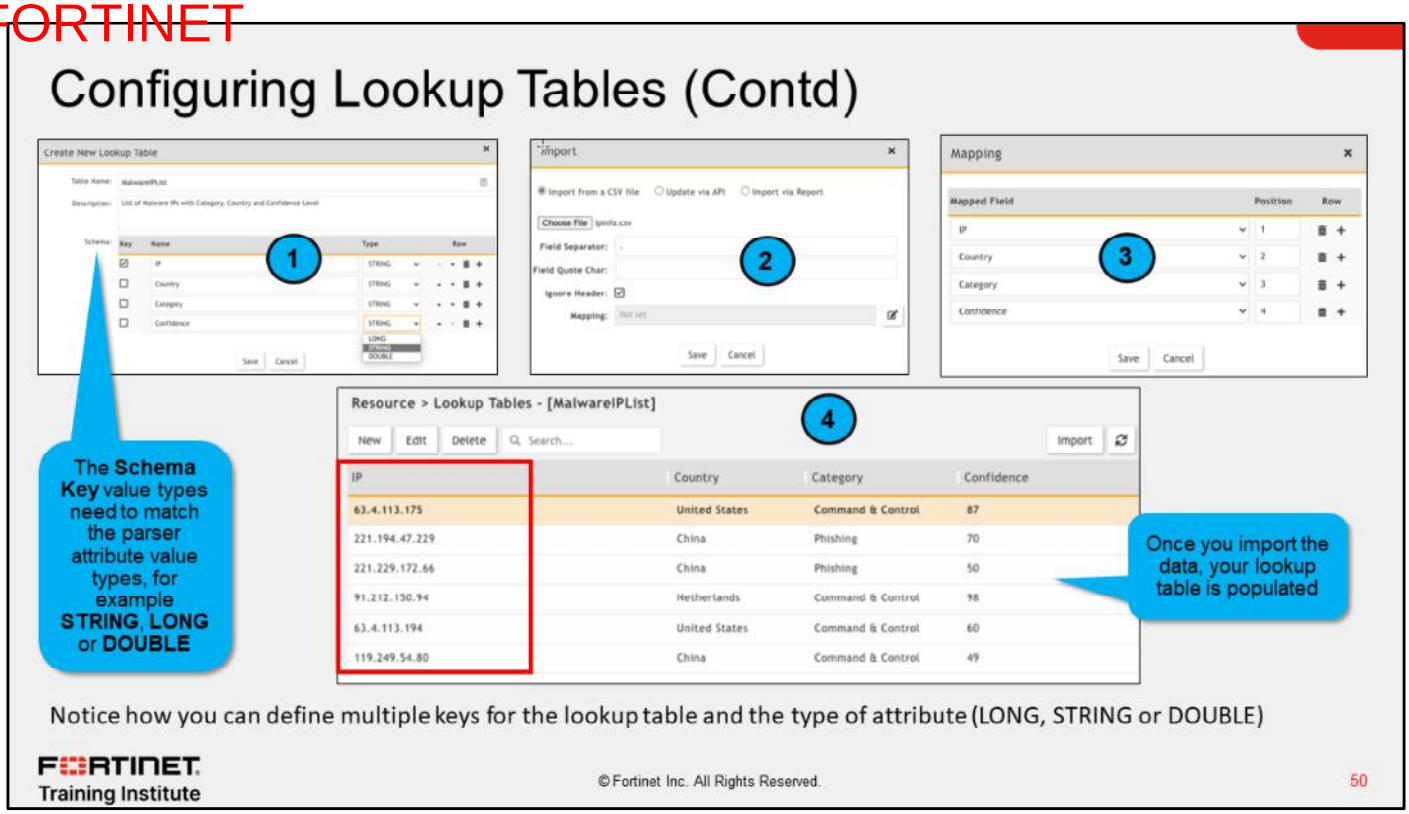
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

49

You can populate lookup tables manually, or run an on-demand or scheduled report, upload a CSV file, or use the new lookup table API.

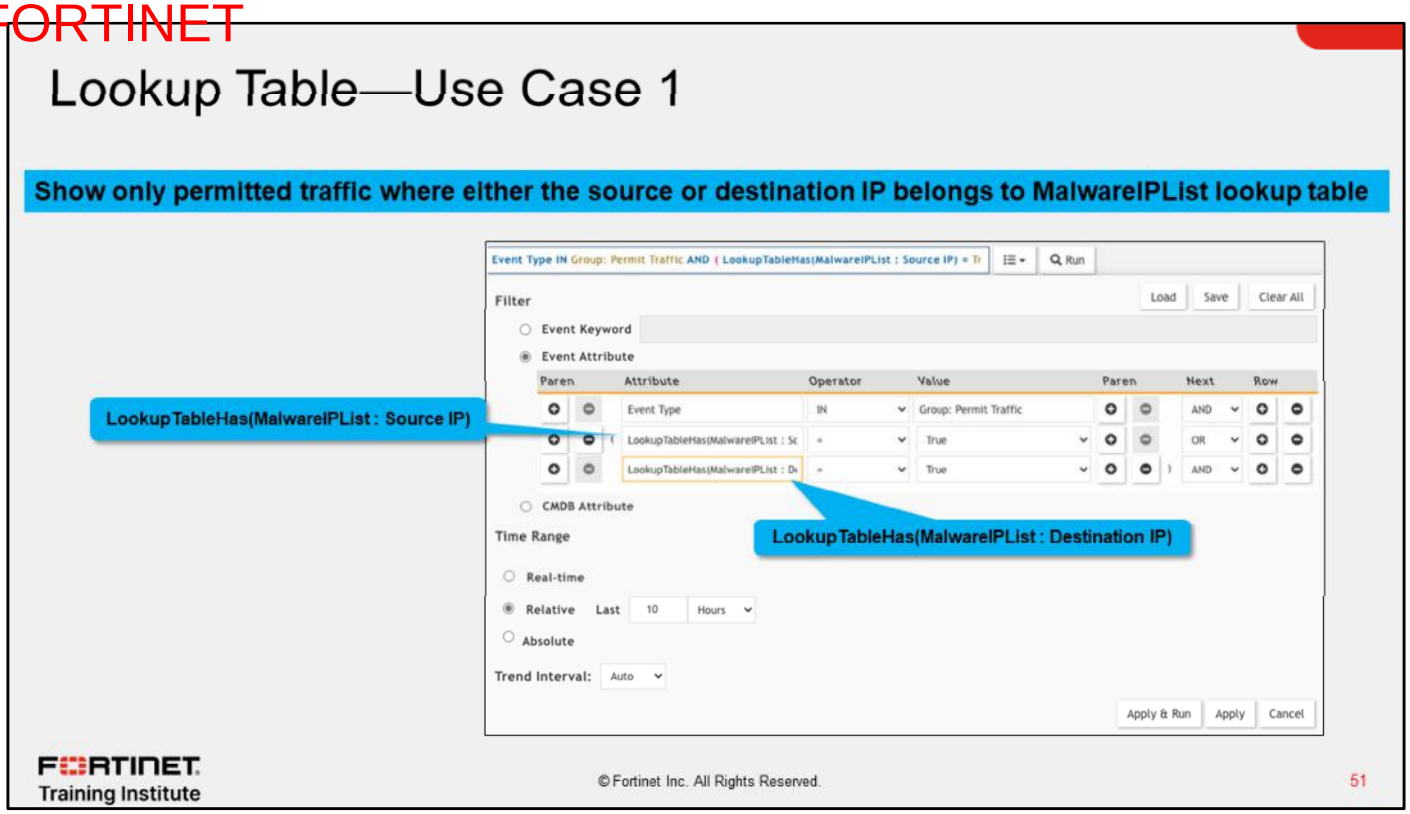
DO NOT REPRINT
© FORTINET



Consider an example of creating a malware IP lookup table manually in the GUI with IP, category, country and confidence information and importing a CSV file with a list of such IPs.

After you configure the lookup table **Malware IP List**, you can analyze permitted traffic events from the firewall and see how you can leverage the new lookup table in a few use cases.

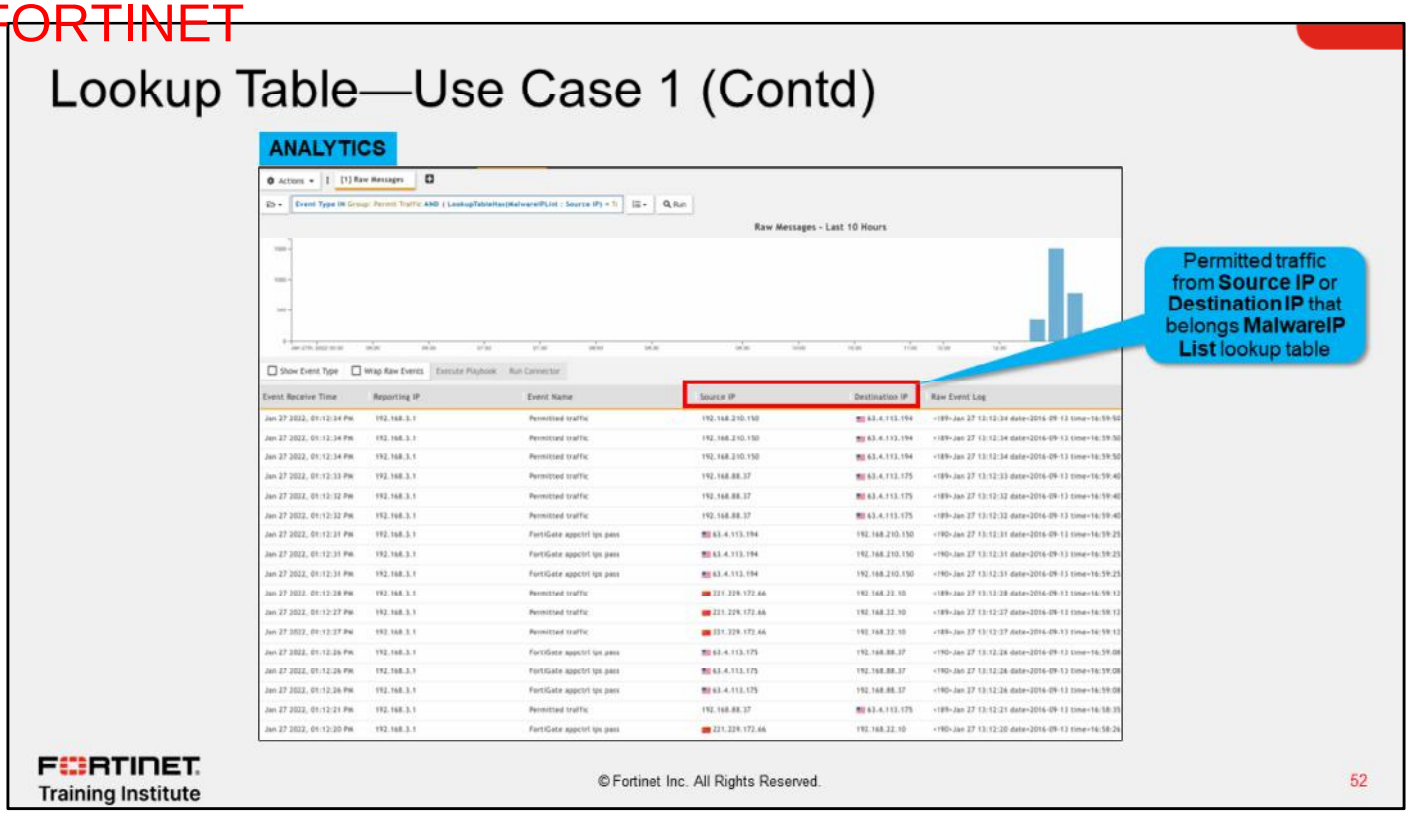
DO NOT REPRINT
© FORTINET



Consider a use case where you want to see only permitted traffic from either the source or destination IP that belongs to the **MalwareIPList** lookup table you configured.

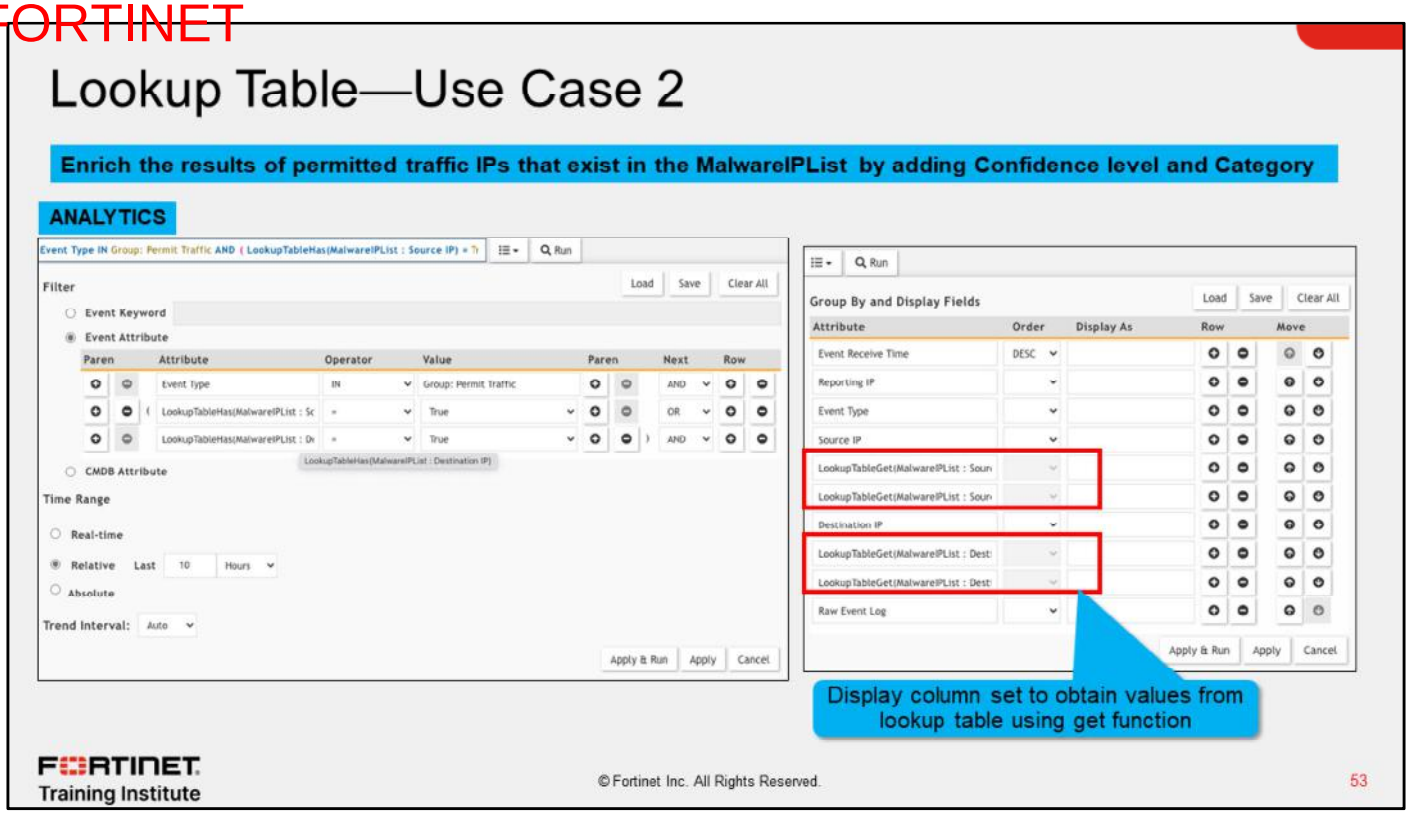
You must specify in the query filter that you want to return results only if either the **Source IP** or **Destination IP** are in the **MalwareIPList** lookup table. You use the lookup table function **LookupTableHas(MalwareIPList : Source IP) = True OR LookupTableHas(MalwareIPList : Destination IP) = True** as this slide shows.

DO NOT REPRINT
© FORTINET



When you run the query with lookup table functions, it returns only events where either the source or destination IP are in the lookup table.

DO NOT REPRINT
© FORTINET



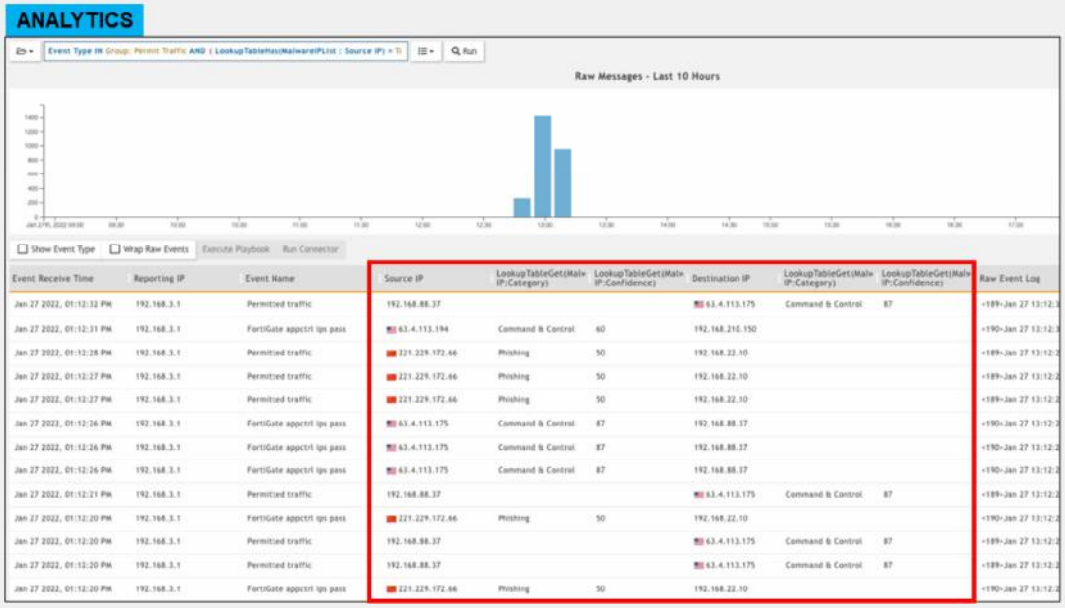
First, specify in the query filter that you only want to return results if either the **Source IP** or **Destination IP** are in the **MalwareIPList** lookup table. Use the function **LookupTableHas(MalwareIPList : Source IP) = True OR LookupTableHas(MalwareIPList : Destination IP)** for this scenario.

To add the category and confidence level of IPs, set the display column to use the following get functions for the lookup table :

- `LookupTableGet( MalwareIPList : Source IP : Category)`
- `LookupTableGet( MalwareIPList : Source IP : Confidence)`
- `LookupTableGet( MalwareIPList : Destination IP : Category)`

DO NOT REPRINT
© FORTINET

Lookup Table—Use Case 2 (Contd)



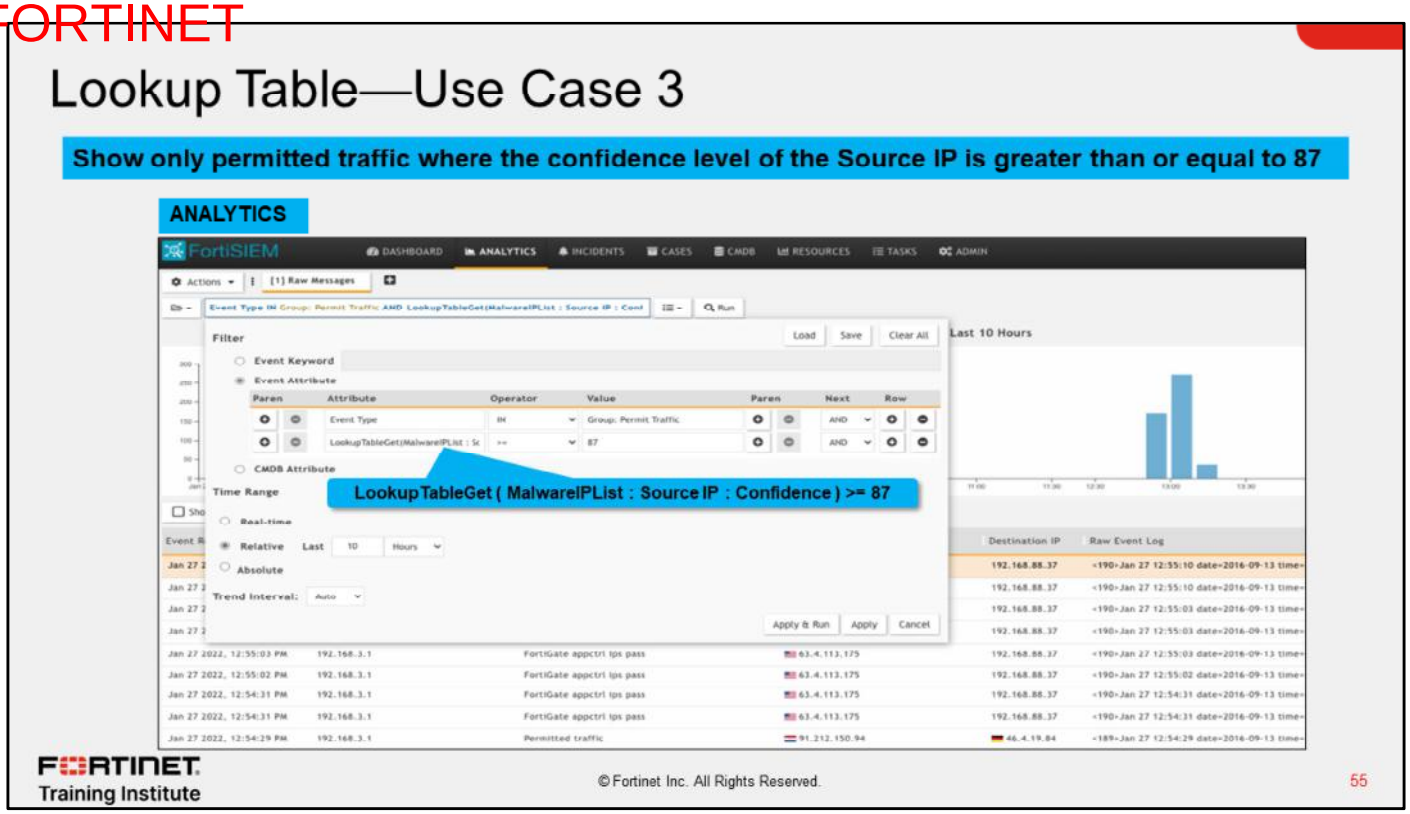
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

54

When you run the query, you should see that the **Source IP** and **Destination IP** addresses are enriched with information from the confidence level and category of IP addresses from the **MalwareIPList** lookup table.

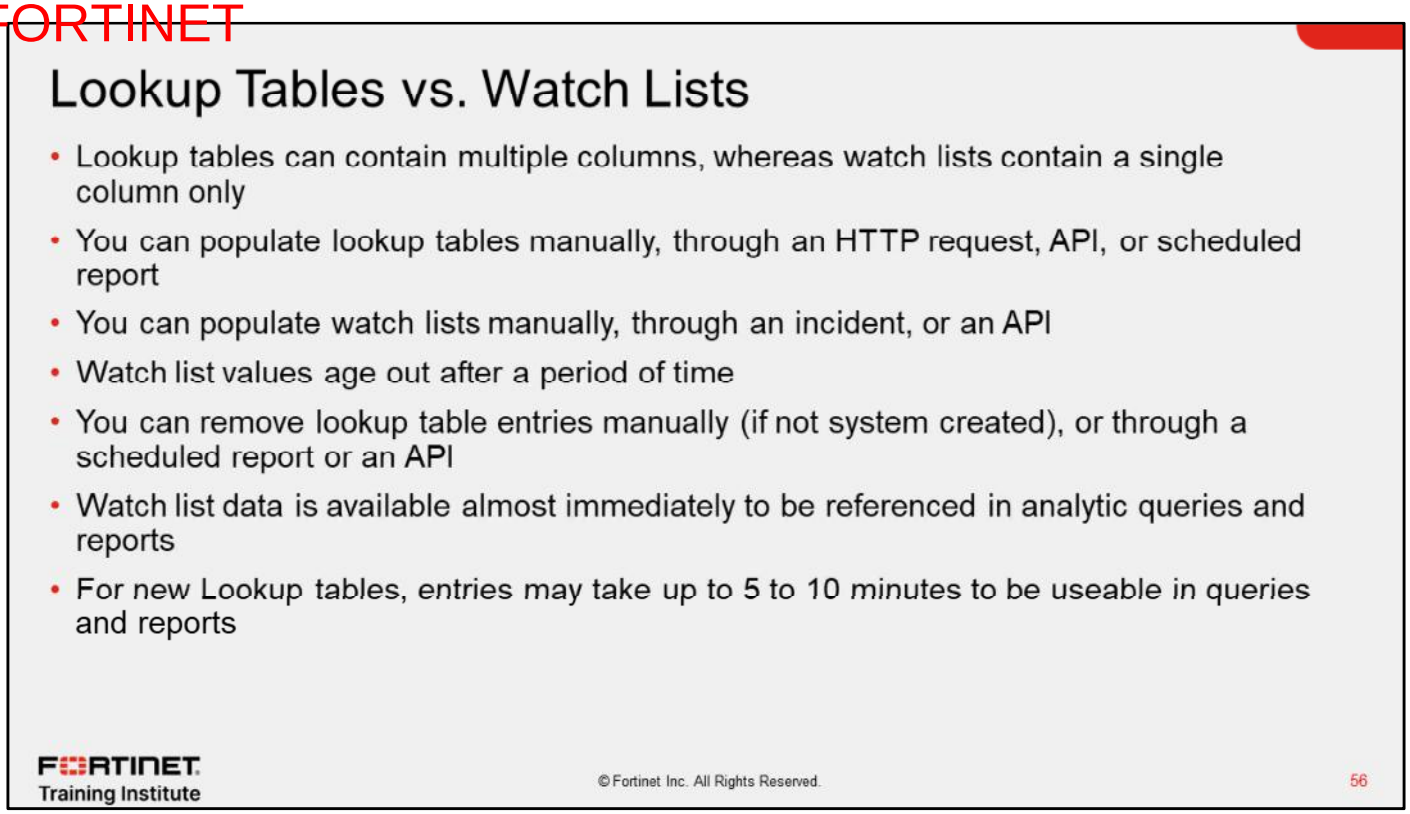
DO NOT REPRINT
© FORTINET



Consider another use case if you want only permitted traffic where the confidence level of the **Source IP** is greater than or equal to **87**.

To achieve this use case, you must set the search filter to leverage the `LookupTableGet` function by specifying `LookupTableGet ( MalwareIPList : Source IP : Confidence ) >= 87`. When you run the query, you should see only the **Source IP** with a confidence level of **87**.

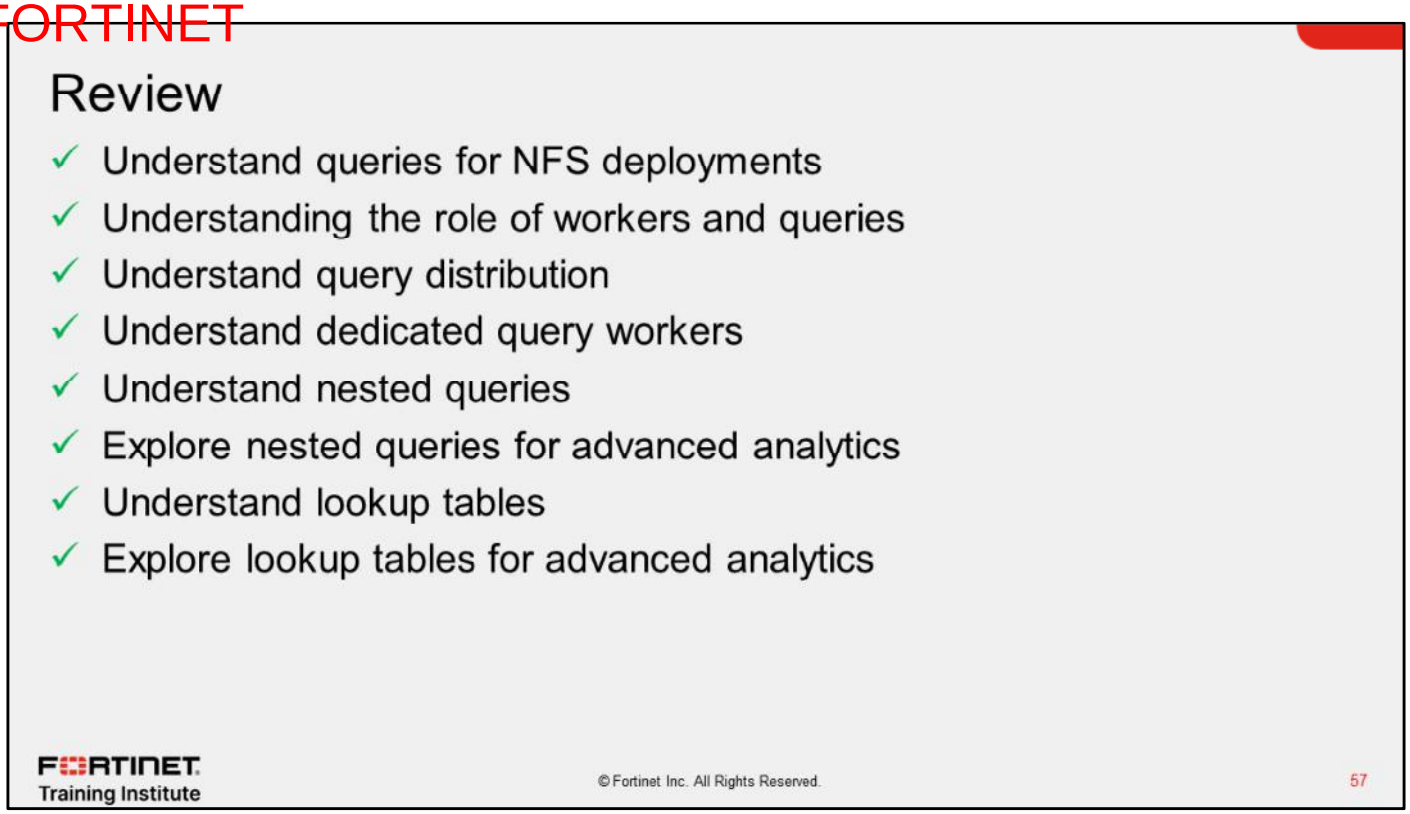
DO NOT REPRINT
© FORTINET



A brief comparison between lookup tables and watchlists:

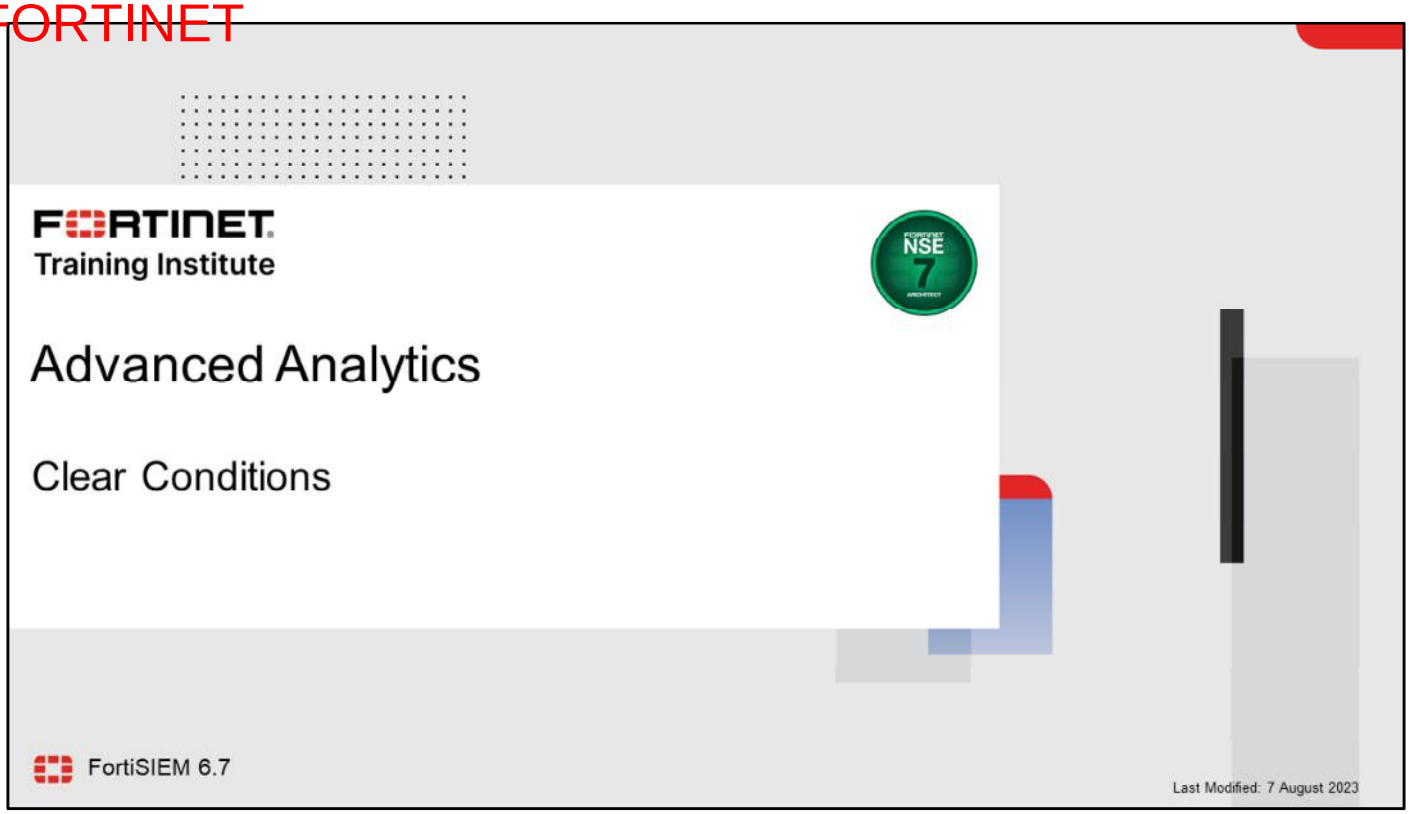
- Lookup tables can contain multiple columns, whereas watch lists contain a single column only.
- You can populate lookup tables manually, or through an http request, API, or scheduled report.
- You can populate watch lists manually, through an incident, or through an API.
- Watch list values age out after a period of time.
- You can remove lookup table entries (if not system created), or through a scheduled report or an API.
- Watch list data is available almost immediately to be referenced in analytic queries and reports.
- For new lookup tables, entries may take up to 5 to 10 minutes to be useable in queries and reports.

DO NOT REPRINT
© FORTINET



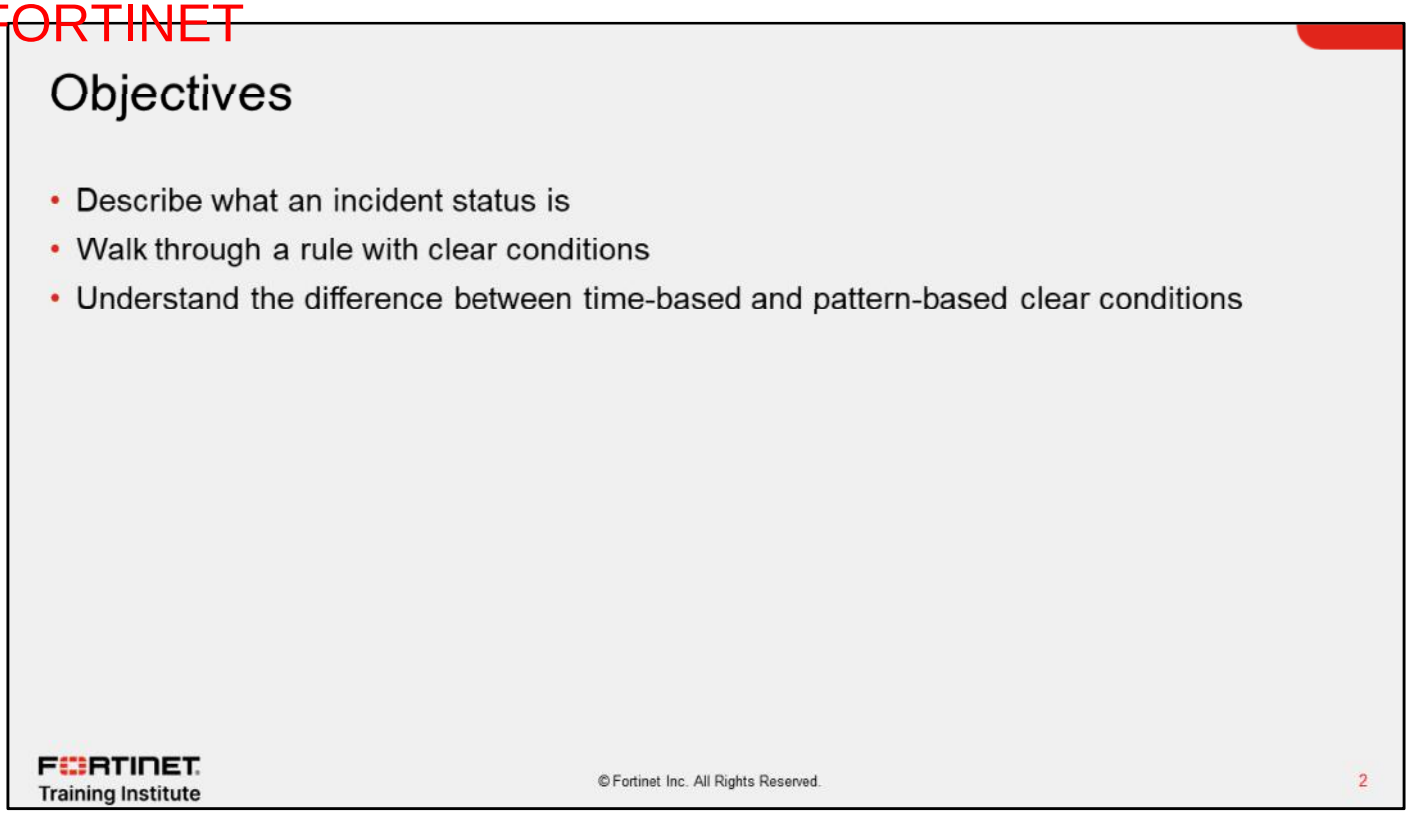
This slide shows the objectives that you covered in this lesson. By mastering the objectives covered in this lesson, you learned how to use the nested queries and lookup tables for advanced analytics.

DO NOT REPRINT
© FORTINET



In this lesson, you will learn about the clear conditions that the system requires to clear an incident automatically.

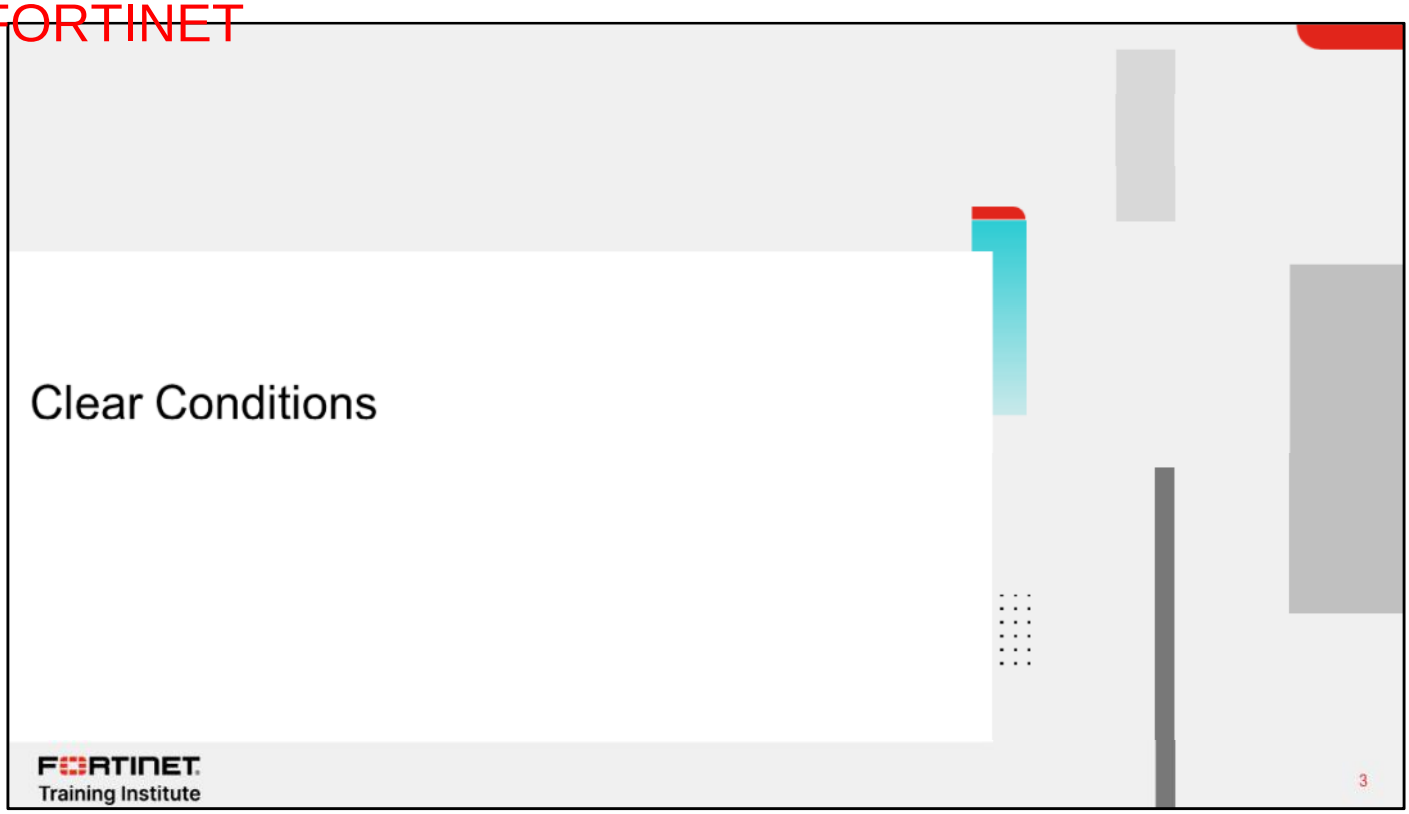
DO NOT REPRINT
© FORTINET



After completing this lesson, you should be able to achieve the objectives shown on this slide.

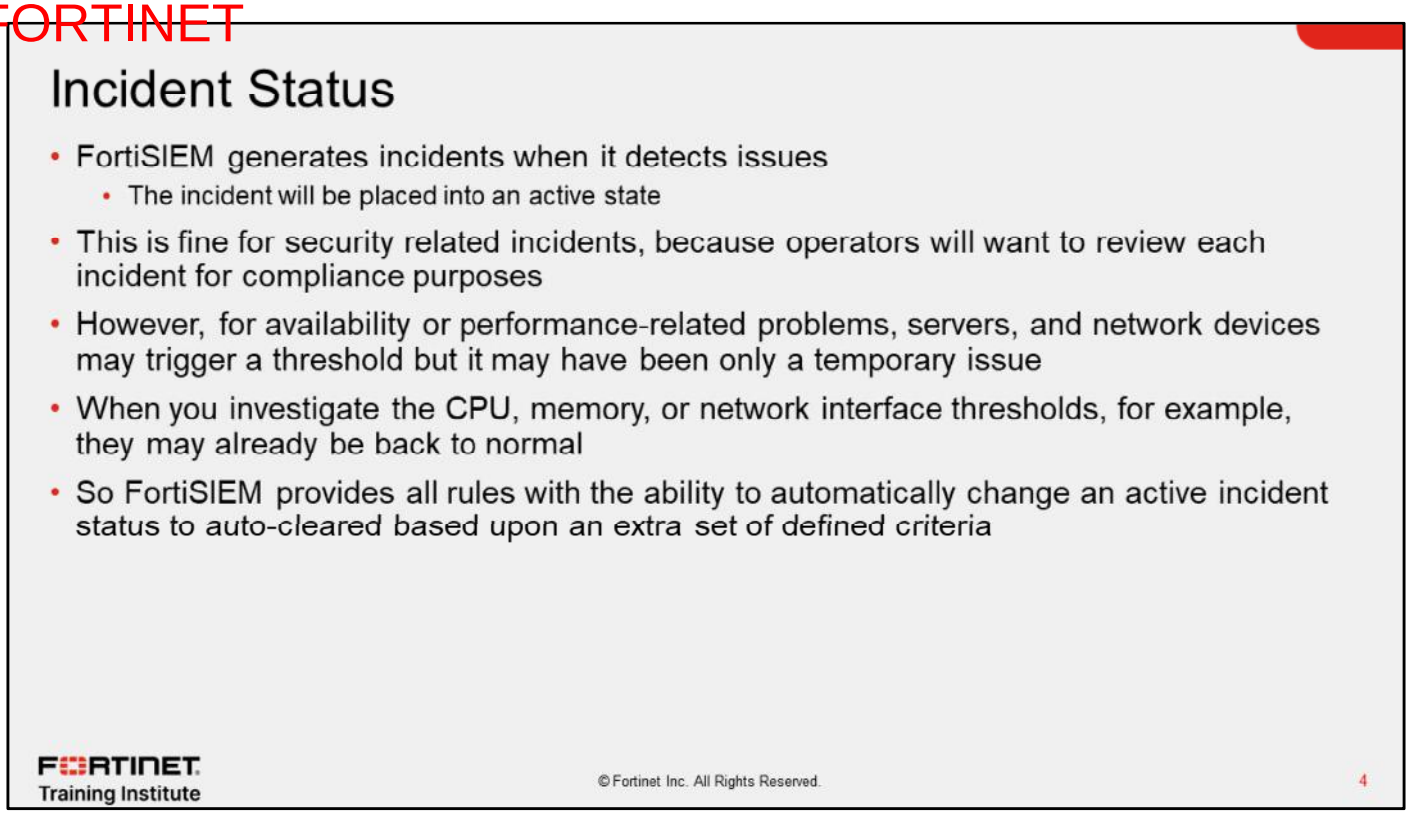
By demonstrating competence in understanding the generation of incident procedures, you will be able to define conditions the system requires to clear an incident automatically.

DO NOT REPRINT
© FORTINET



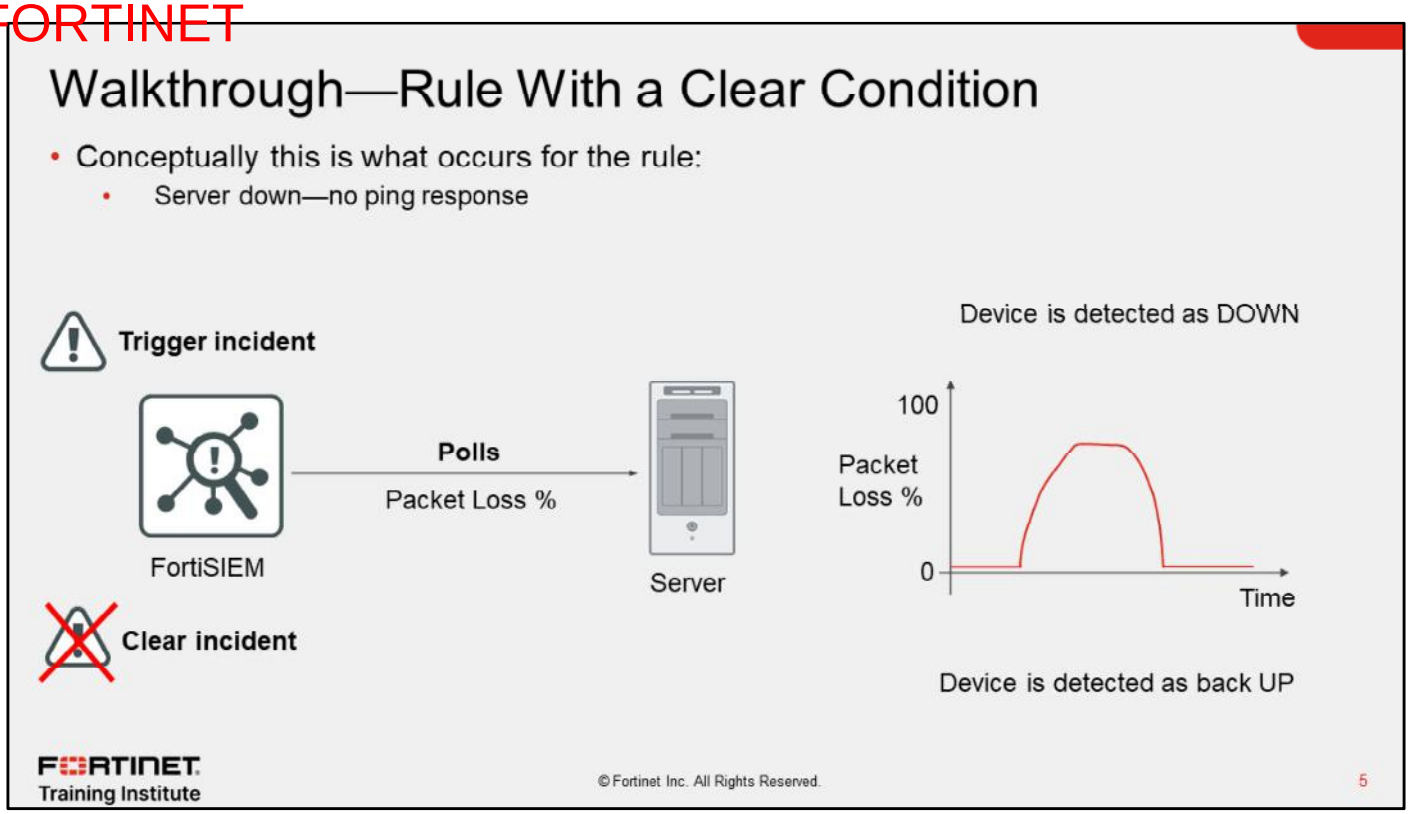
In this section, you will learn about incident status and how you can clear incidents by defining clear conditions. You will also walk through the process of clearing an incident for a server that is not responding to a ping response.

DO NOT REPRINT
© FORTINET



FortiSIEM generates incidents when it detects issues. The incident will be placed into an active state. This is fine for security-related incidents because operators will want to review each incident for compliance purposes. However, for availability or performance-related problems, servers and network devices may experience only temporary issues. For example, when you come to investigate the CPU, memory, or network interface thresholds, they may be back to normal already. Therefore, FortiSIEM provides all rules, with the ability to automatically change an active incident status to auto-cleared, based on an extra set of defined criteria.

DO NOT REPRINT
© FORTINET

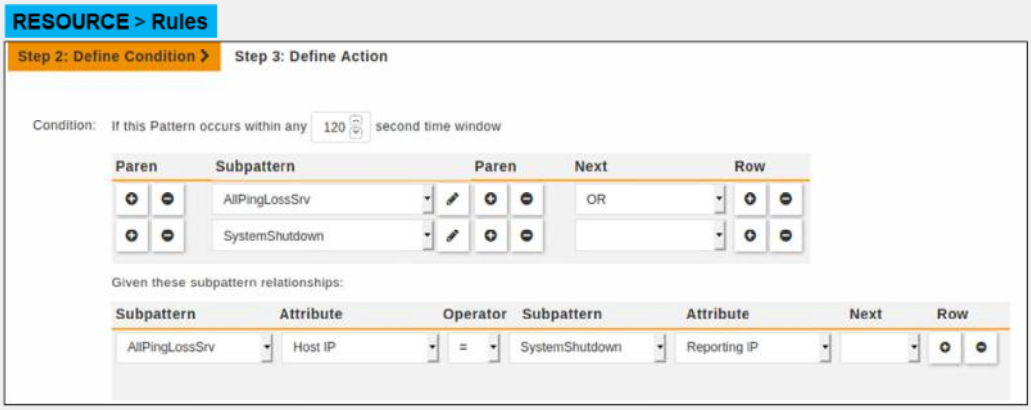


This slide shows an example of a rule with a defined clear condition. This rules is tracking a server using ICMP packets, and when there is no response to ICMP packets from the server, FortiSIEM triggers an incident. The status of the incident is set as active. After the server comes back up and starts responding to ICMP packets again, the incident will be cleared automatically by the system because of the defined clear condition.

DO NOT REPRINT
© FORTINET

Server Down Rule—No Ping Response

- Consists of two subpatterns with an OR condition



- The rule is evaluated over a 2-minute time window

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

6

In the following slides, you will walk through the rule that detects if a device is responding to pings. During a 120-second window, if 10 out of 10 ping packets are lost, then either the host is down or there is a routing problem.

There are two subpatterns for this rule. If any one of the subpatterns meets the conditions, then the rule will trigger an incident.

DO NOT REPRINT
© FORTINET

AllPingLossSrv Subpattern

- Filters for the PING STAT event reported by server devices only
- Groups the results by the same host IP and host name
- Aggregates based on receiving one or more events where the average packet loss is 100%—meaning the device is down

RESOURCE > Rules

Name: AllPingLossSrv

Filters:

Paren	Attribute	Operator	Value	Paren	Next	Row
	Event Type	=	PH_DEV_MON_PING_STAT		AND	
	Host IP	IN	Devices: Server		AND	

Aggregate:

Paren	Attribute	Operator	Value	Paren	Next	Row
	AVG(Packet Loss Pct)	=	100		AND	
	COUNT(Matched Events)	>=	1		AND	

Group By:

Attribute	Row	Move
Host IP		
Host Name		

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

7

Take a look at the subpattern shown on this slide. The filter is tracking event types, which is specifically for monitoring device performance and availability using pings. The filter also ensures that the ping monitoring events are coming from server devices.

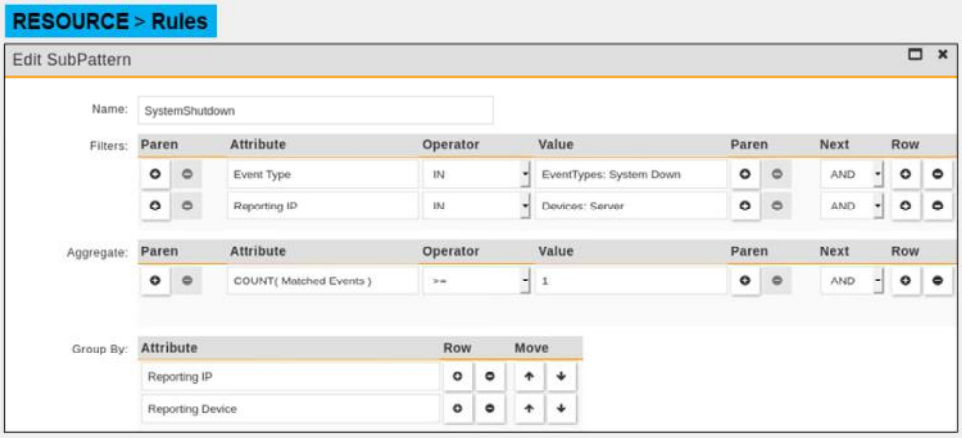
FortiSIEM groups all the events by their host IP and host name, which are the server IP and server name.

After the events are grouped, FortiSIEM performs the aggregate function. If the average packet loss percentage is equal to 100% and the matched events count is greater than or equal to 1, then FortiSIEM triggers an incident.

DO NOT REPRINT
© FORTINET

System Shutdown Subpattern

- Filters any events reported by a server device that are categorized as system down
- Groups the results by the same reporting IP and reporting device
- Aggregates based on receiving one or more events



FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

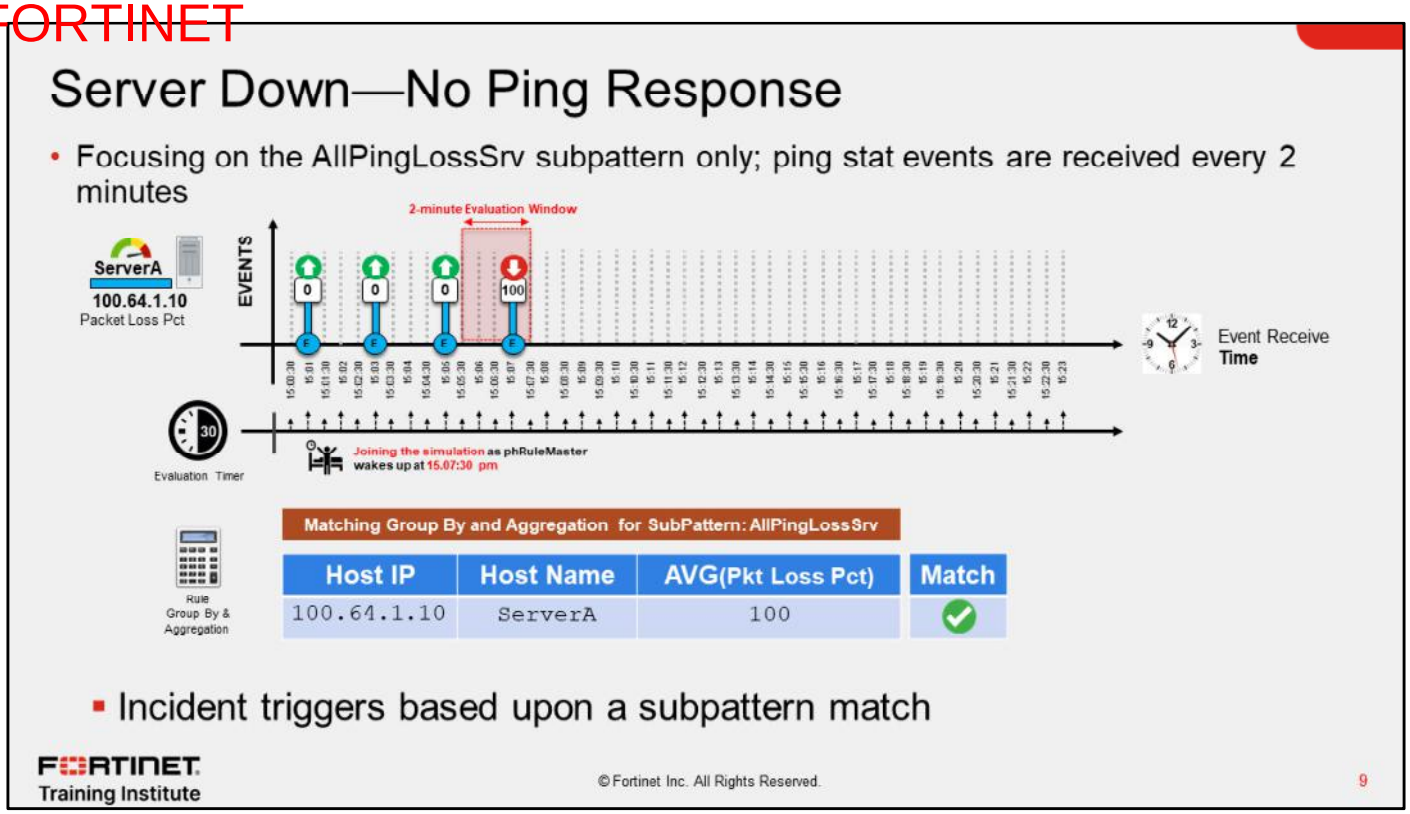
8

Take a look at the system shutdown subpattern shown on this slide. The filter is tracking event types that belongs to the system down event type category, which is defined in the CMDB. The filter also ensures that the reporting IP belongs to the server devices.

FortiSIEM groups all the events by their reporting IP and reporting device, which are the server IP and server name.

After the events are grouped, FortiSIEM will perform the aggregate function. If the matched events count is greater than or equal to 1, then FortiSIEM will trigger an incident.

DO NOT REPRINT
© FORTINET



Let's focus on the pattern shown on this slide.

The pattern on this slide shows that FortiSIEM will poll the server every 2 minutes and collect the ping stats. Based on those ping stats it creates an event. If only 5 out of 10 pings fail, then it is considered a 50% packet loss and FortiSIEM will not generate an incident.

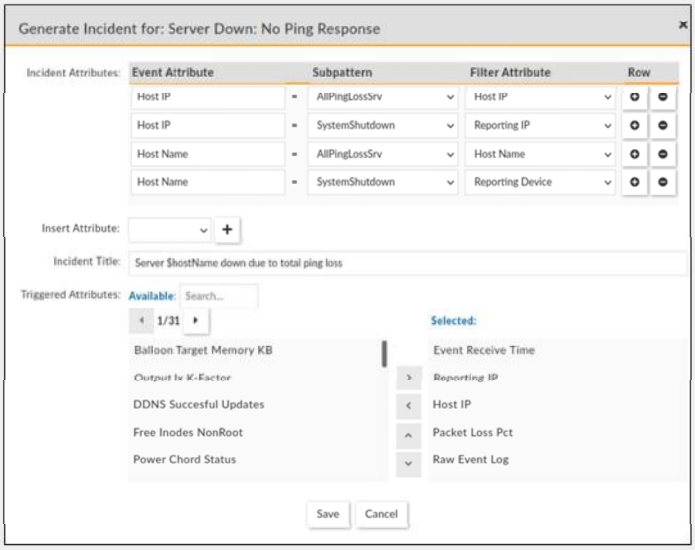
In this example, it's a 100% packet loss, which means all 10 pings failed, which will trigger an incident.

DO NOT REPRINT
© FORTINET

Server Down—No Ping Response Incident

- The incident action is based upon filter attributes in the **AllPingLossSrv** subpattern

RESOURCE > Rules



FORTINET Training Institute

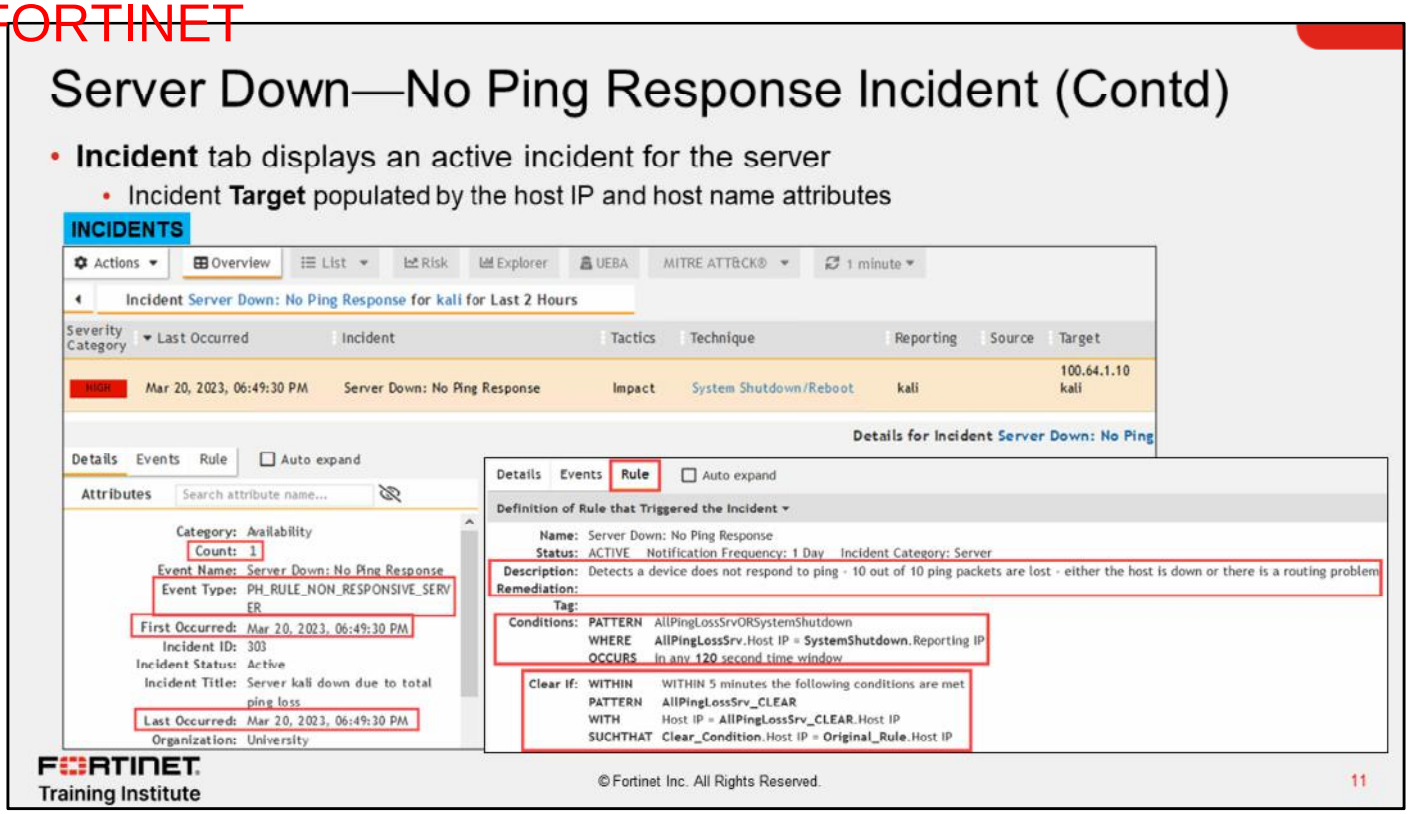
© Fortinet Inc. All Rights Reserved.

10

The action on the rule generates the incident that is based upon the filter attributes in the subpattern.

In the example shown on this slide, the **Host IP** and **Reporting IP** defined in the subpattern filter are the same so they are mapped to the incident as the **Host IP**. The **Host Name** and the **Reporting Device** defined in the subpattern filter are similar, so they are mapped to the incident as the **Host Name**.

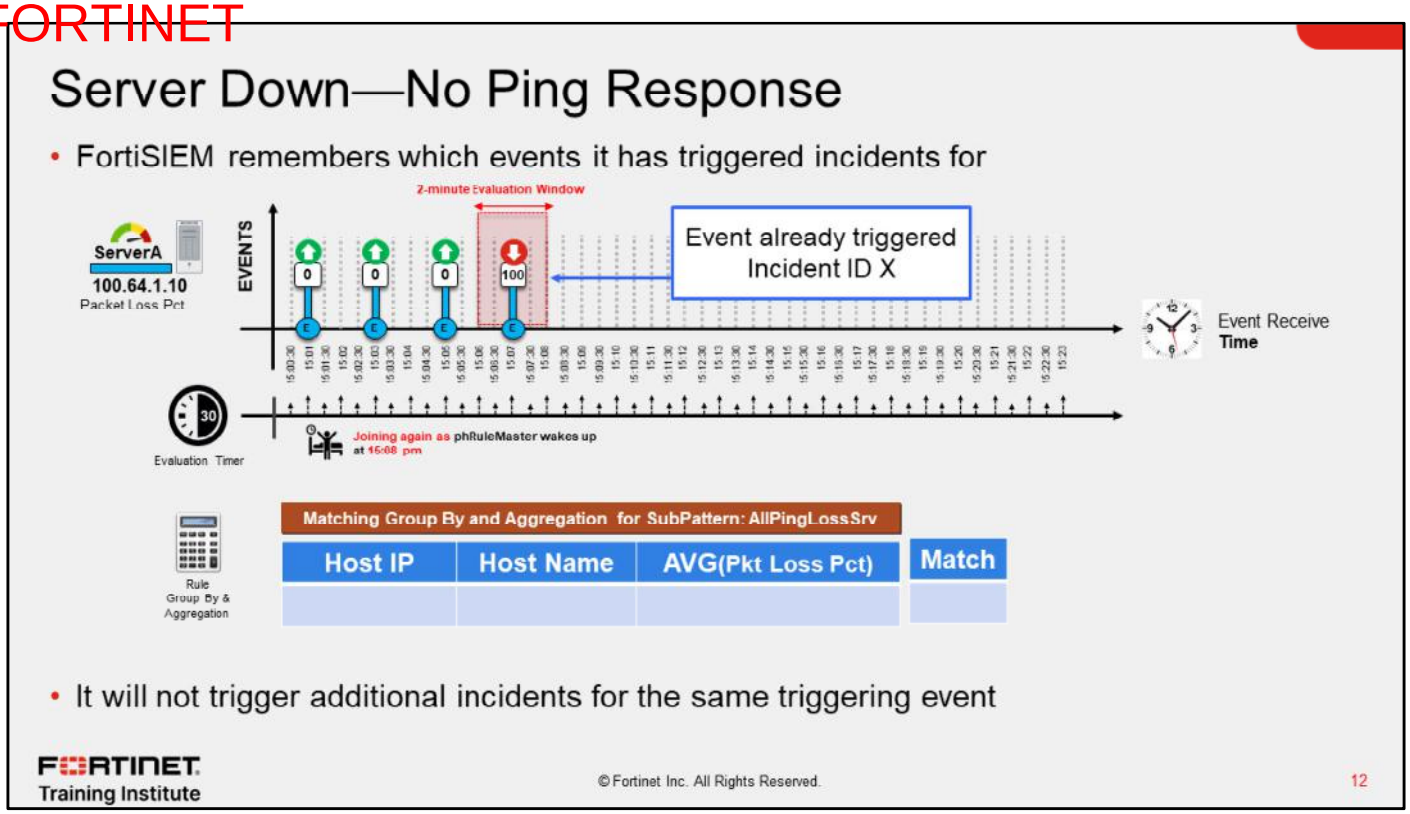
DO NOT REPRINT
© FORTINET



The incident is generated when there is no ping response from the server. From the **Incident Type** you can identify which subpattern triggered the incident. In the example shown on this slide, the incident type is **PH_RULE_NON_RESPONSIVE_SERVER**, which means that the **AllPingLossSrv** subpattern triggered the incident.

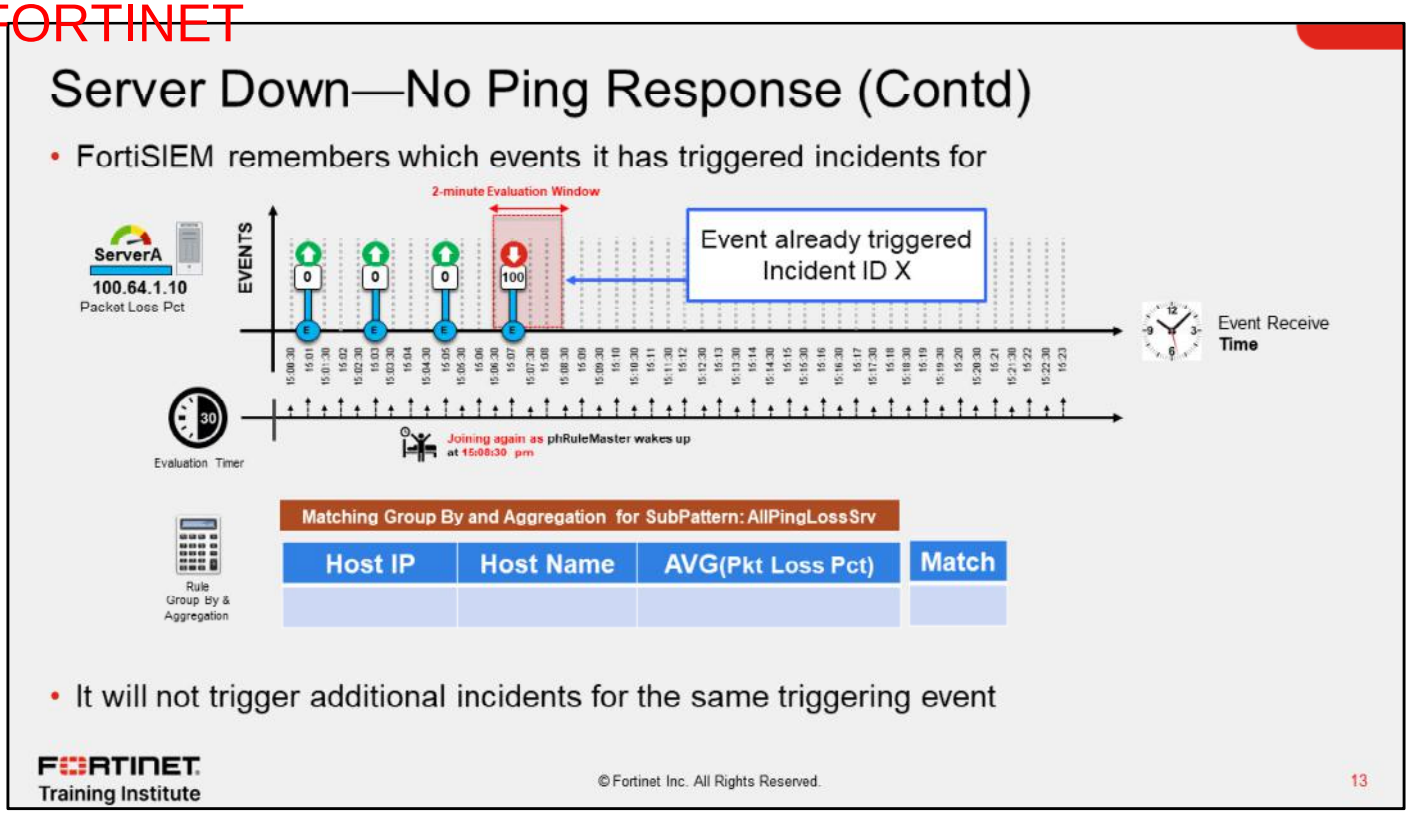
Because this is the first occurrence of the incident, the count will be 1 and the **First Occurred** and **Last Occurred** times will be the same.

DO NOT REPRINT
© FORTINET



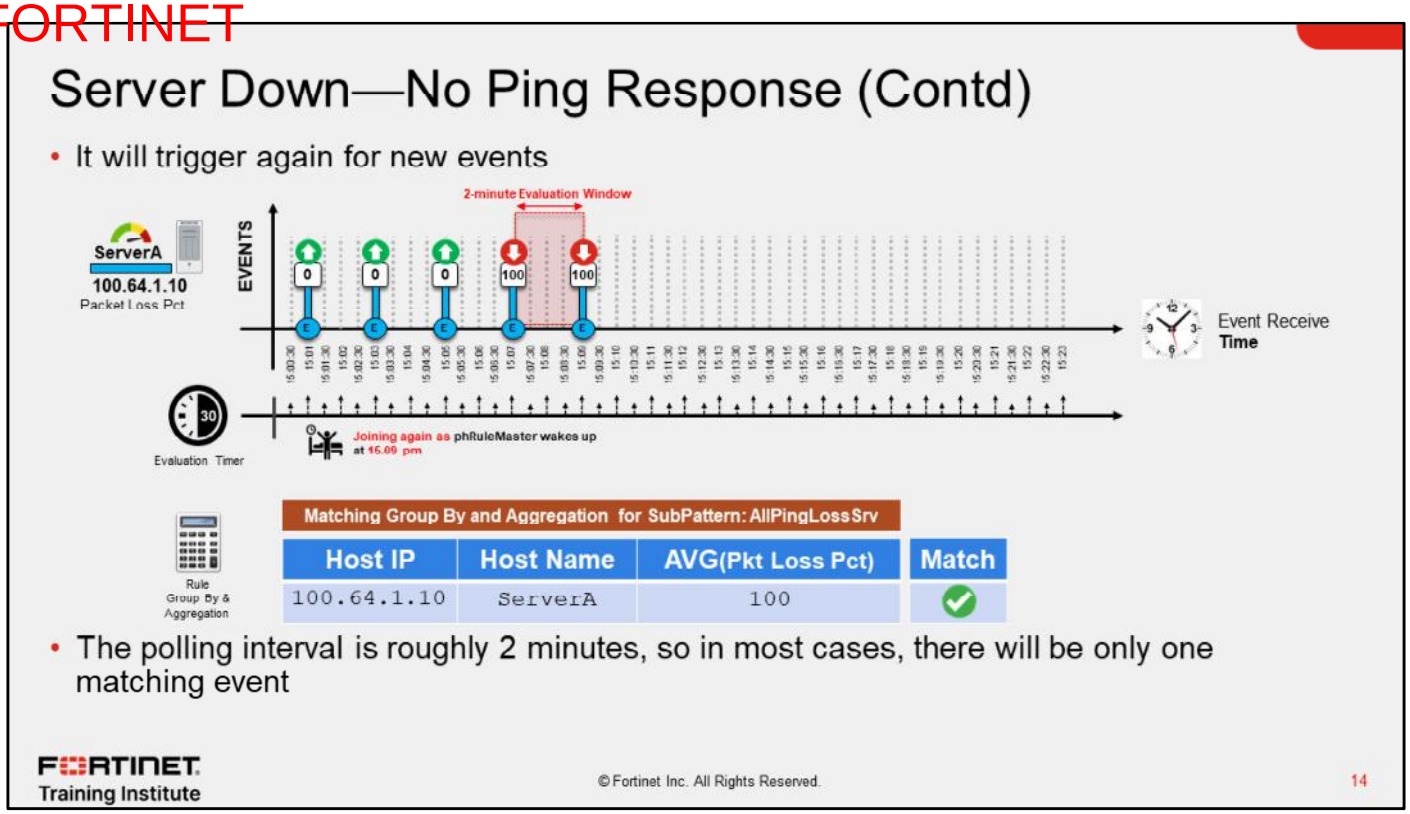
FortiSIEM remembers which events the incident was triggered for. When the `phRuleMaster` wakes up 30 seconds later and evaluates the 2-minute window, it will not trigger additional incidents for the same triggering event.

DO NOT REPRINT
© FORTINET



When the `phRuleMaster` wakes up again after 30 seconds and evaluates the 2-minute window, it will not trigger additional incidents for the same triggering event.

DO NOT REPRINT
© FORTINET



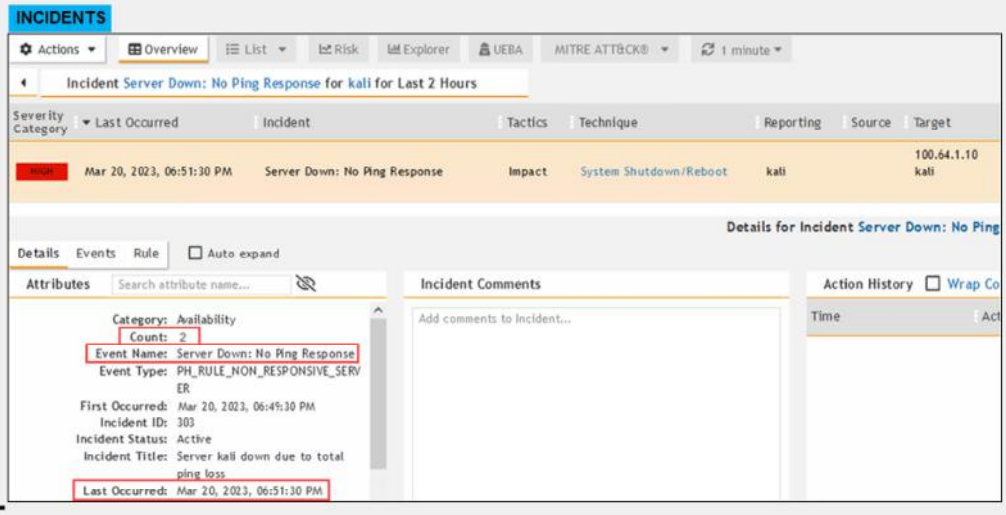
Because this is a performance-related metric, FortiSIEM polls the server periodically and, in the example shown on this slide, it is 2 minutes. The evaluation window is also for 2 minutes and for that reason, in most cases, there will be only one matching event.

When FortiSIEM polls the server again in 2 minutes, a new event will be generated with 100% average packet loss. When the `phRuleMaster` wakes up again and evaluates the 2-minute window, it will trigger the incident caused by the new triggering event. The incident ID will be the same as before. It will just update the count, last occurred, and event details for the same incident ID.

DO NOT REPRINT
© FORTINET

Server Down—No Ping Response Incident

- The new incident data is merged with the current active incident
  - The **Count** is increased to **2**
  - The **Last Occurred** time is updated



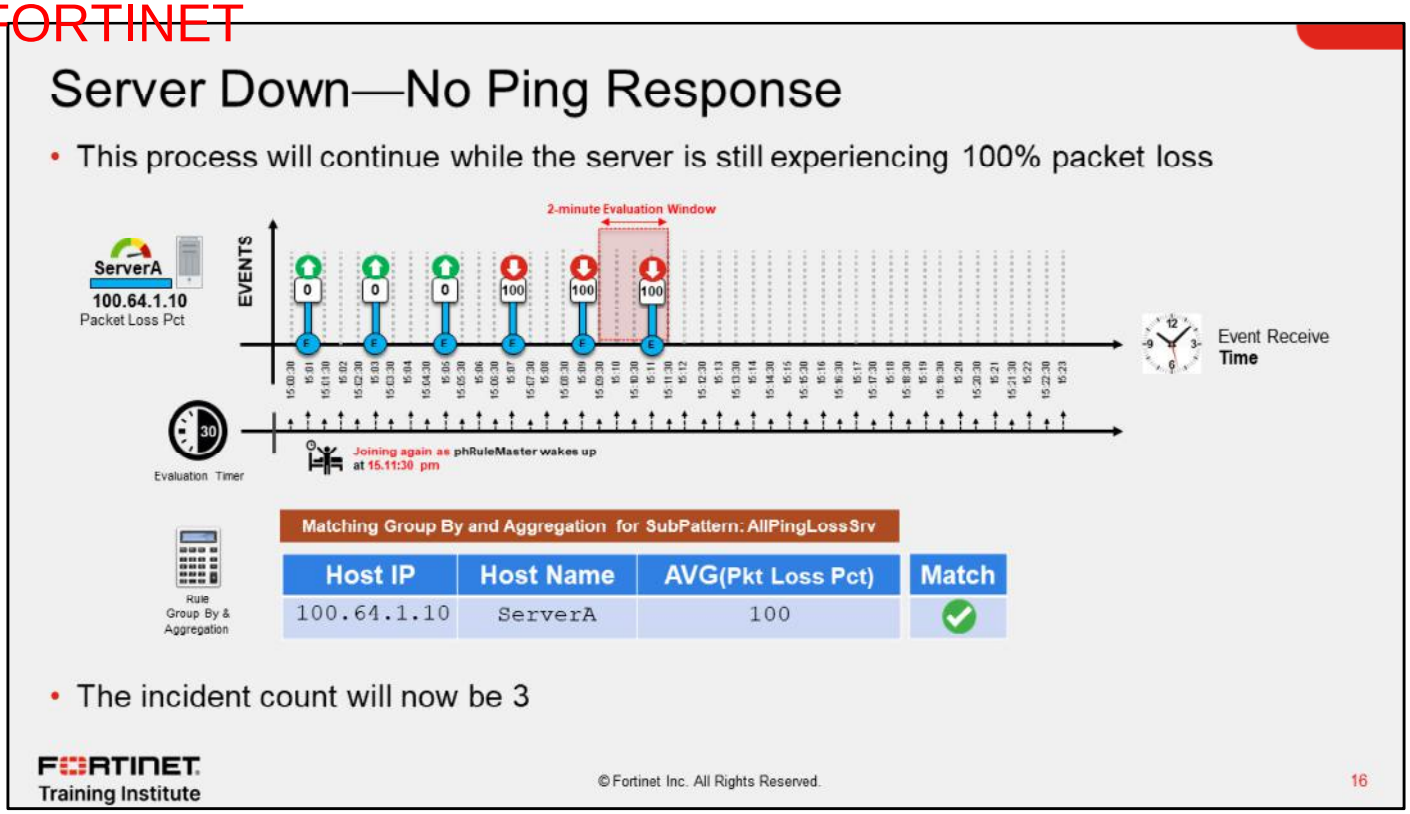
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

15

The new incident data is merged with the current active incident. The **Count** is increased to **2** and the **Last Occurred** time is updated.

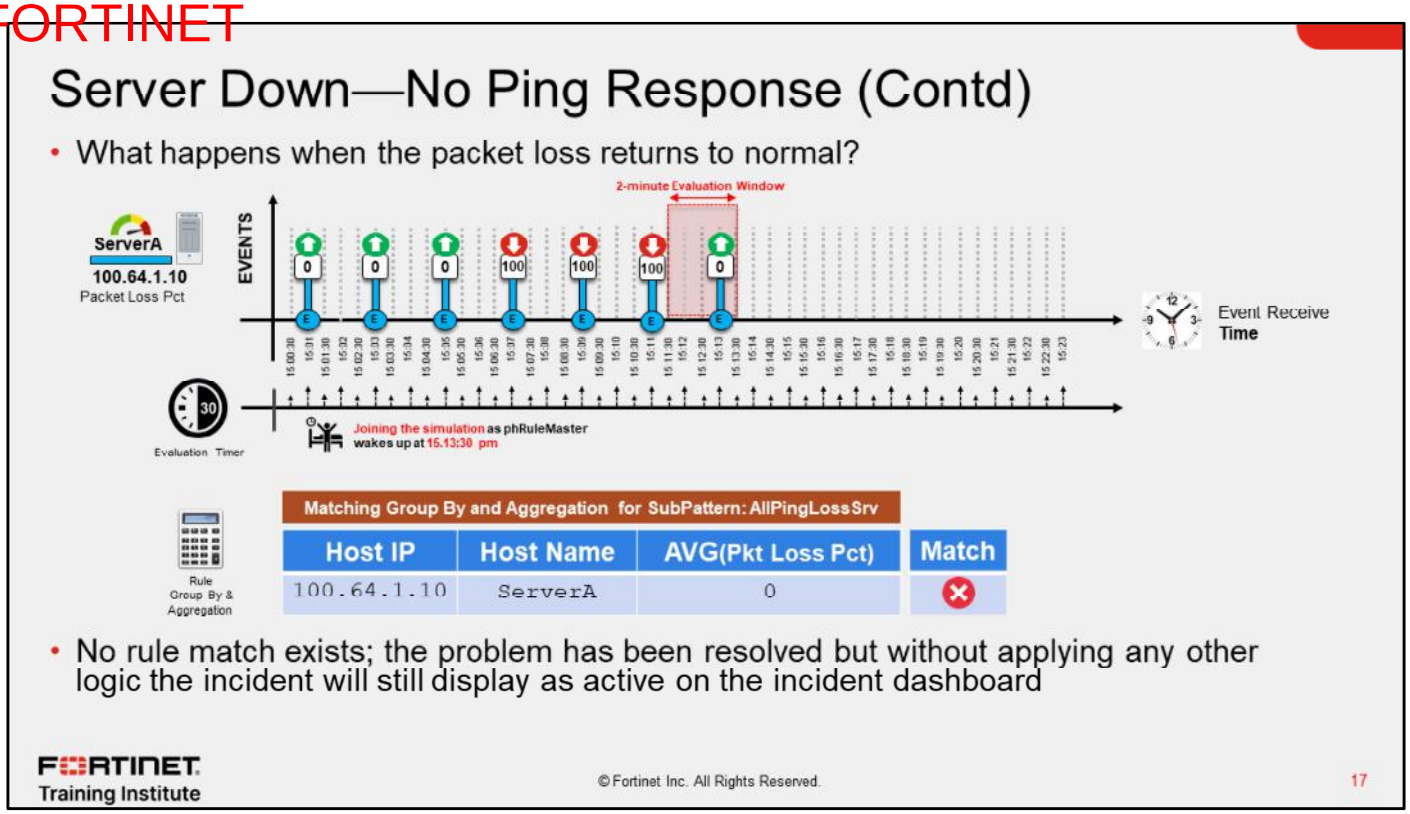
DO NOT REPRINT
© FORTINET



This process will continue while the server is still at 100% packet loss. After 2 minutes, when the `phRuleMaster` wakes up again and evaluates the 2-minute window, it will detect a new triggering event with 100% packet loss. The incident count will be updated to 3.

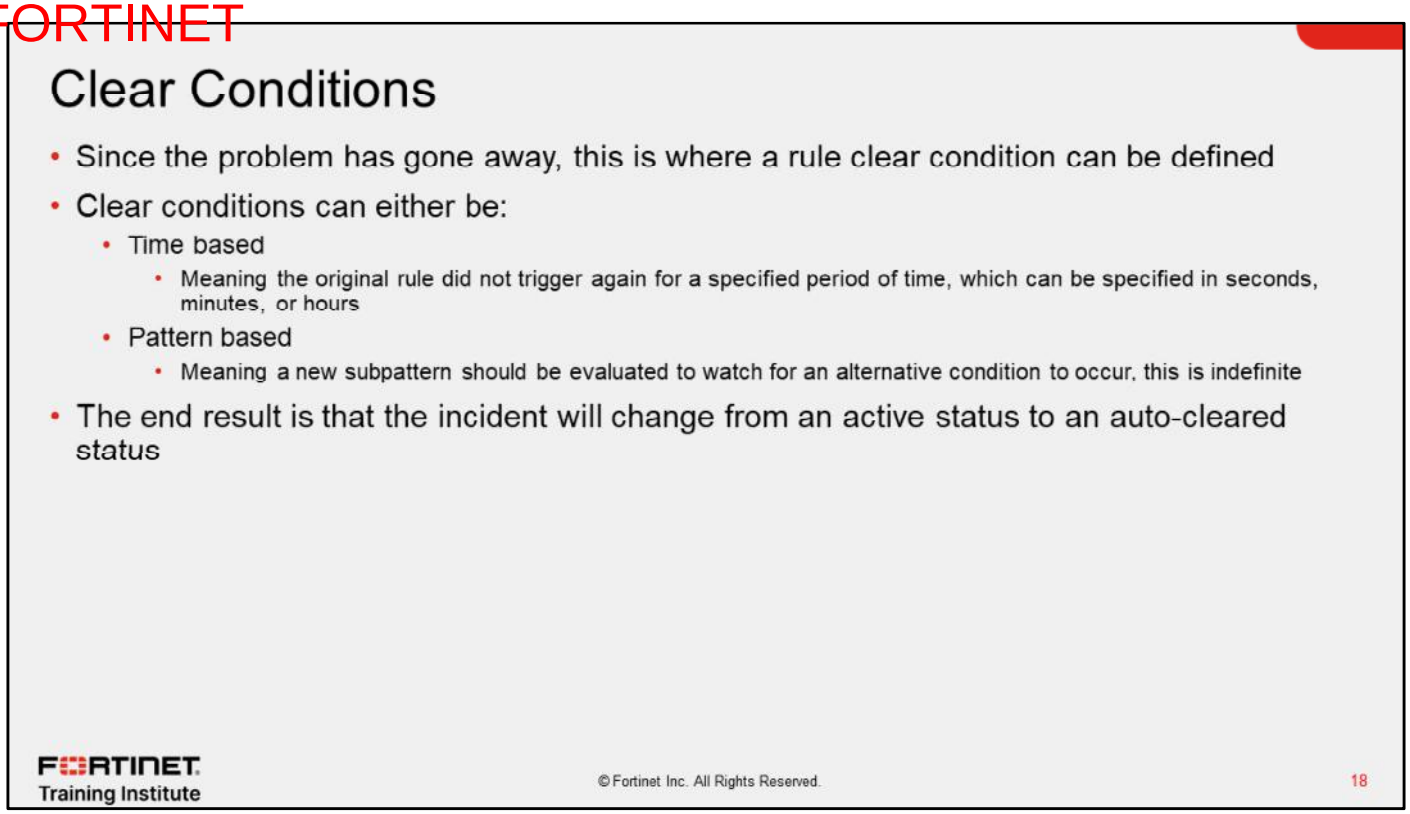
Remember that the `phRuleMaster` is actually waking up every 30 seconds to evaluate all the active rules. It will update the incident only if there is a new triggering event.

DO NOT REPRINT
© FORTINET



Again, after 2 minutes when the `phRuleMaster` wakes up and evaluates the 2-minute window, it will detect a new event with 0% packet loss. At this point, the rule conditions do not match the event and no incident will be triggered or updated. The problem has been resolved, but without applying any other logic, the incident will still display as active on the incident dashboard.

DO NOT REPRINT
© FORTINET



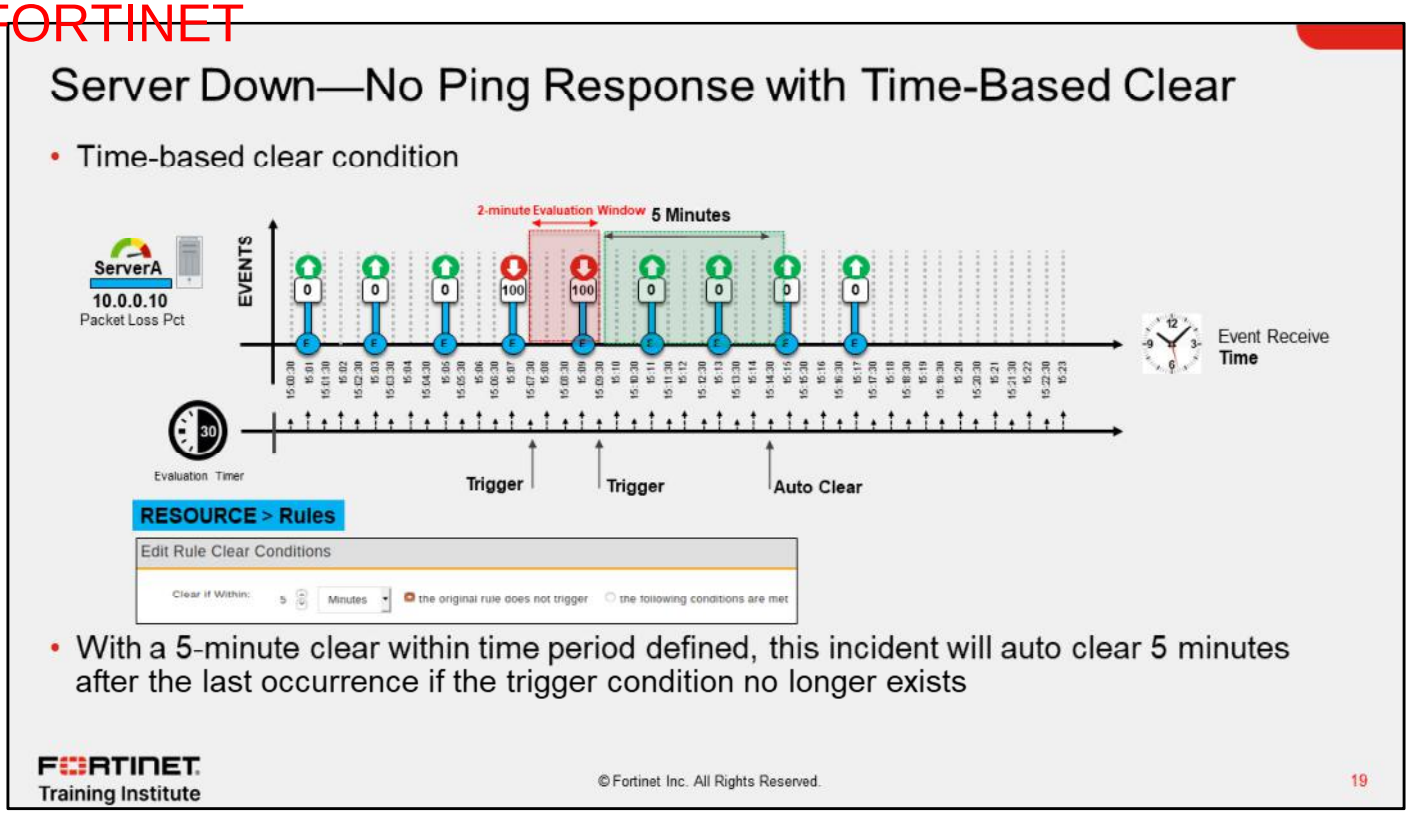
Since the problem has gone away, this is where you can define a rule clear condition. Clear conditions can either be time based or pattern based.

Time based means that the original rule did not trigger again for a specified period of time, which can be specified in seconds, minutes, or hours.

Pattern based means that a new subpattern should be evaluated to watch for an alternative condition to occur.

The end result is that the incident will change from an active status to an auto-cleared status.

DO NOT REPRINT
© FORTINET



If the clear condition is time-based, then with a 5-minute clear window, the incident will auto clear after the last occurrence of the incident. The clearance window begins from the last time the incident was triggered and, if the incident triggers again during that window, then the window is reset.

DO NOT REPRINT
© FORTINET

Pattern-Based Clear Condition

- Define a clear condition subpattern for a pattern-based clear condition



- It could be a single or multiple-pattern definition
- Usually it is almost an exact mirror of the original pattern in the rule, but with a different aggregation calculation

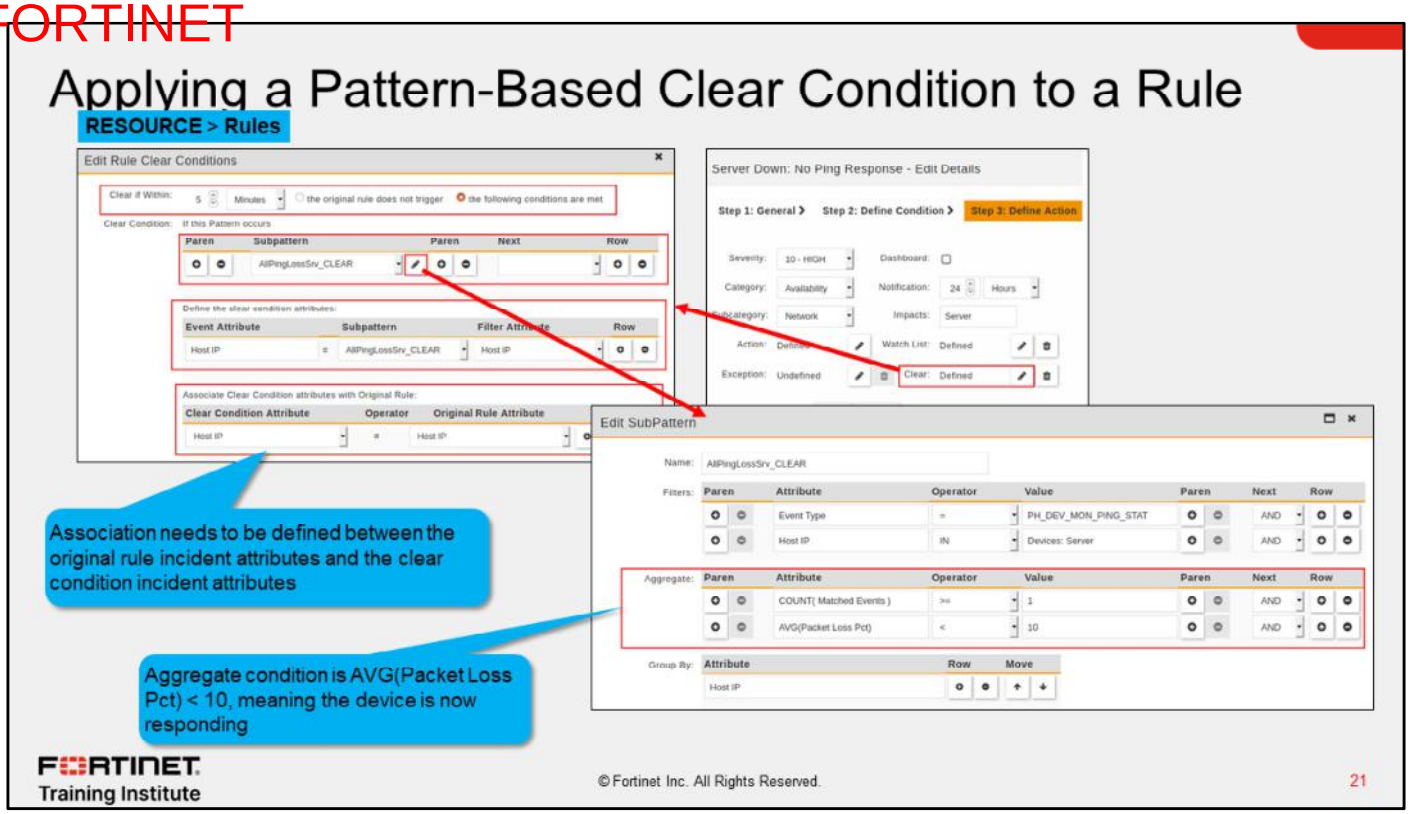
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

20

With a pattern-based clear condition, a subpattern needs to be defined, which could be a single-pattern or multiple-pattern definition. Usually, it is almost an exact mirror of the original pattern in the rule but with a different aggregation calculation.

DO NOT REPRINT
© FORTINET

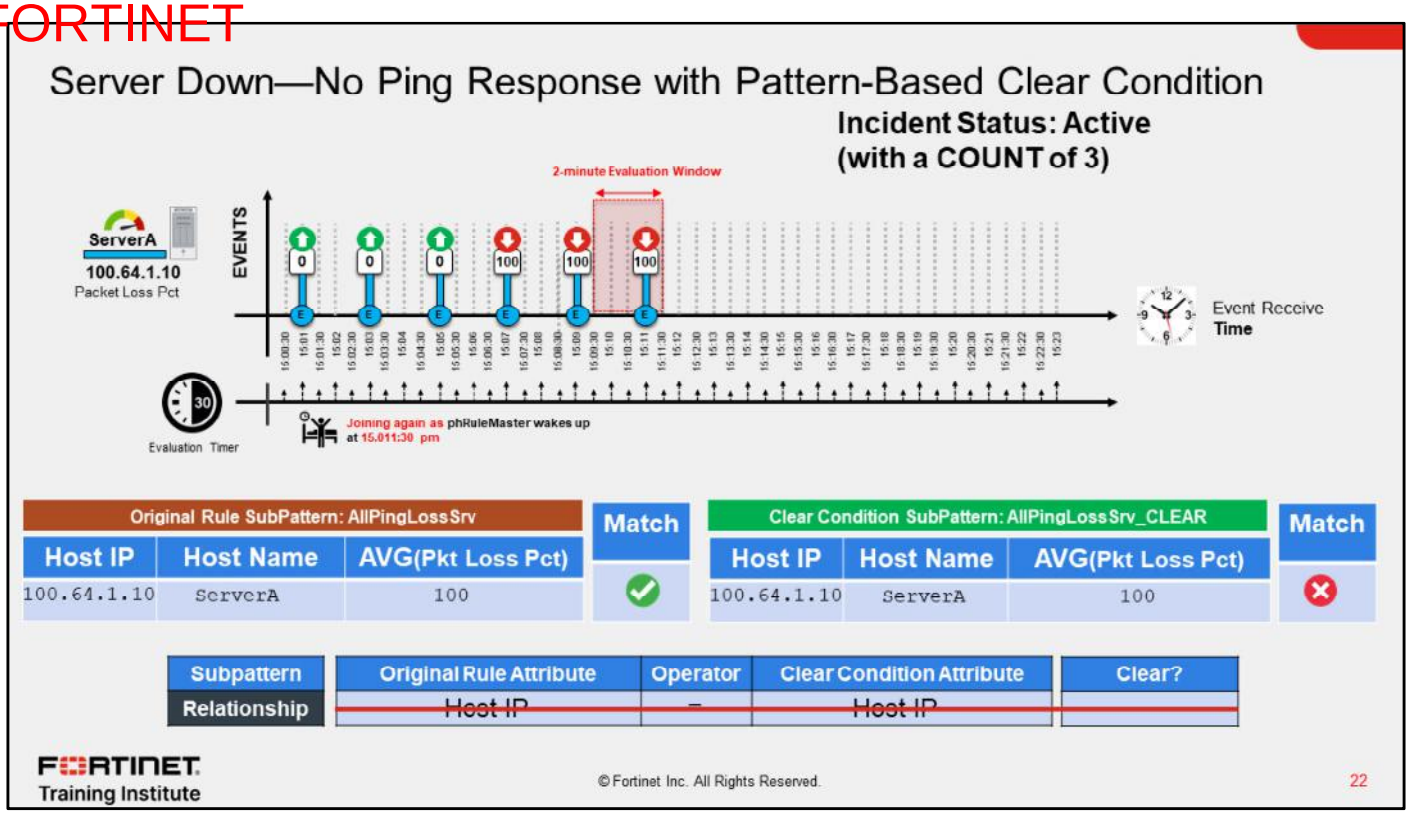


The clear condition is defined using the **Clear** setting in the rule definition. The clear condition is available on all rules but you need to define the conditions. However, availability and performance rules have a default clear condition predefined.

You need to define an association between the original rule's incident attributes and the clear condition incident attributes to ensure that the incident that is being cleared is associated with the original rule.

In the example shown on this slide of a pattern-based clear condition, if the average packet loss is less than 10% and if it matches at least one event, then the incident will be cleared.

DO NOT REPRINT
© FORTINET

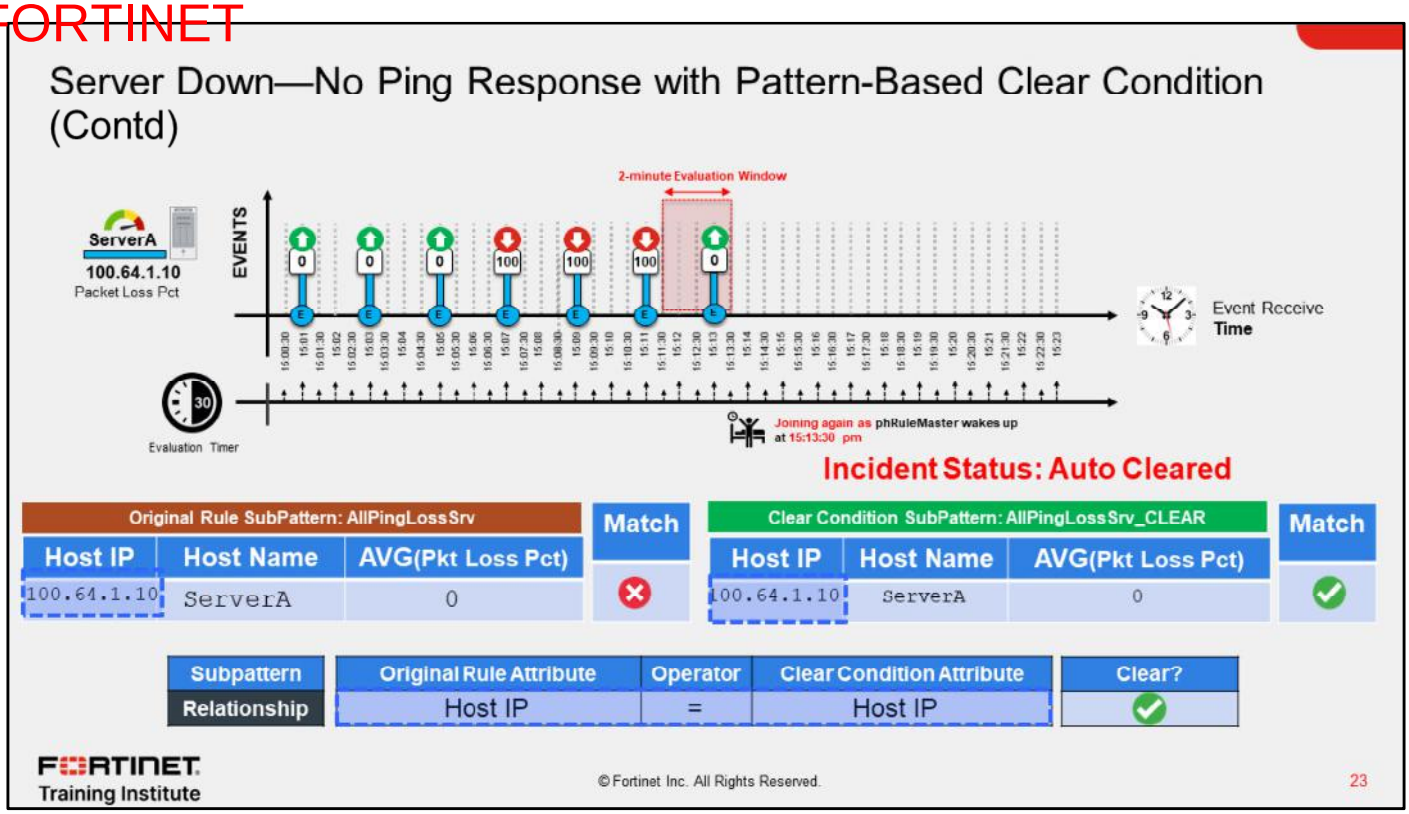


This is a conceptual explanation of a pattern-based clear condition.

The `phRuleMaster` wakes up to evaluate the events for the past 2 minutes. Based on the rule subpattern, there is 100% packet loss. The rule will update the count value of the previous incident because it is the third occurrence of the event.

Since the average packet loss is still at 100%, it does not match the clear condition defined for that rule.

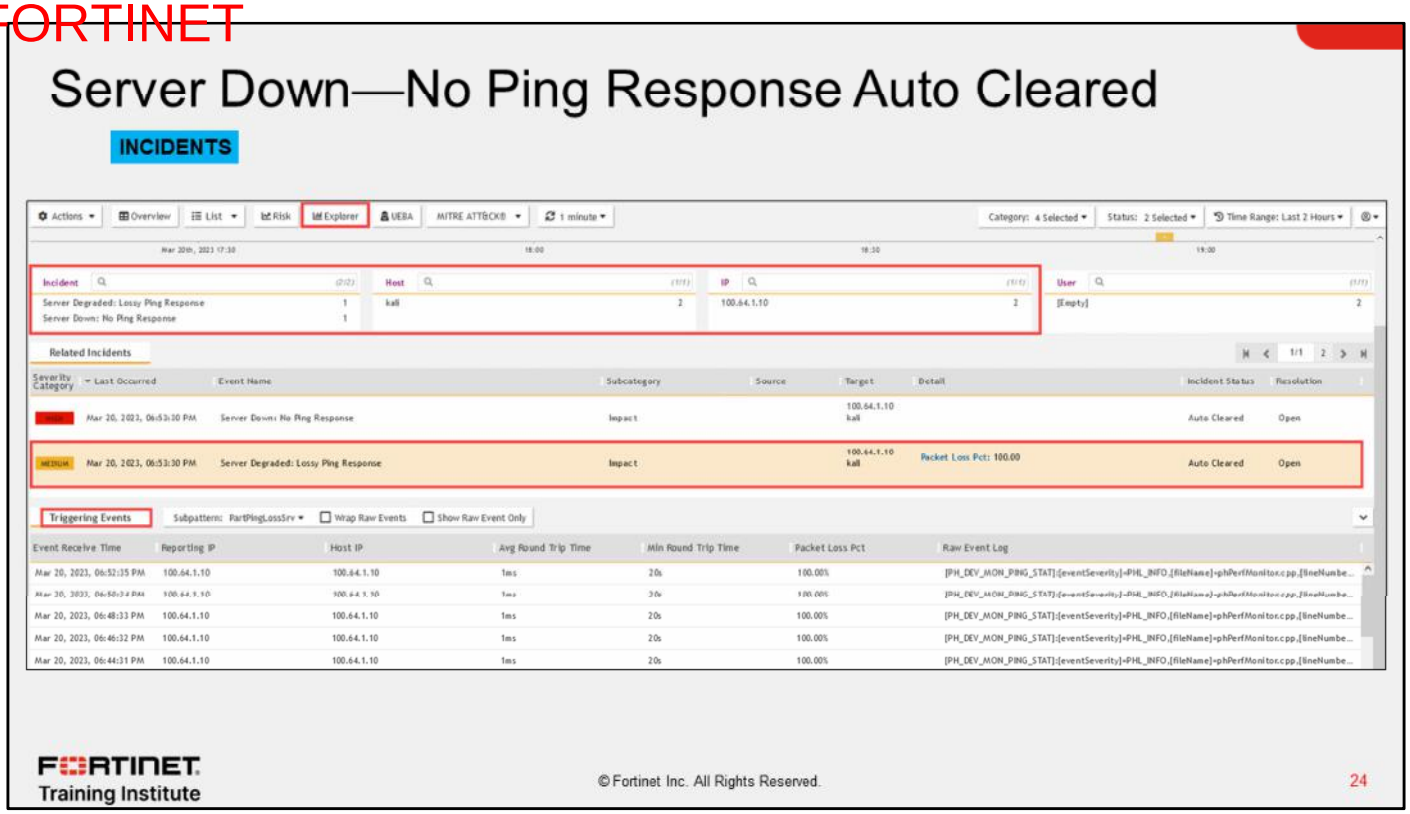
DO NOT REPRINT
© FORTINET



The `phRuleMaster` will wake up again in 2 minutes and evaluate the events for the past 2 minutes. This time, the packet loss percentage is 0% and that will trigger the clear condition subpattern.

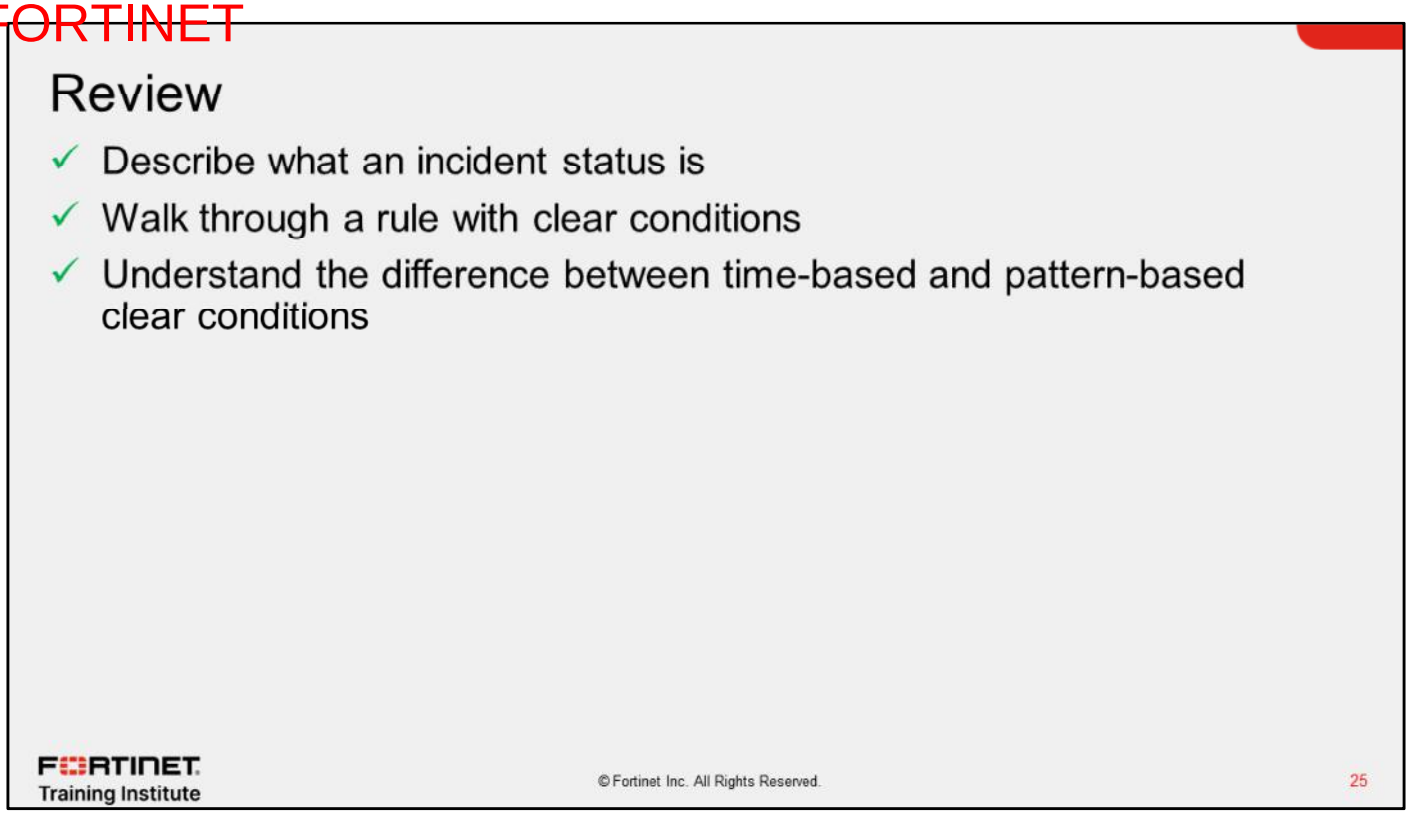
The rule will compare the host IP from the event that triggered the clear condition with the host IP of the original rule that triggered the incident. If the host IP addresses match, that means that the same server that was not responding to the ping is now responding, and the incident status will be set to auto cleared.

DO NOT REPRINT
© FORTINET



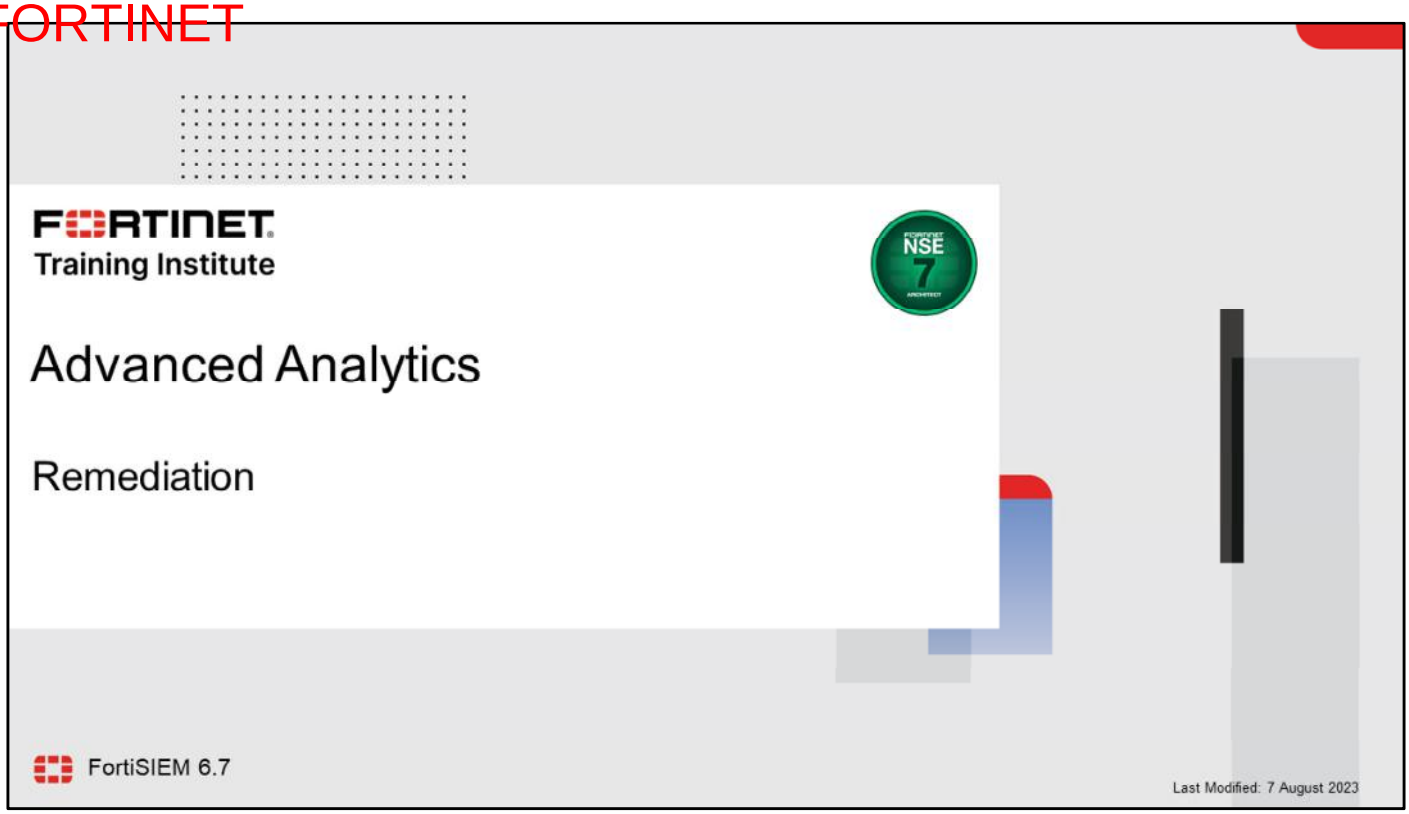
You can verify the status of the incident on the **INCIDENTS** tab. You will see that the incident status for those incidents that have been cleared by the system. Because the incident satisfies clear conditions, the status will be set to auto cleared.

DO NOT REPRINT
© FORTINET



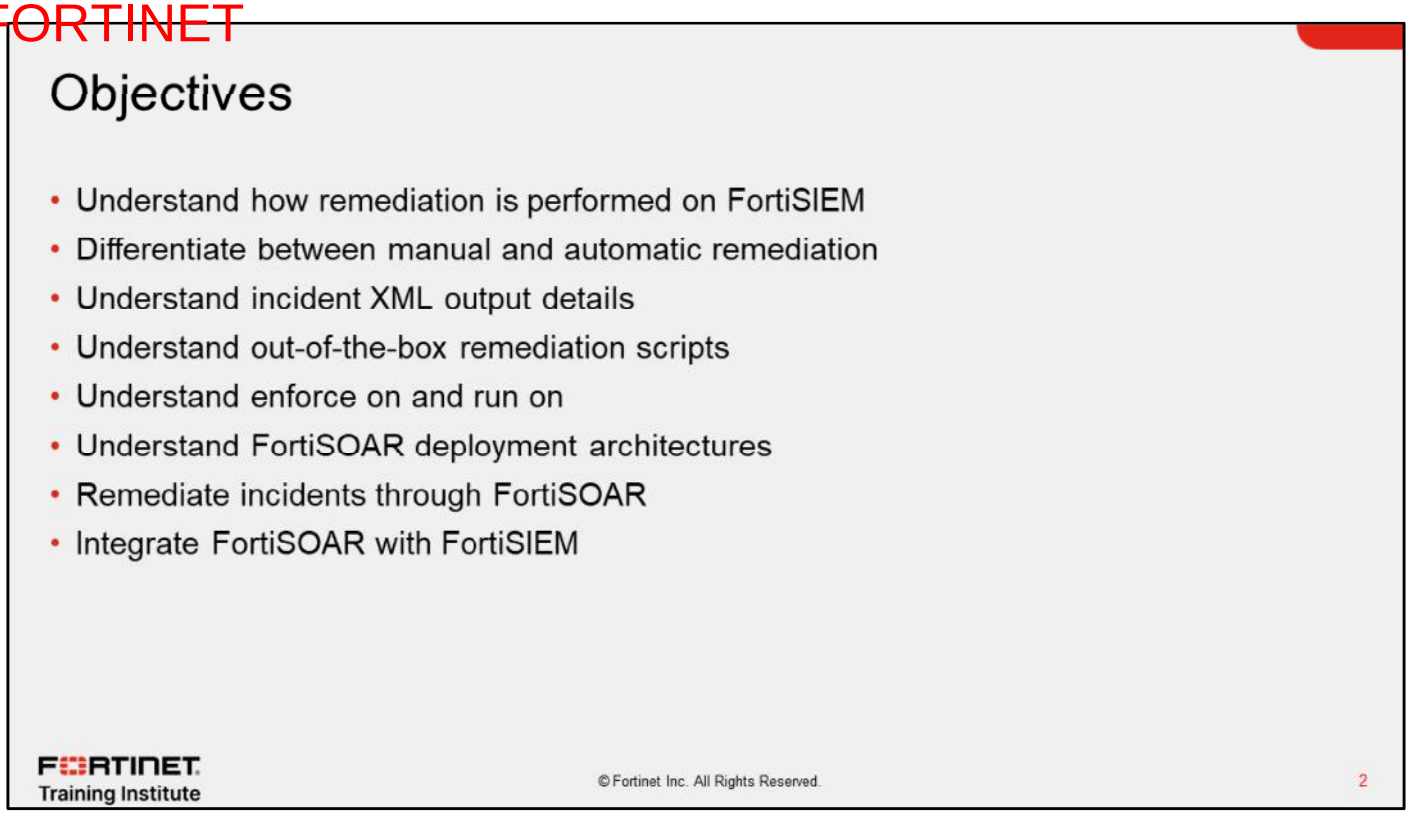
By mastering the objectives covered in this lesson, you learned how to define the conditions that the system requires to clear an incident automatically.

DO NOT REPRINT
© FORTINET



In this lesson, you will learn about remediation on FortiSIEM, which can be done either on an as needed basis or using a notification policy, where the system takes the remediation action when an incident happens. You will also learn to remediate incidents through FortiSOAR, and about various deployment architectures of FortiSOAR.

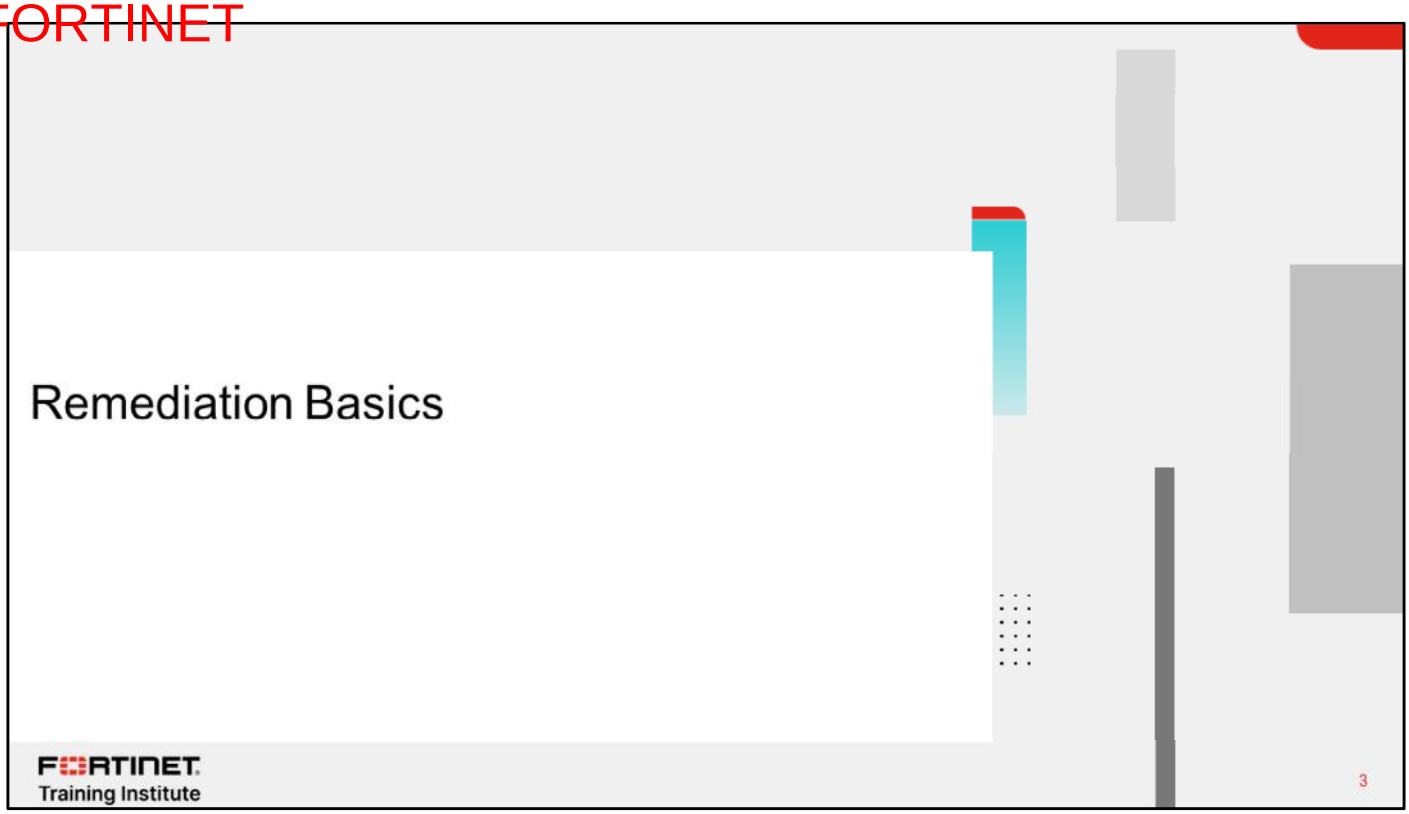
DO NOT REPRINT
© FORTINET



After completing this lesson, you should be able to achieve the objectives shown on this slide.

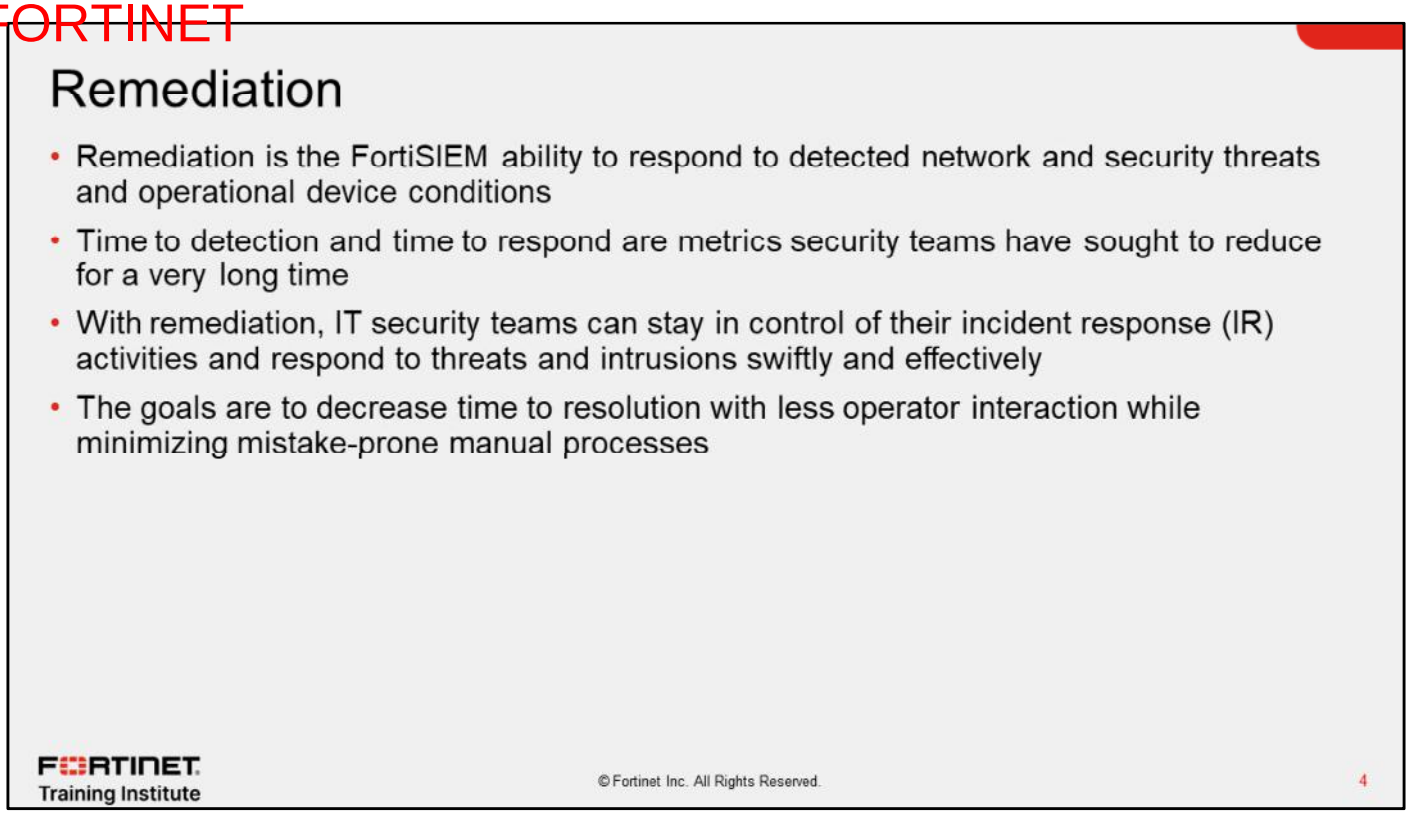
By demonstrating competence in understanding remediation, you will be able to remediate incidents manually or automatically through FortiSIEM or FortiSOAR.

DO NOT REPRINT
© FORTINET



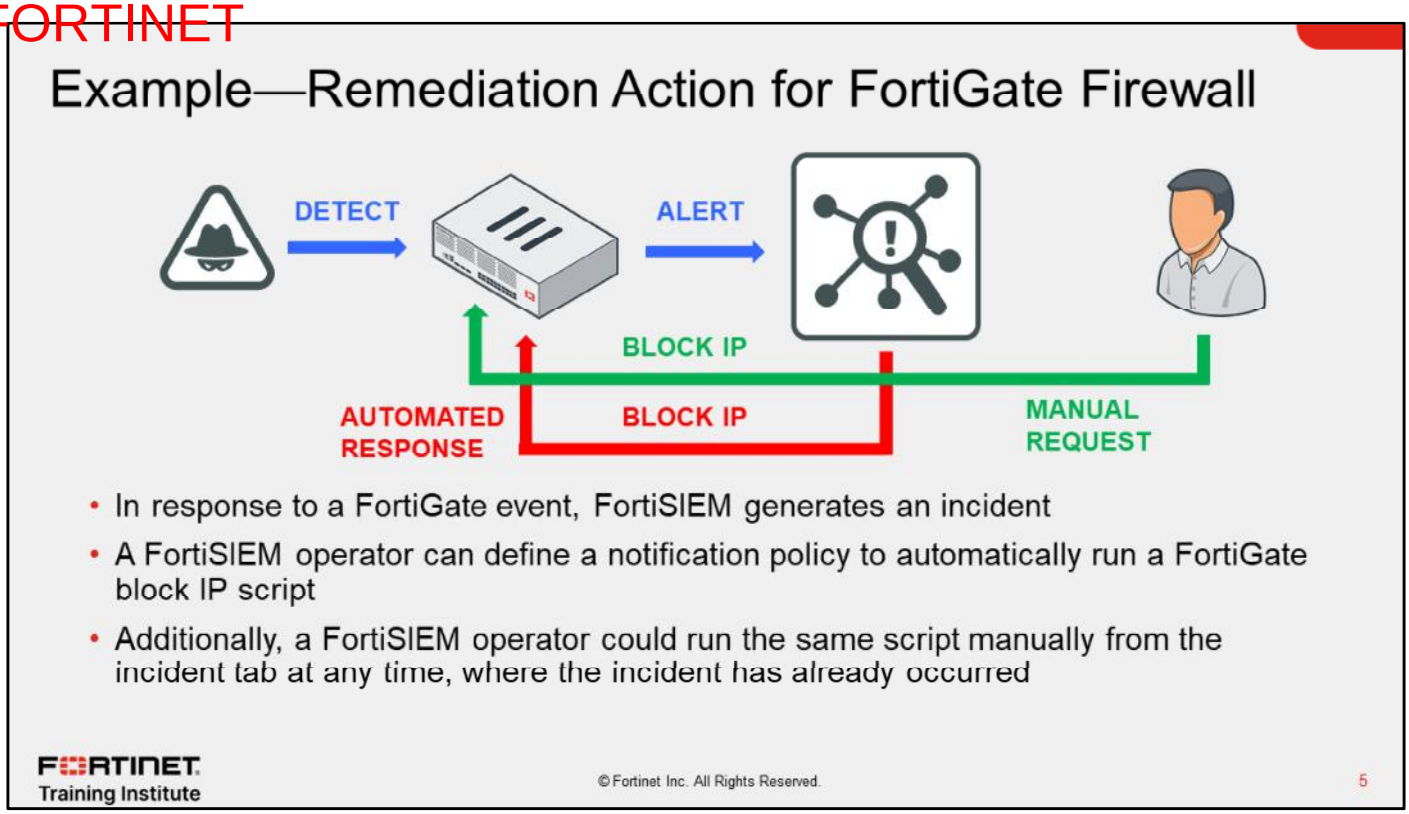
In this section, you will learn the basics of remediation, and how remediation is performed on FortiSIEM.

DO NOT REPRINT
© FORTINET



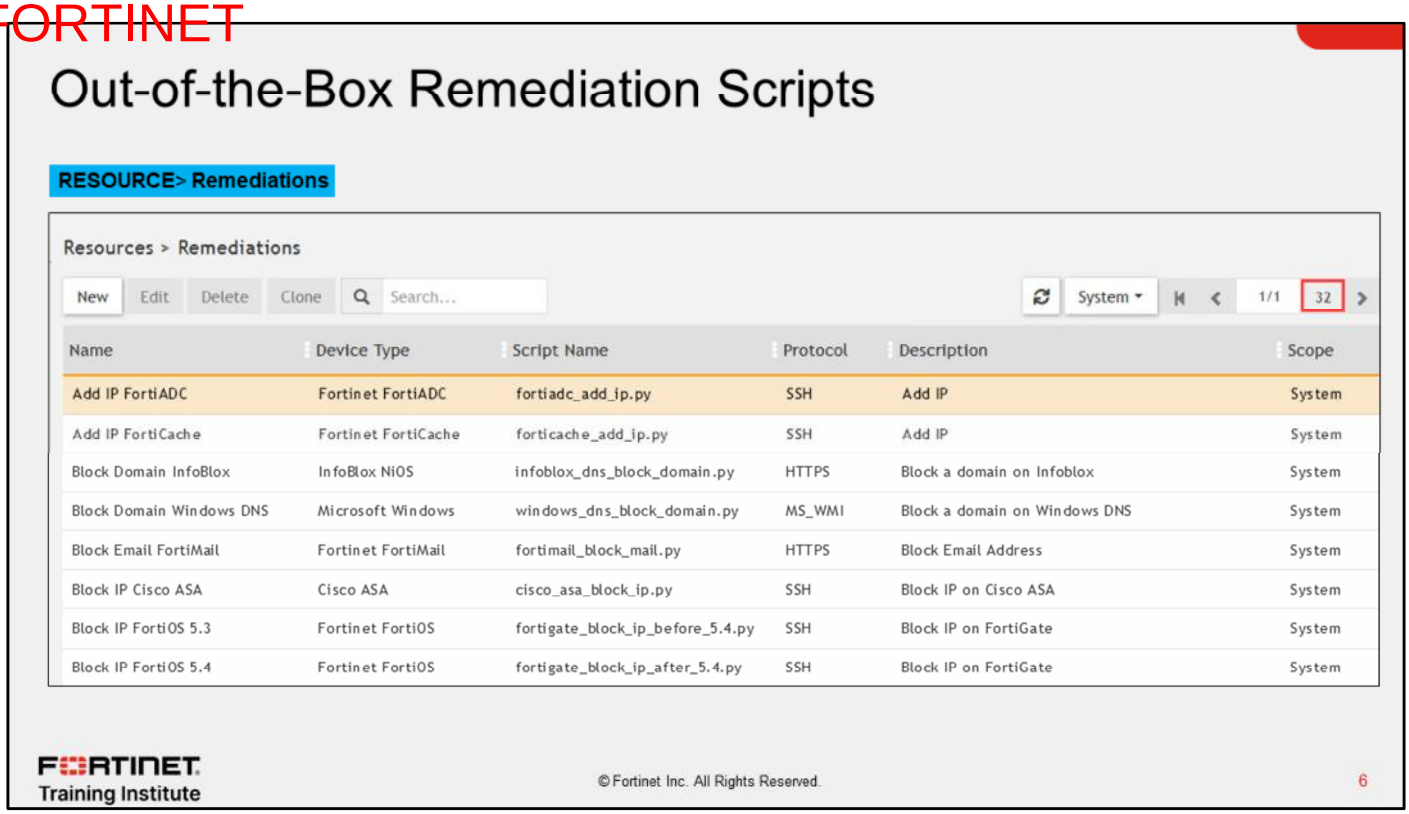
Remediation is the FortiSIEM ability to respond to detected network and security threats and operational device conditions. Time to detection and time to respond are metrics security teams have sought to reduce for a very long time. With remediation, IT security teams can stay in control of their IR activities and respond to threats and intrusions swiftly and effectively. This decreases time to resolution, with less operator interaction while minimizing mistake-prone manual processes.

DO NOT REPRINT
© FORTINET



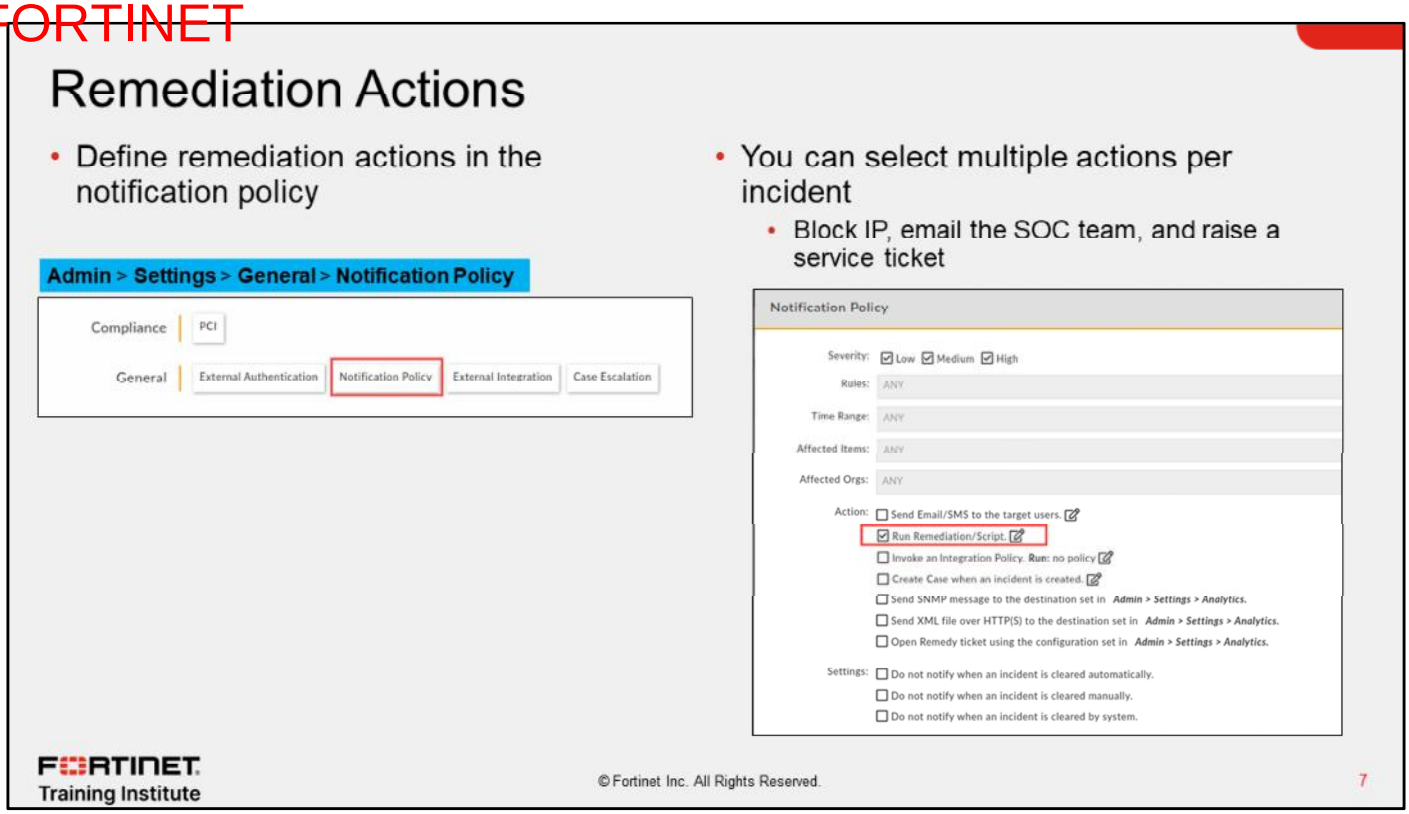
In this conceptual case, if FortiGate generates an event that triggers a rule, which in turn triggers an incident, then you have the option to define the way you want to remediate the incident. You can configure a notification policy, which can be defined to automatically run a FortiGate block IP script. You could also run the same script manually from the incident tab at any time, where the incident has already occurred.

DO NOT REPRINT
© FORTINET



This slide shows some of the out-of-the-box remediation scripts that are available on FortiSIEM. As of release 6.7.4, there are 32 remediation scripts for devices such as FortiGate, Linux, and Windows.

DO NOT REPRINT
© FORTINET

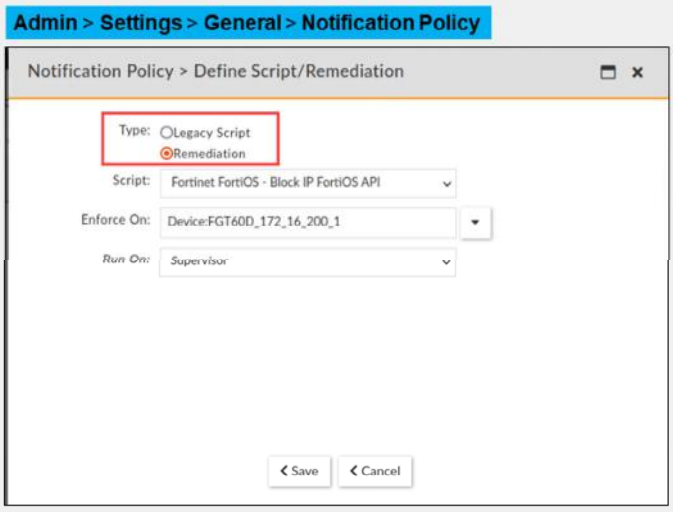


You define the The automatic remediation actions in the notification policy. You can select multiple actions for every incident, such as blocking the IP, emailing the SOC team, and also raising a service ticket. Select **Run Remediation/Script** in the notification policy configuration.

DO NOT REPRINT
© FORTINET

Script Types

- Two types of script options
  - Legacy script is the scripting option supported in previous FortiSIEM and AccelOps releases
  - Remediation is using the new built-in scripts in FortiSIEM 5.0 and beyond



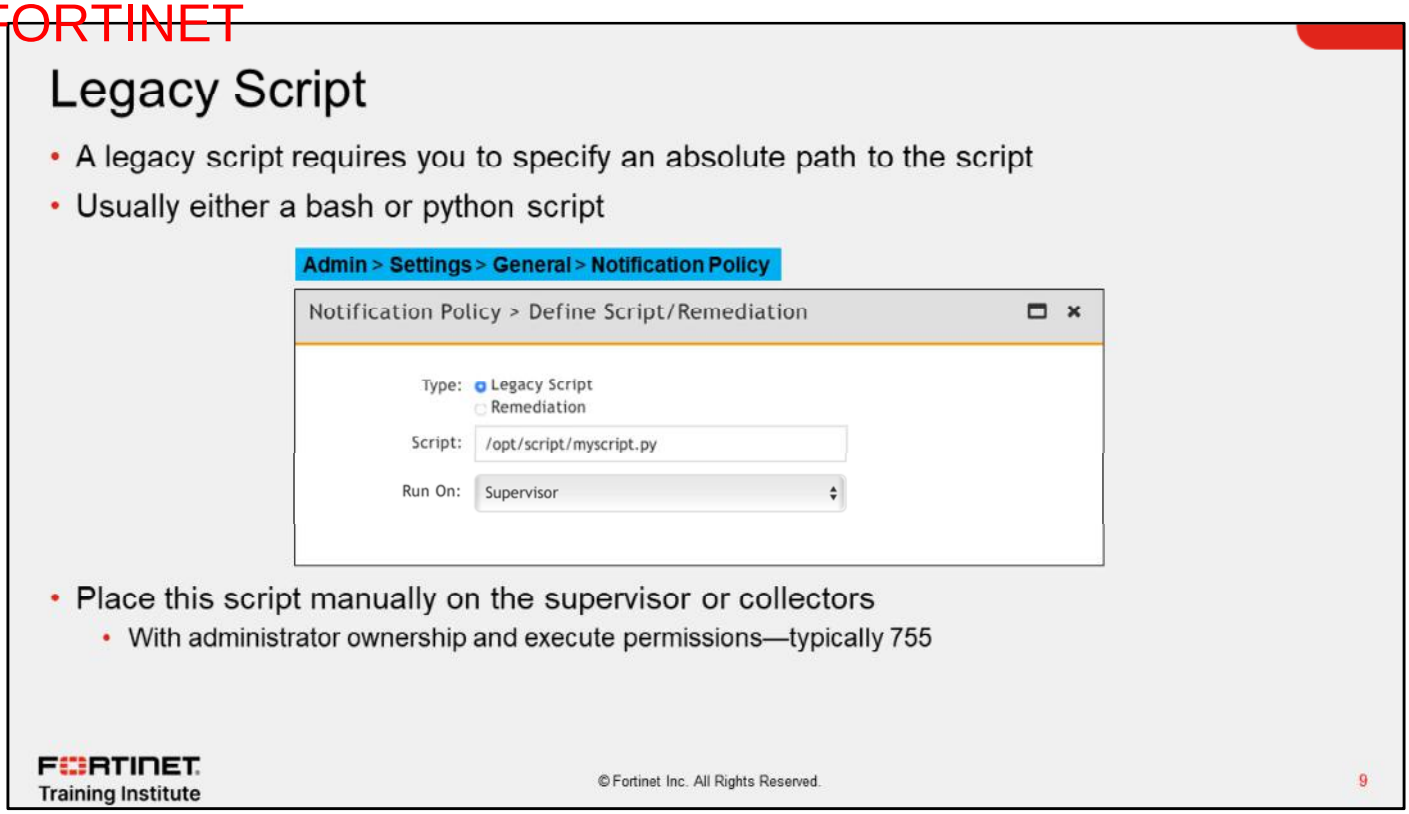
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

8

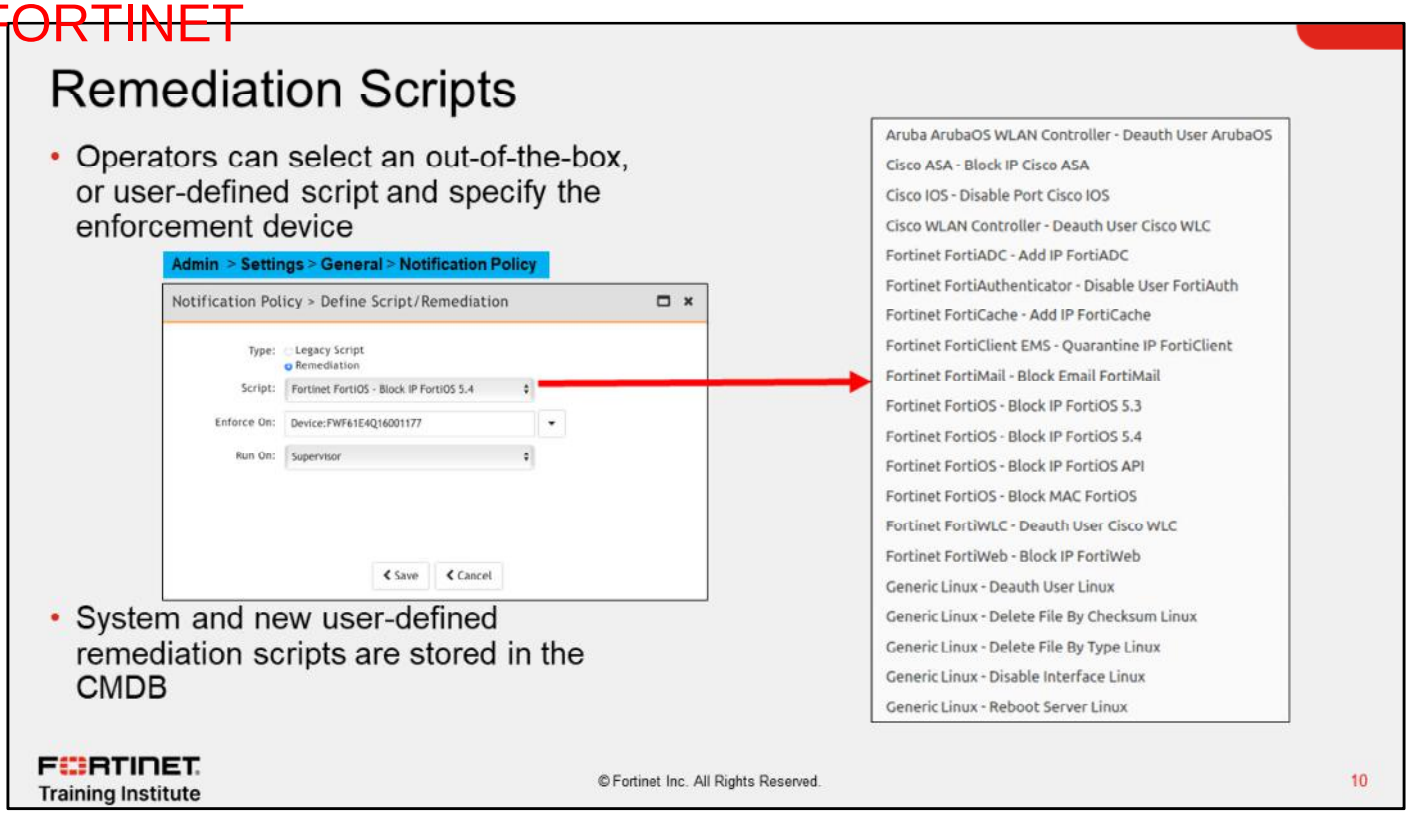
There are two types of script options. Legacy script is the scripting option supported in previous FortiSIEM and AccelOps releases. Remediation uses the new, built-in scripts included FortiSIEM 5.0 and later.

DO NOT REPRINT
© FORTINET



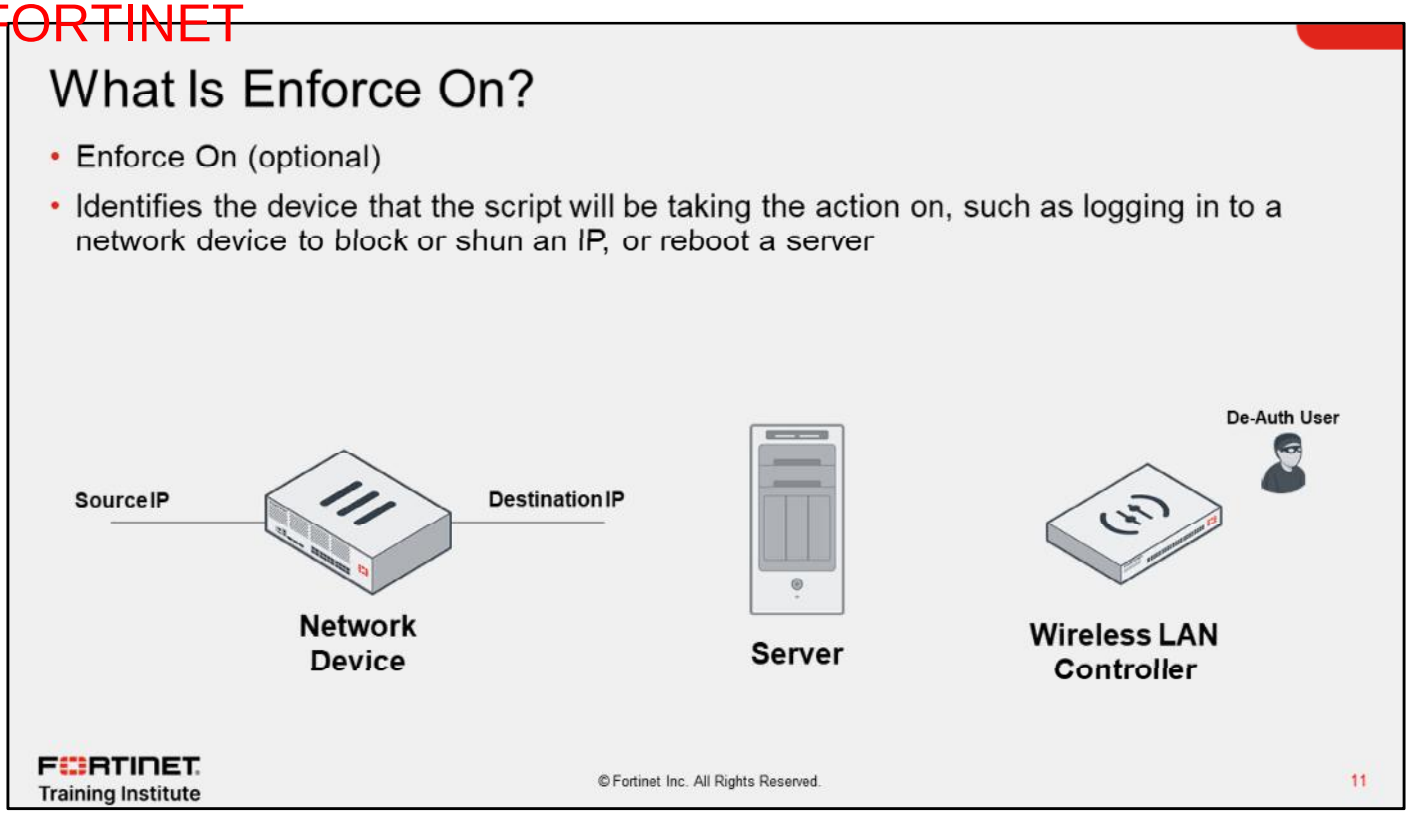
A legacy script requires you to specify an absolute path to the script. The script is usually either a bash or python script. A FortiSIEM administrator who has administrative ownership and execute permissions would manually place this script on the supervisor or collectors.

DO NOT REPRINT
© FORTINET



If you choose to use a remediation script, then you can select an out-of-the-box or user-defined script and specify the enforcement device. Starting from version 5.0 and later, FortiSIEM stores system and new user-defined remediation scripts are stored in the **CMDB**.

DO NOT REPRINT
© FORTINET

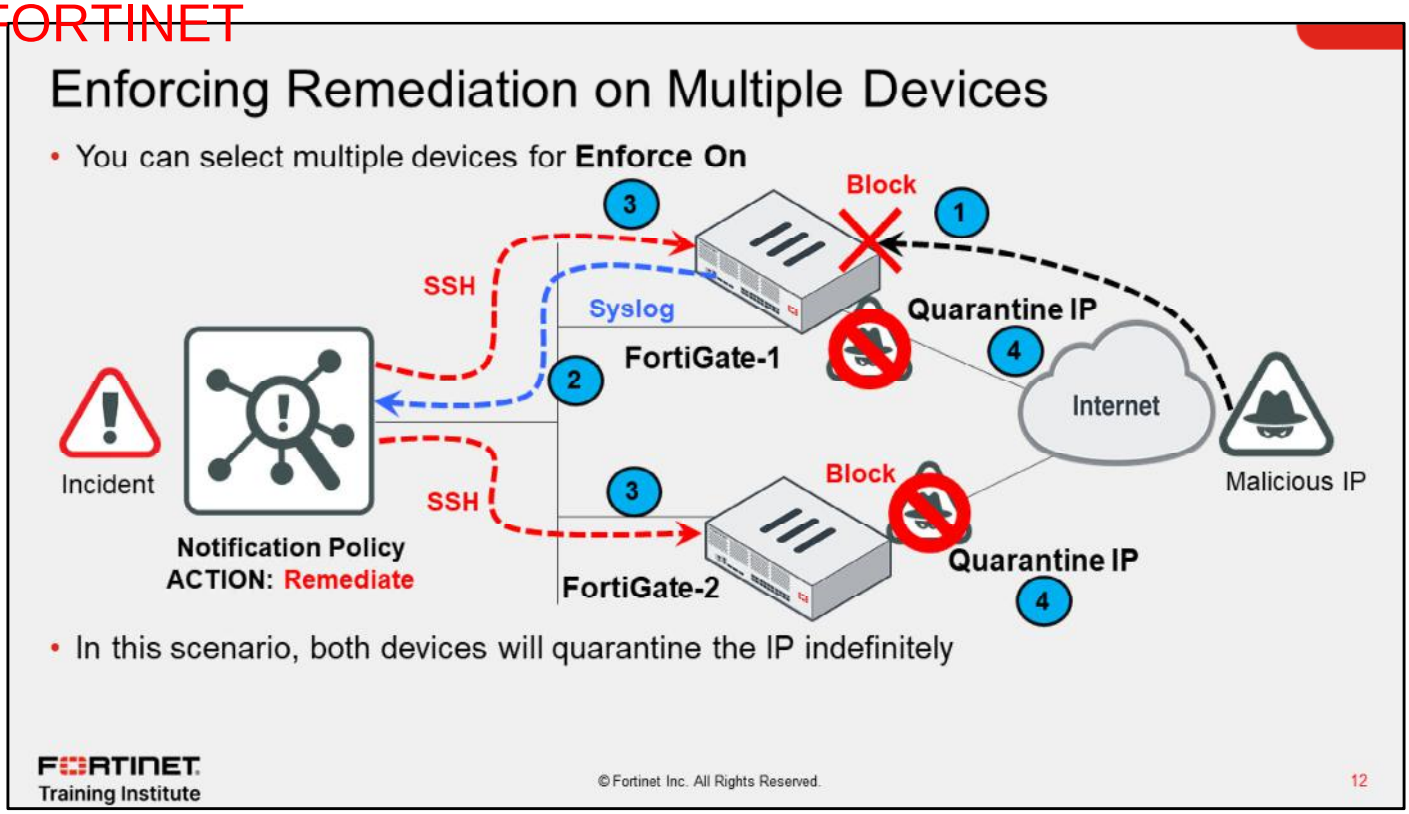


The **Enforce On** setting identifies the device that the script will be taking the action on, such as logging in to a network device to block or shun an IP, or rebooting a server.

If the **Enforce On** value is defined then the remediation script uses it as a variable on the FortiSIEM back end. Otherwise, if left blank, the out-of-the-box scripts try to identify the **Enforce On** device through the incident XML.

Enforce On is optional for both automatic remediation in the notification policy or manual remediation on the **Incident** tab. However, it is recommended that you always specify a value.

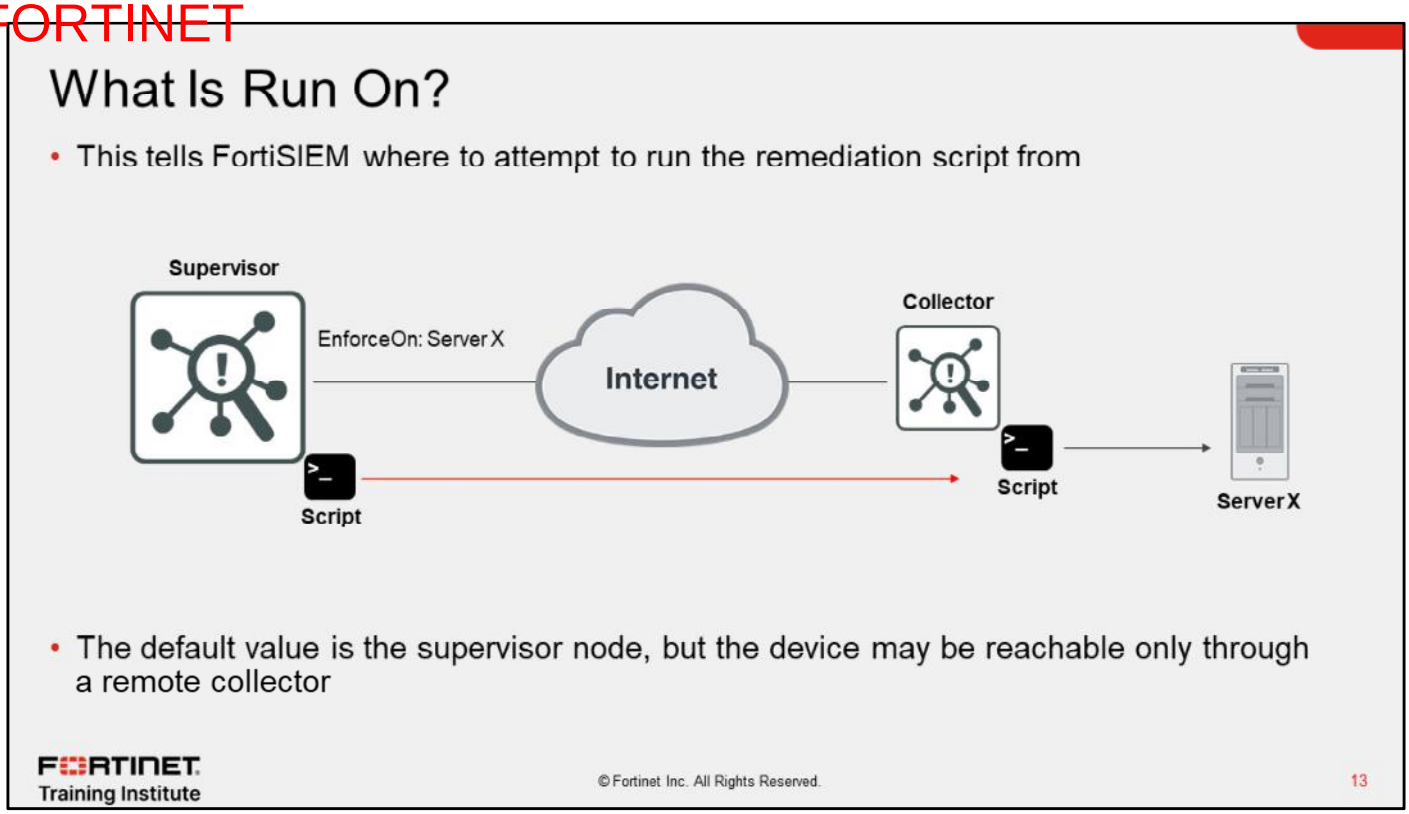
DO NOT REPRINT
© FORTINET



You can select one or more devices from the CMDB on which the script needs to be enforced. In that case, when an incident is generated, the script will run on multiple devices. If an incident, such as a zero-day attack, is detected on one device, then you can run a script across multiple devices to block the attacker before the attack occurs on other devices.

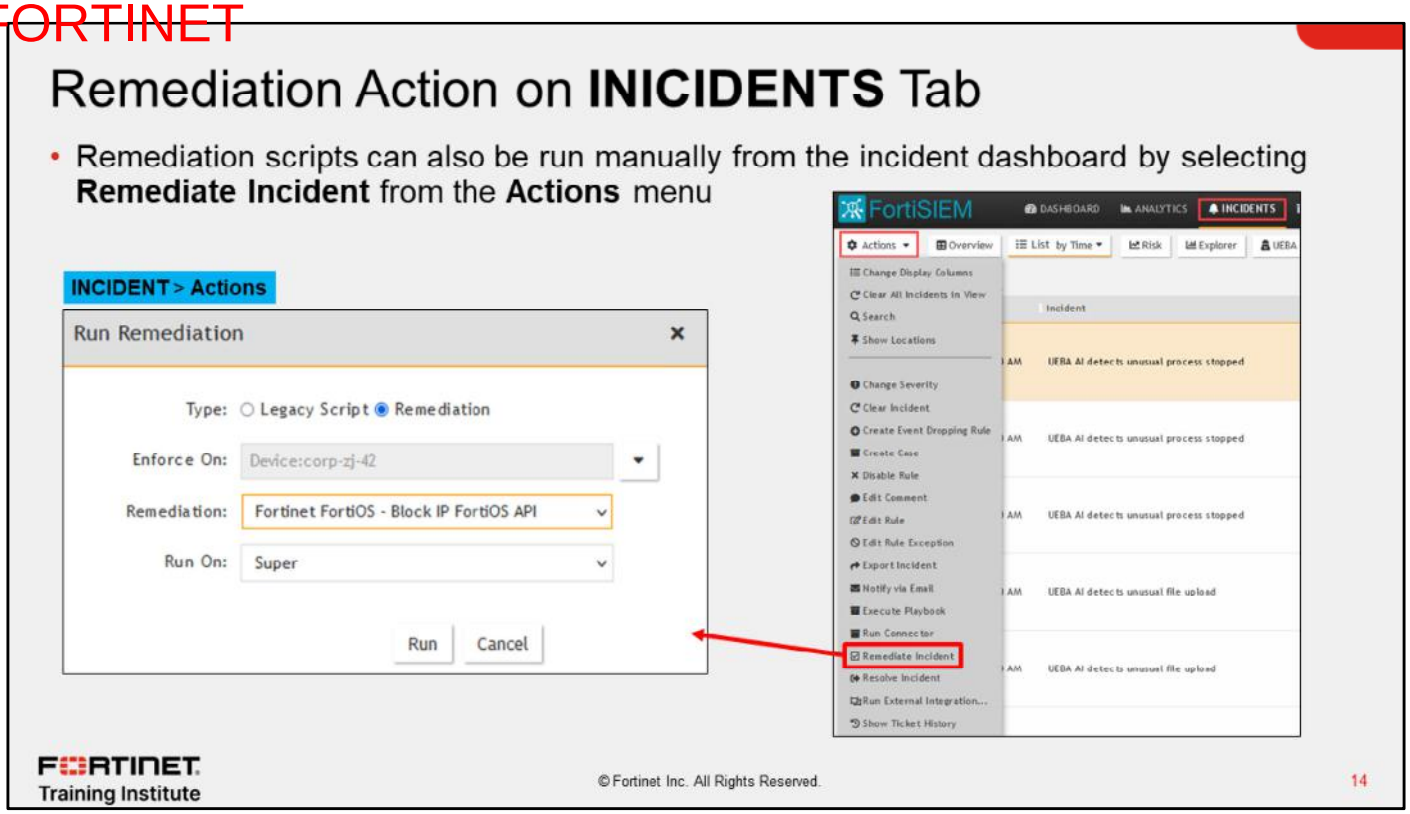
In the example shown on this slide, a malicious IP is detected by the first FortiGate, which reports that event to FortiSIEM. The remediation script then connects over SSH to both FortiGate devices, and then executes the remediation script which, in this case, quarantines the malicious IP address.

DO NOT REPRINT
© FORTINET



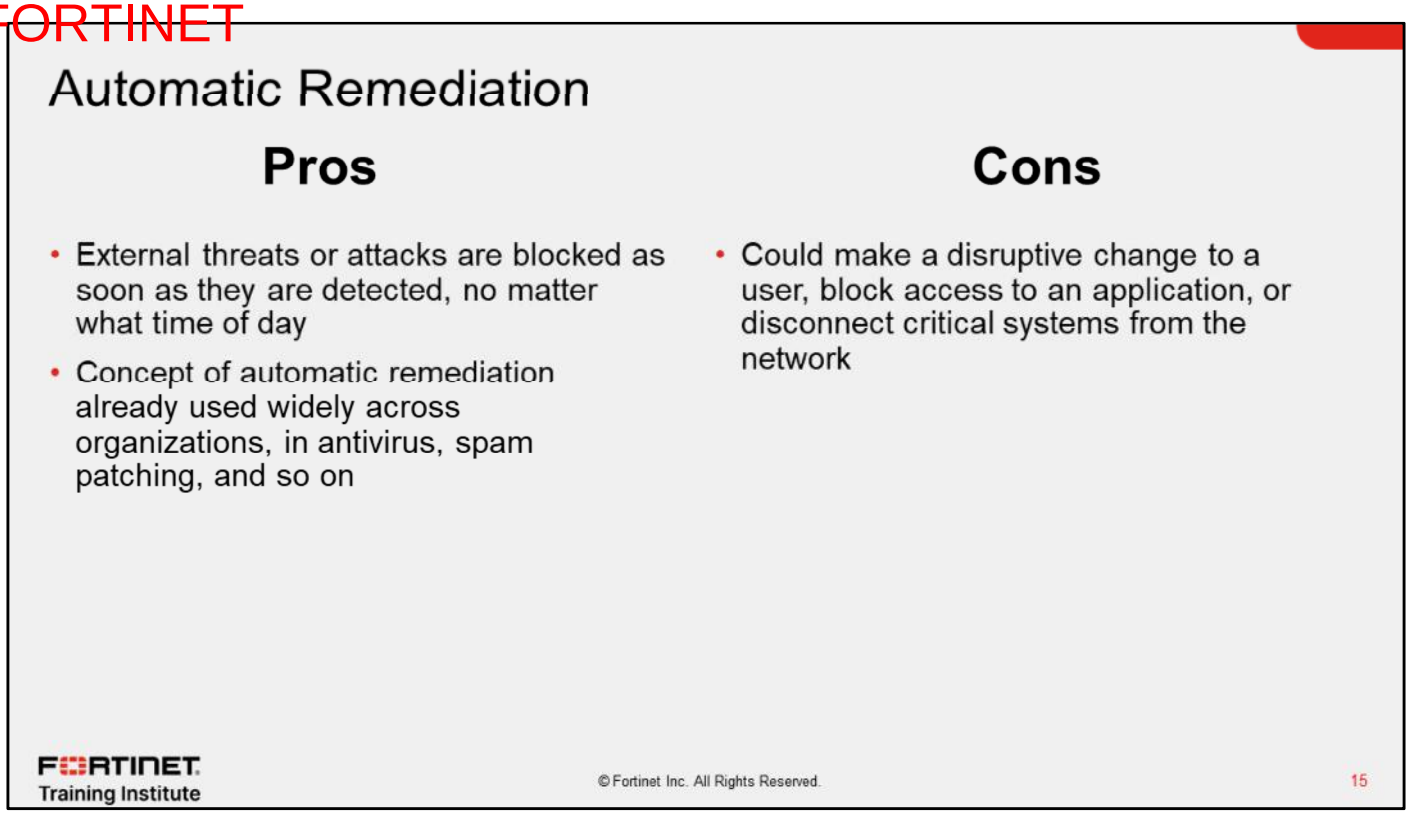
The **Run On** setting tells FortiSIEM where to attempt to run the remediation script from. The default value is the supervisor node, but the device may be reachable through a remote collector. If you have defined collectors for any organization, then by default the run on resource is the collector for that organization.

DO NOT REPRINT
© FORTINET



You can run remediation scripts manually from the incident dashboard by selecting **Remediate Incident** from the **Actions** menu.

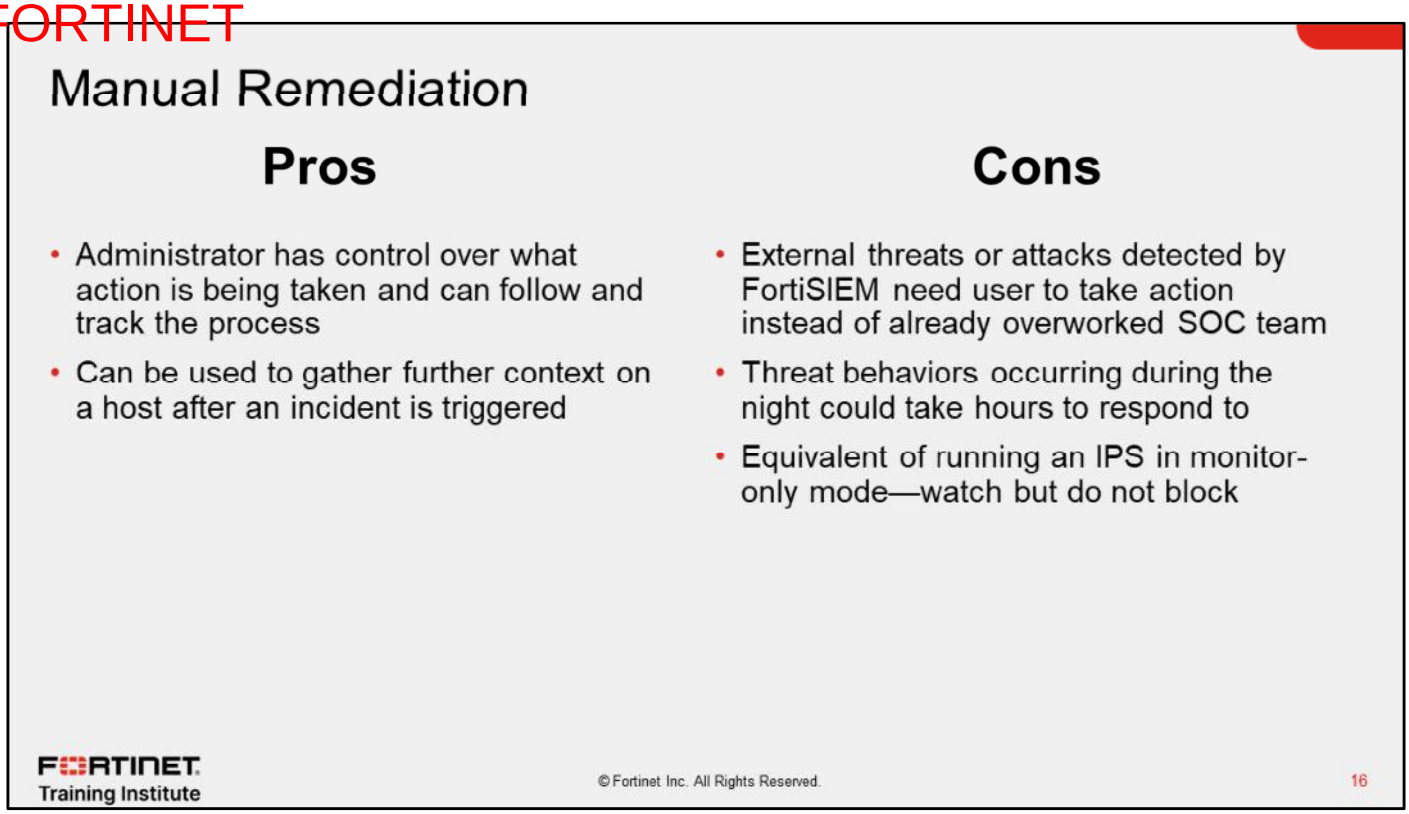
DO NOT REPRINT
© FORTINET



The benefit of using automatic remediation is that external threats or attacks are blocked as soon as they are detected, no matter what time of day. The concept of automatic remediation is already used widely across organizations in antivirus, spam patching, and so on. So, it is just another layer of security on top of your existing automated security framework.

The disadvantage of automatic remediation is that the script can potentially make disruptive changes to a user, blocking access to an application or disconnecting critical systems from the network, causing more harm than good.

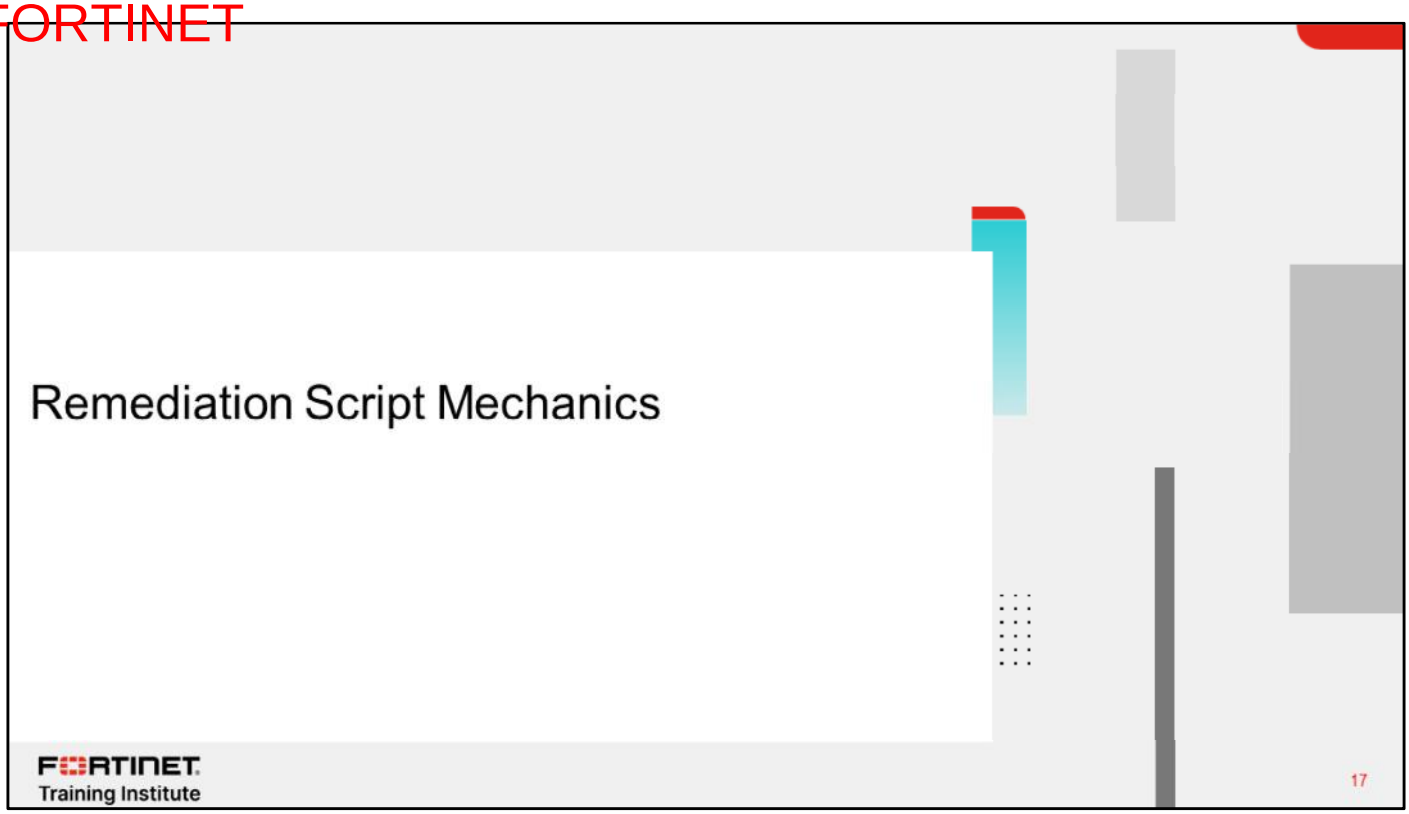
DO NOT REPRINT
© FORTINET



The advantages of manual remediation are that the administrator has control over what action is being taken and can follow and track the process of the remediation action. You can gather further useful context on a host after an incident is triggered if you are performing manual remediation.

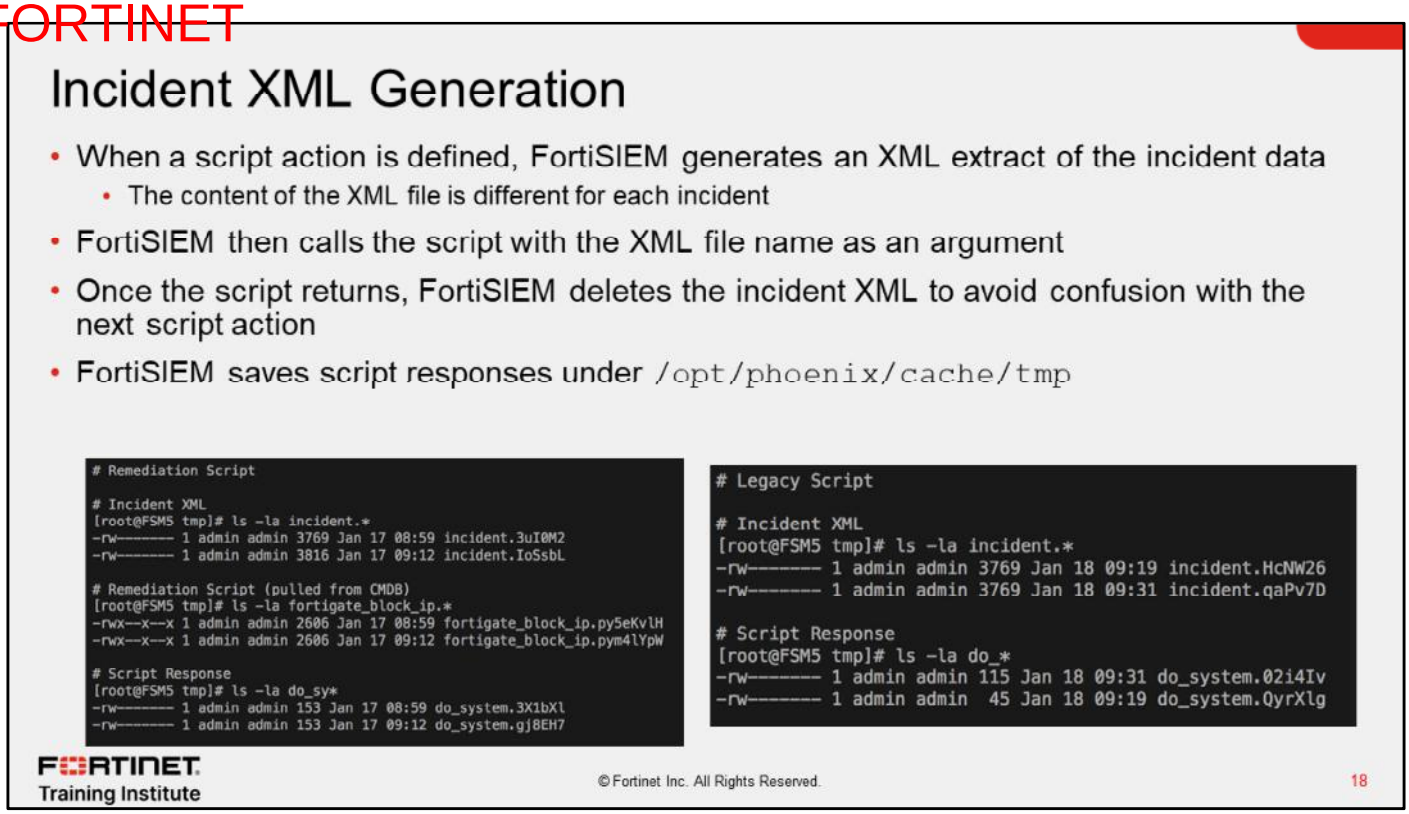
The disadvantage of manual remediation is that external threats or attacks that are detected by FortiSIEM need the user to apply the action, which can be time consuming for an overworked SOC team. Threats that occur at night can take hours to respond to, and during that time more systems could be compromised. Running manual remediation is equivalent to running an IPS in monitor-only mode—watch but do not block.

DO NOT REPRINT
© FORTINET



In this section, you will learn how remediation scripts work on FortiSIEM.

DO NOT REPRINT
© FORTINET



When you define a script action on FortiSIEM, the device will generates an XML extract of the incident data. The contents of the XML file are different for each incident. FortiSIEM then calls the script with the XML file name as an argument. After the script returns, FortiSIEM deletes the incident XML file is deleted to avoid any confusion with the next script action.

By default, FortiSIEM saves the incident output files, remediation scripts, and script responses it generates in a temporary folder. The incident files have the name `incident` followed by some random characters. If you are running a remediation script, then FortiSIEM copies the script to the device that is executing the script with the script name followed by random characters. The response from the script is saved with the filename `do_system` followed by random characters.

If you are running only a legacy script, then FortiSIEM runs the script from the location where the script is stored. For legacy scripts, you see only the incident and script response file. Because FortiSIEM deletes these files almost immediately deleted, it can be hard to catch the output.

DO NOT REPRINT
© FORTINET

Incident XML Example 1

```
<?xml version="1.0" encoding="UTF-8" ?>
<incident incidentId="723" ruleType="PH_RULE_VIRUS_BY_FIREWALL_NON_REMEDY" severity="9"
  repeatCount="1" organization="Aviation" status="0">
  <name>Malware found by firewall but not remediated</name>
  <remediation></remediation>
  <description>Detects that firewall content inspection devices found a virus but could not remediate it</description>
  <policyID></policyID>
  <displayTime>Thu Feb 06 13:56:00 EST 2020</displayTime>
  <incidentCategory>Security/Persistence</incidentCategory>
  <incidentSource>
  <entry attribute="srcIpAddr" name="Source IP">10.0.3.10
 (Win_Agent)</entry>
  </incidentSource>
  <incidentTarget>
  </incidentTarget>
  <incidentDetails>
  <entry attribute="virusName" name="Malware Name">EICAR_TEST_FILE</entry>

  </incidentDetails>
  <affectedBizSrvc>null</affectedBizSrvc>
  <identityLocation>
</identityLocation>  </incident>
```

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

19

This is an example of an incident XML file where a firewall found malware, but did not remediate it. You can remediate this incident using a remediation script that extracts relevant information from the incident, such as source IP, file name, and so on. You can use the information that the script gathers from the incident XML file to perform actions, such as configuring or editing a firewall policy, to block any future occurrence of the malware or virus.

DO NOT REPRINT
© FORTINET

Incident XML Output Details

- XML tag and attribute definitions

```
<?xml version="1.0" encoding="UTF-8" ?>
<incident incidentId="723" ruleType="PH_RULE_VIRUS_BY_FIREWALL_NON_REMEDY" severity="9"
   repeatCount="1" organization="Aviation" status="0">
```

XML Tag	XML Attribute	Description
<incident>		
	incidentId	Unique ID of the incident
	ruleType	ID of the rule, which is the rule event type created
	severity	Severity of the incident between 1–10 (low=1–4, medium=5–8, high=9–10)
	repeatCount	How many times this incident has occurred
	organization	In service provider deployments, the organization affected by the incident, which is always Super in enterprise deployments
	status	Incident status, one of New, Update, Cleared Manually, Clear, System Clear

FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

20

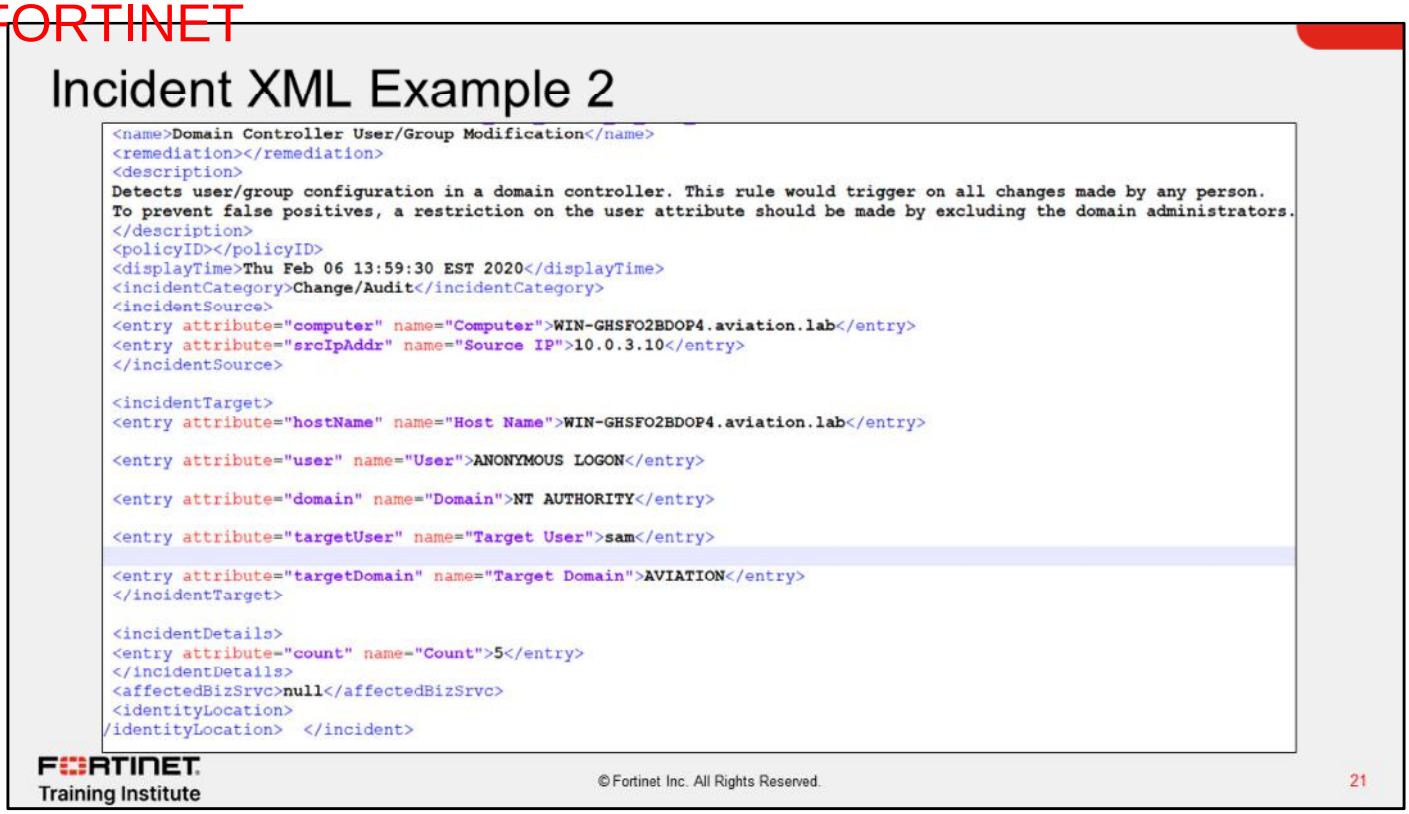
Every incident has a unique ID, which is tagged with XML attributes. There are XML tags and within those tags you can define XML attributes.

The `ruleType` attribute defines the ID of the rule, which is the rule event type that was created. The `severity` attribute defines the value between 1–10 where 1–4 is low severity, 5–8 is medium severity, and 9–10 is high severity. The `repeatCount` attribute defines the number of times the incident has occurred.

The `organization` attribute defines which organization was impacted in service provider deployments. In enterprise deployments this organization is always `Super`.

The `status` attribute defines the incident status where 0 means the incident is active, 1 means the incident was automatically cleared, 2 means the incident was manually cleared, and 3 means the incident was cleared by the system.

DO NOT REPRINT
© FORTINET



This slide shows another example of an incident XML file. This incident was generated because the rule detected that a user made some changes to domain controller user and computer settings. This incident file has `user` and `targetUser` as incident attributes. The user is the person who made the changes and the target user is the one whose account was changed or modified.

DO NOT REPRINT
© FORTINET

Incident XML Output Details

XML Tag	XML Attribute	Description
<name>		Name of the rule that triggered the incident
<remediation>		Any rule-defined remediation instructions
<policyID>		Incident notification policy ID
<displayTime>		Time at which the incident occurred (Last Occur Time)
<incidentSource>	*many possible	Source of the Incident
<incidentTarget>	*many possible	The target of the incident
<incidentDetails>	*many possible	Details that triggered the incident
<affectedBizSrvc>		Lists the business services that the incident source or target is a member of
<identityLocation>		Information associated with user identity and location
<clearedReason>		Present only for manually cleared incidents
<rawEvents>		Contains maximum of 10, raw event logs for the incident

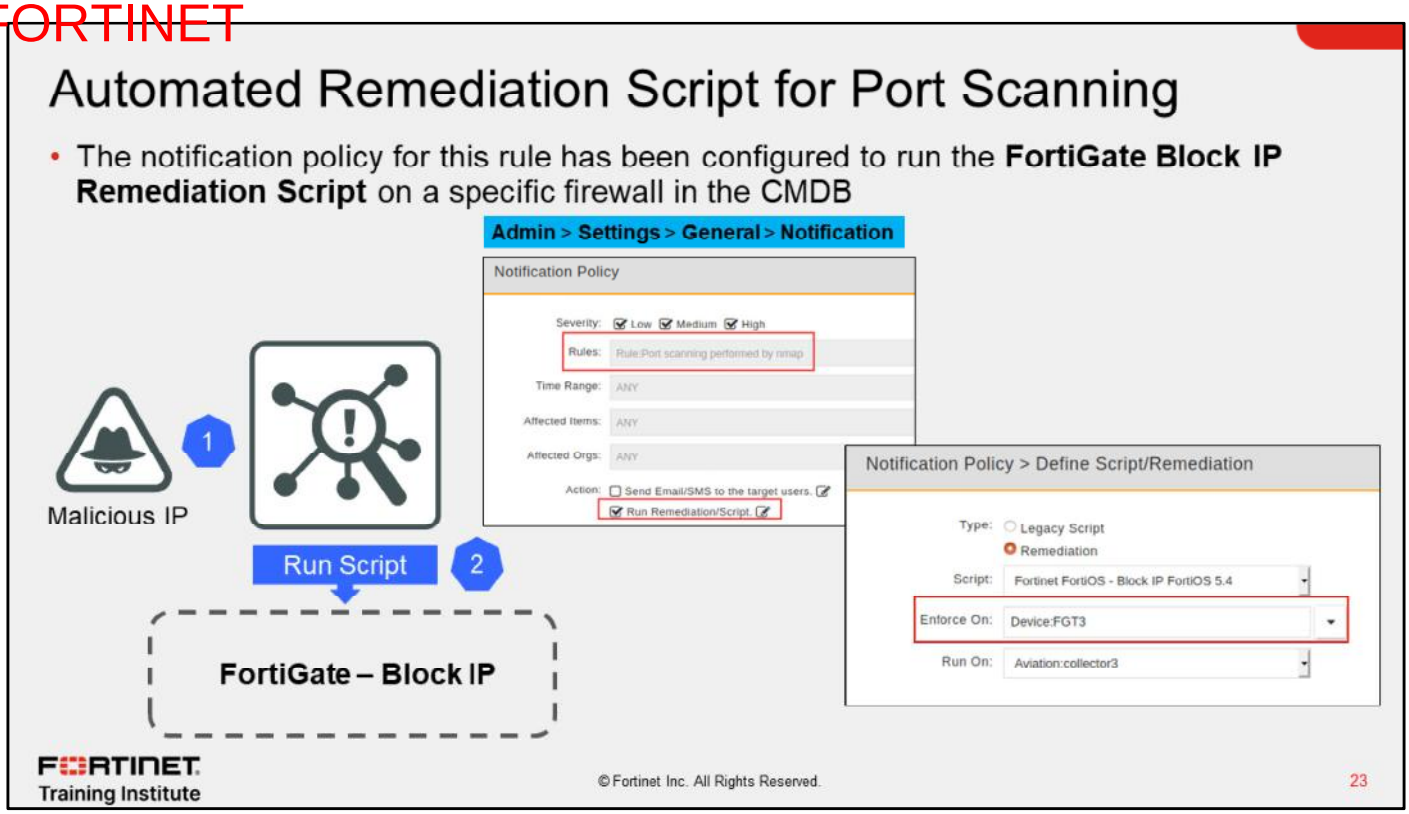
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

22

This slide shows some XML tags and their definitions.

DO NOT REPRINT
© FORTINET

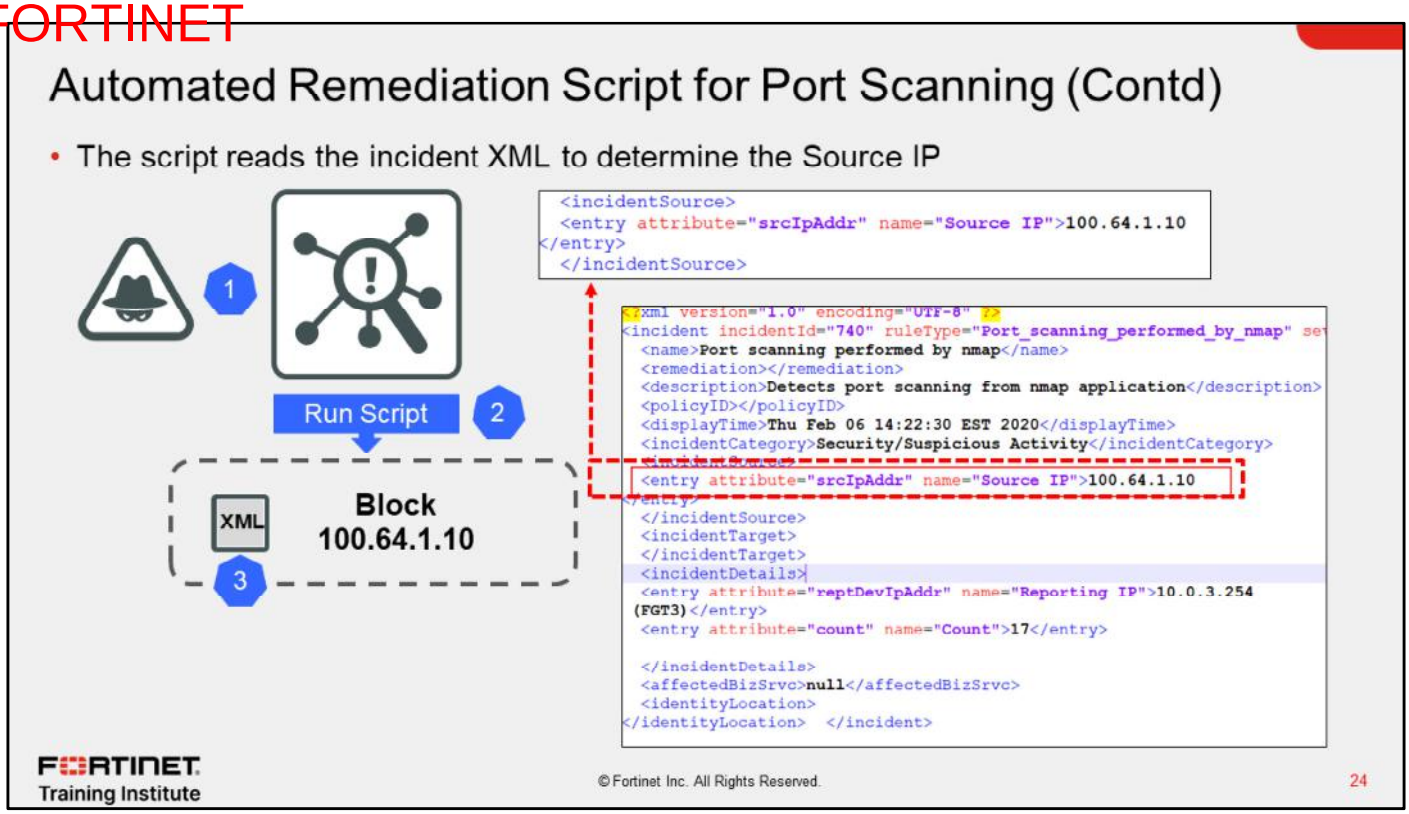


This slide shows an example of an automated remediation script for port scanning.

In the notification policy, you select the rules for which the remediation will be performed. You must define the script before you can enable **Run Remediation Script**. While defining the script, you can select the **Enforce On** and **Run On** options. **Enforce On** is optional, but it is recommended that you configure this setting. If you don't configure **Enforce On**, then the script picks the **Enforce On** setting from the incident XML file.

In the example shown on this slide, the script is selected to block a malicious IP on FortiGate.

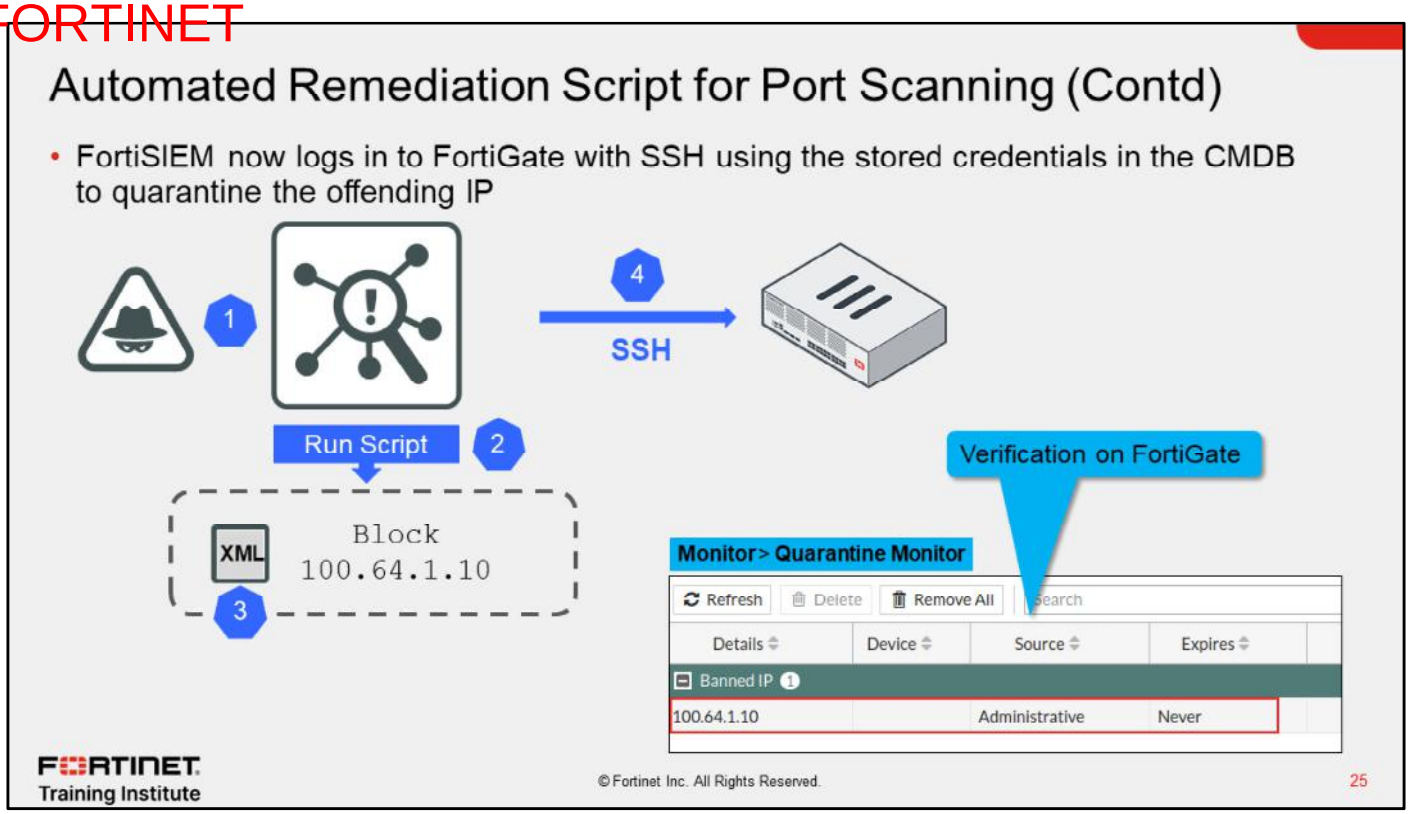
DO NOT REPRINT
© FORTINET



If an incident is generated, then the script reads the source IP from the XML file. The source IP is specified with the `incidentSource` tags.

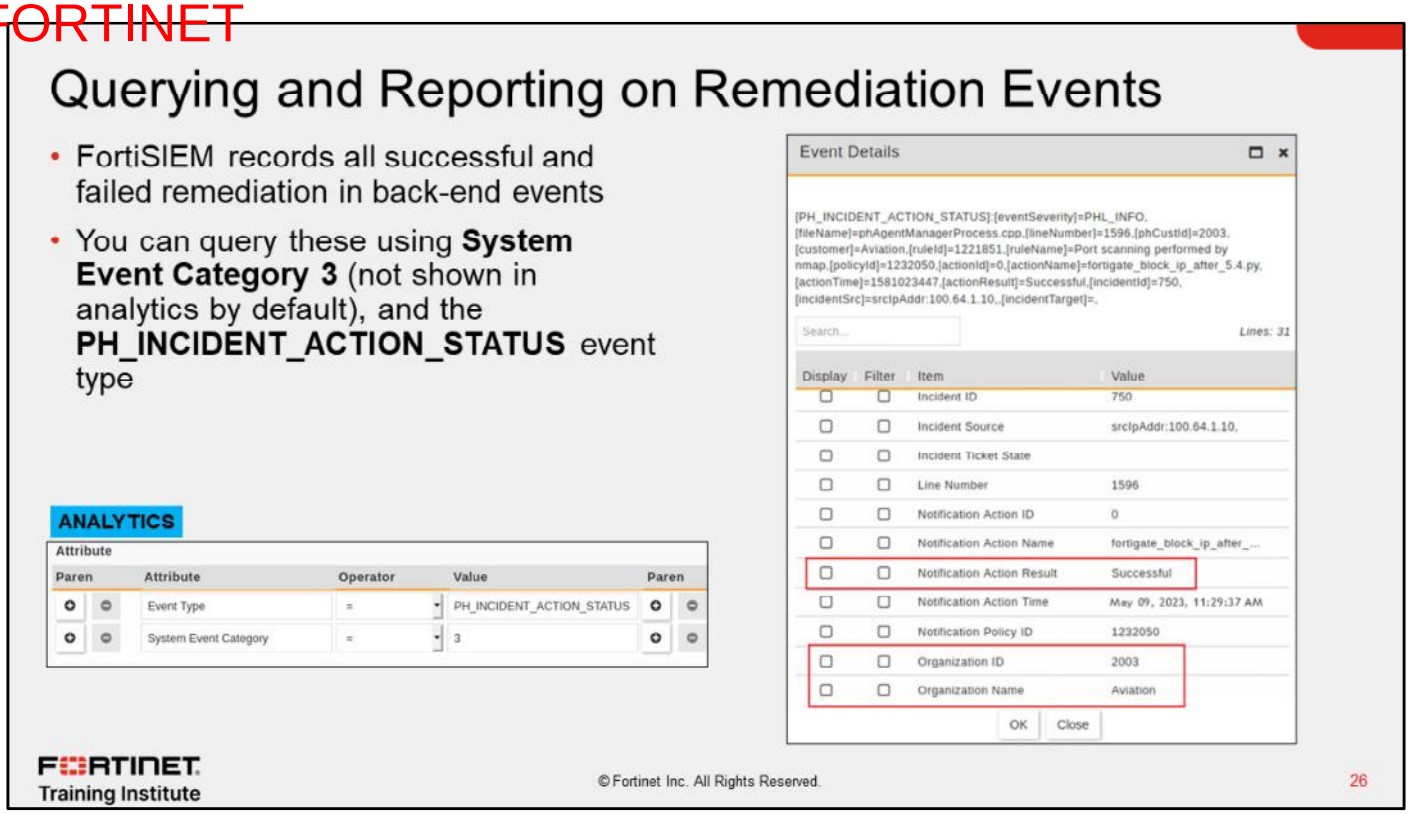
There can be different types of XML attributes depending on the type of source. In the example shown on this slide, the source IP address will be extracted from the XML file and action will be taken on that IP address, which is considered a malicious IP.

DO NOT REPRINT
© FORTINET



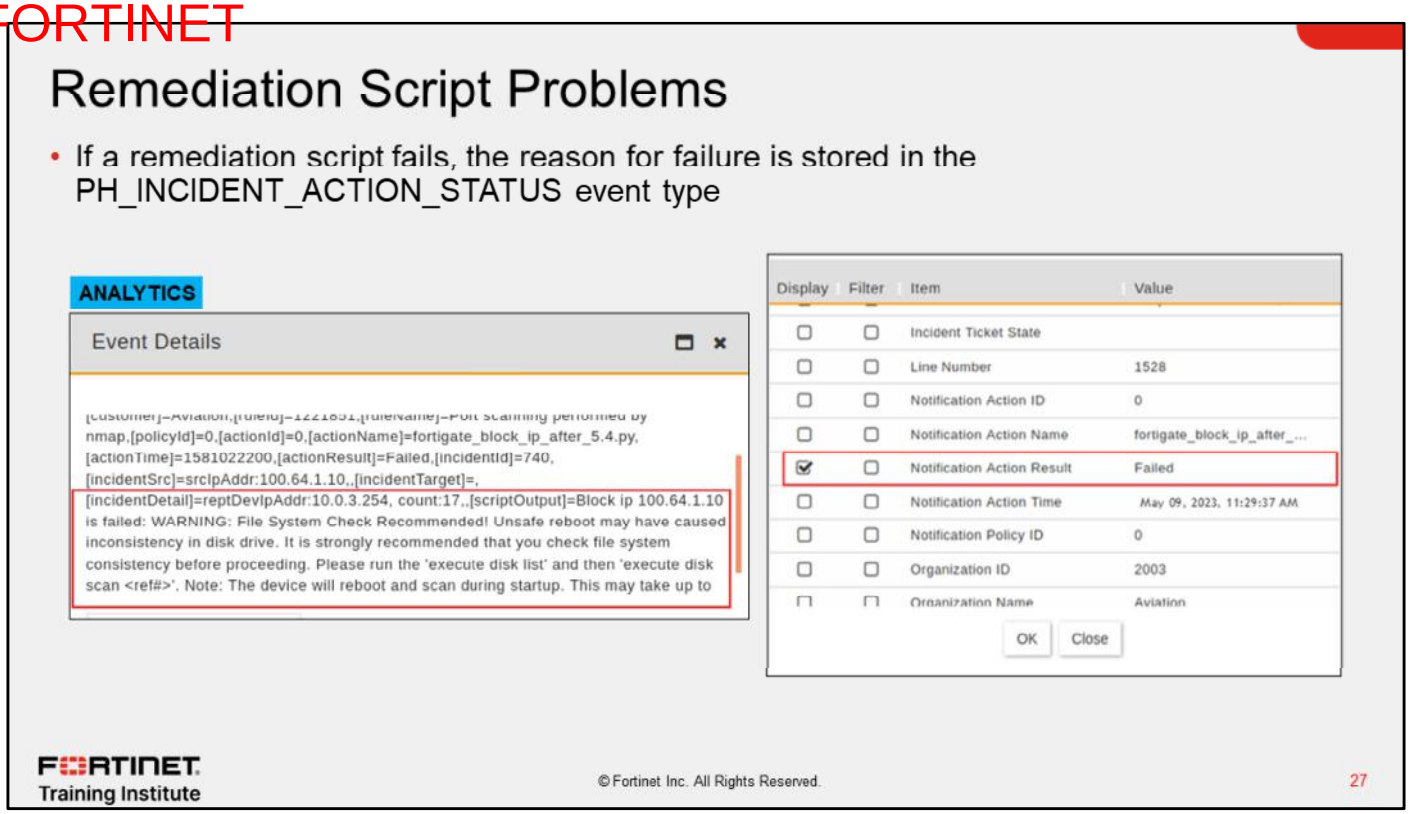
FortiSIEM now logs in to FortiGate with SSH using the stored credentials on the CMDB to execute the commands necessary to quarantine the offending IP address. You can view the banned IP on FortiGate by navigating to the **Quarantine Monitor** page. You can also view the banned IP in the CLI by running the command `diagnose user quarantine list`.

DO NOT REPRINT
© FORTINET



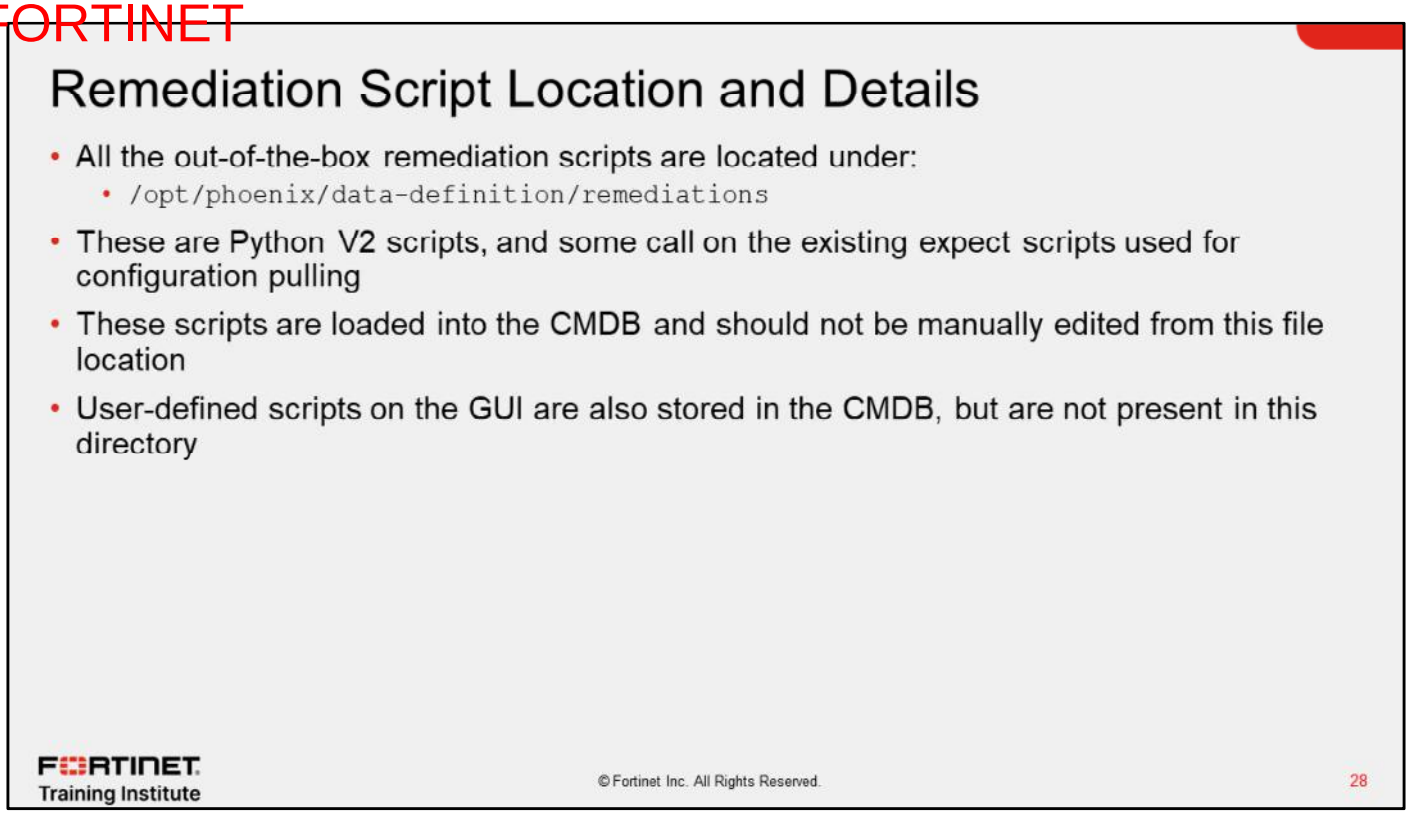
FortiSIEM records all successful and failed remediation in back-end events. You can query these using **System Event Category 3**, which is not shown in analytics by default and the `PH_INCIDENT_ACTION_STATUS` event type.

DO NOT REPRINT
© FORTINET



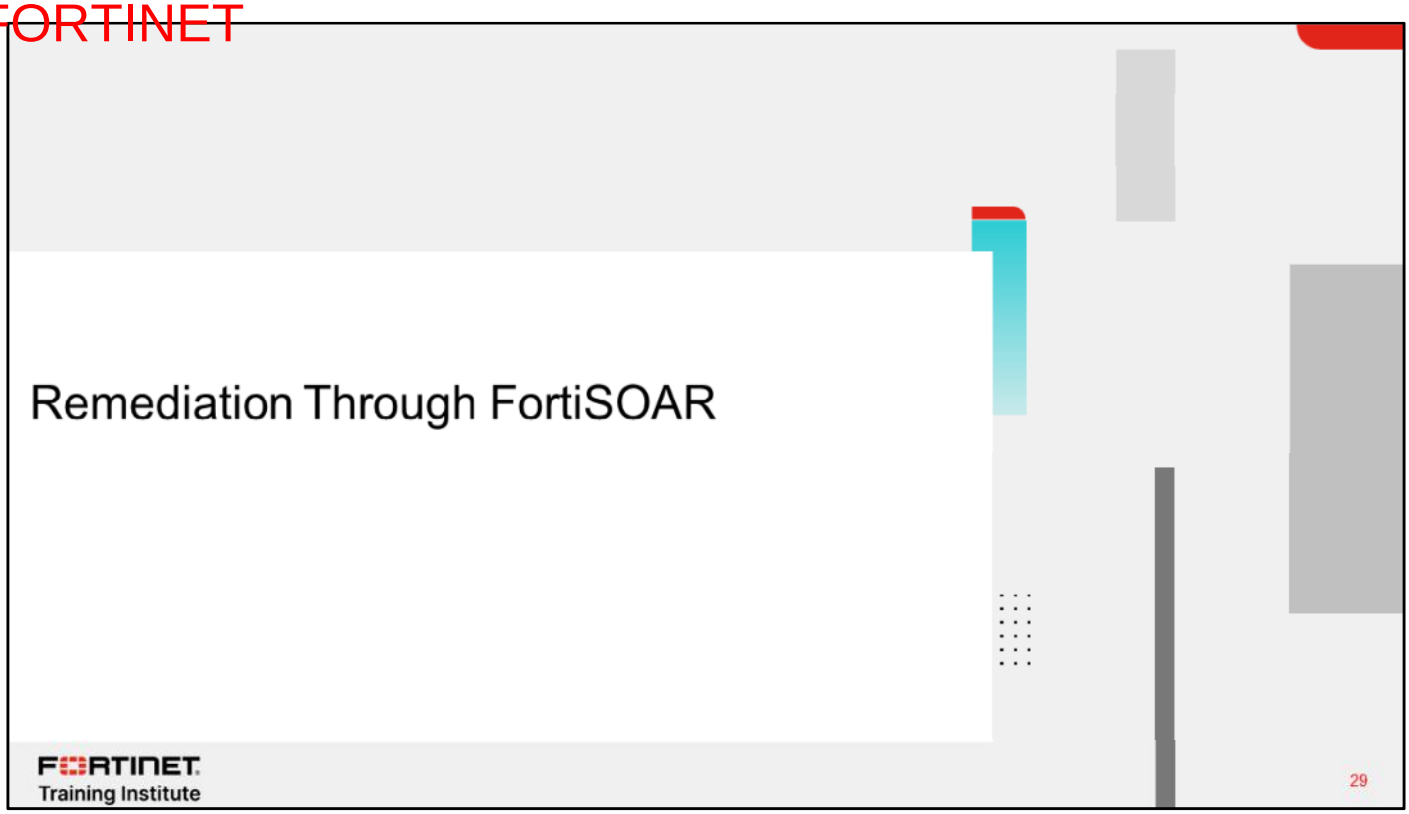
In the example shown on this slide, the remediation script fails for some reason, and FortiSIEM records the reason for the failure. As shown, the script failed to execute because the file system on FortiGate was corrupt. Sometimes, FortiSIEM also provides a solution to the problem, such as suggesting that you run commands on FortiGate to resolve the issue.

DO NOT REPRINT
© FORTINET



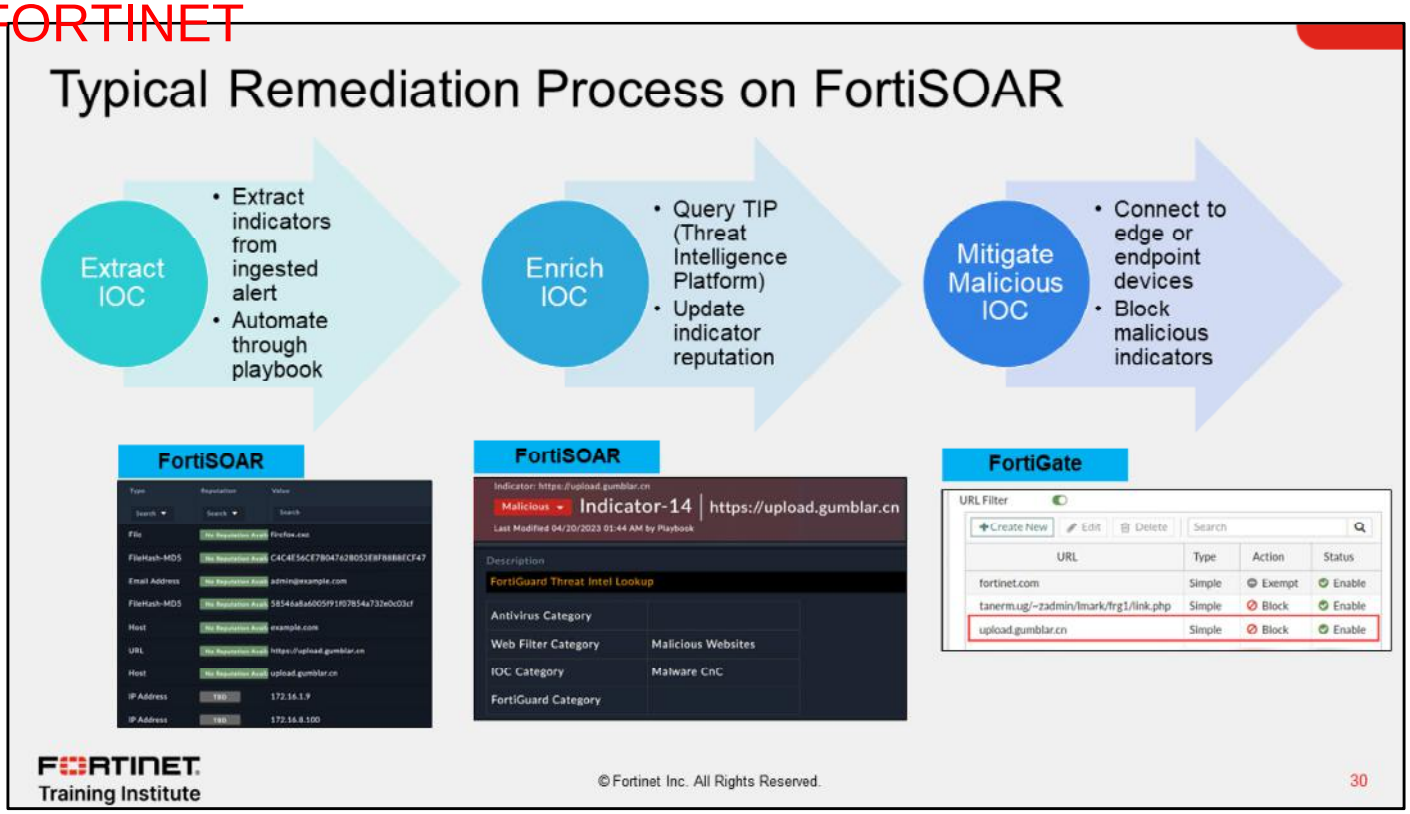
All out-of-the-box remediation scripts are located in a remediations directory. These are Python V2 scripts, and some call on the existing expect scripts used for configuration pulling. These scripts are loaded into the CMDB and should not be manually edited from this file location. User-defined scripts on the GUI are also stored in the CMDB, but are not present in this directory.

DO NOT REPRINT
© FORTINET



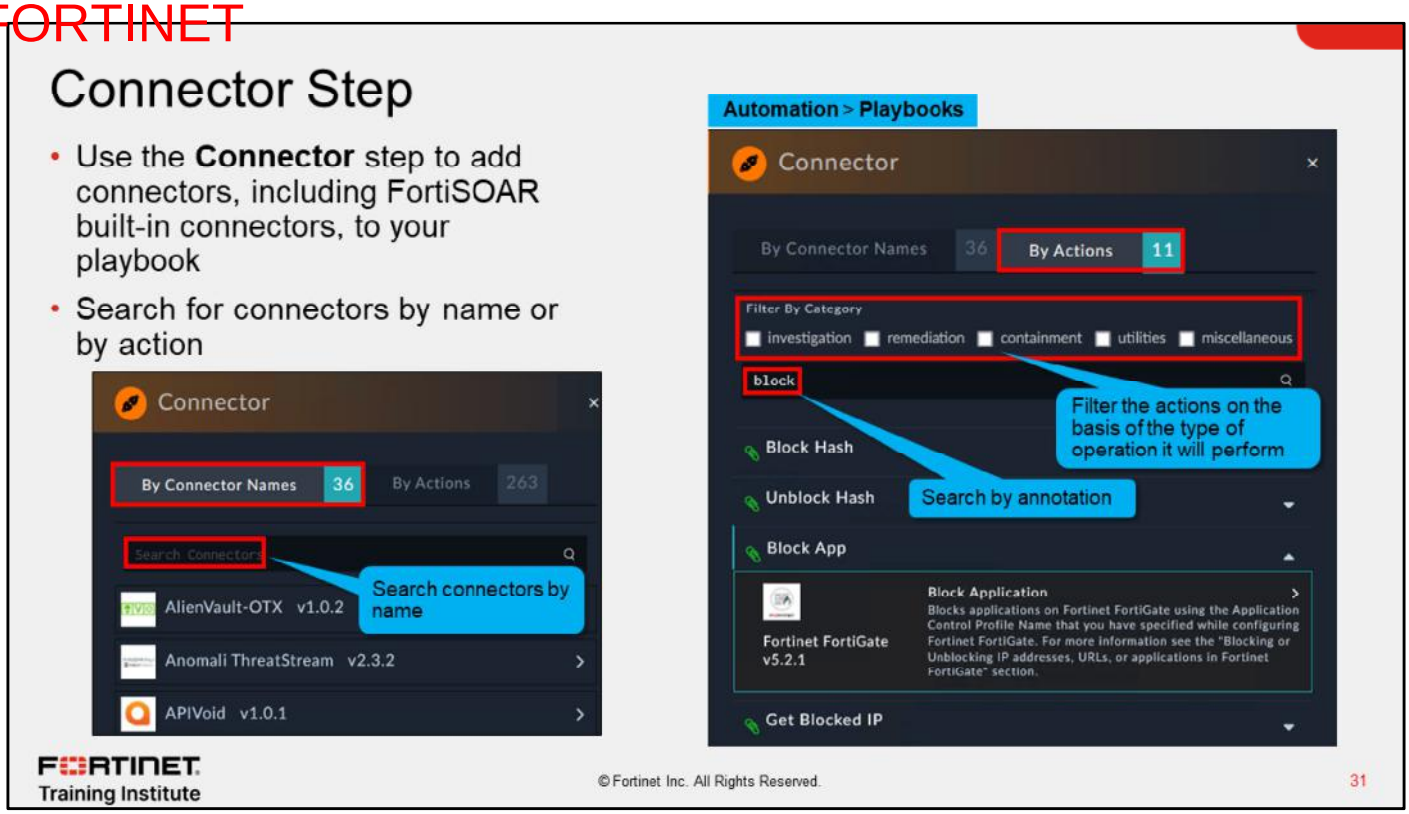
In this section, you will learn about mitigating incidents through FortiSOAR.

DO NOT REPRINT
© FORTINET



Typically, the remediation process on FortiSOAR begins with the extraction of indicators of compromise (IOC) from alerts. There could be several different types of indicators in an alert. Through the indicator extraction playbook, you can extract those indicators. After the indicators are extracted, you must run the enrichment playbook to enrich those indicators by querying a third-party threat intelligence platform. Finally, you can block any malicious indicators through a playbook by connecting to an edge device or to an end-point device.

DO NOT REPRINT
© FORTINET

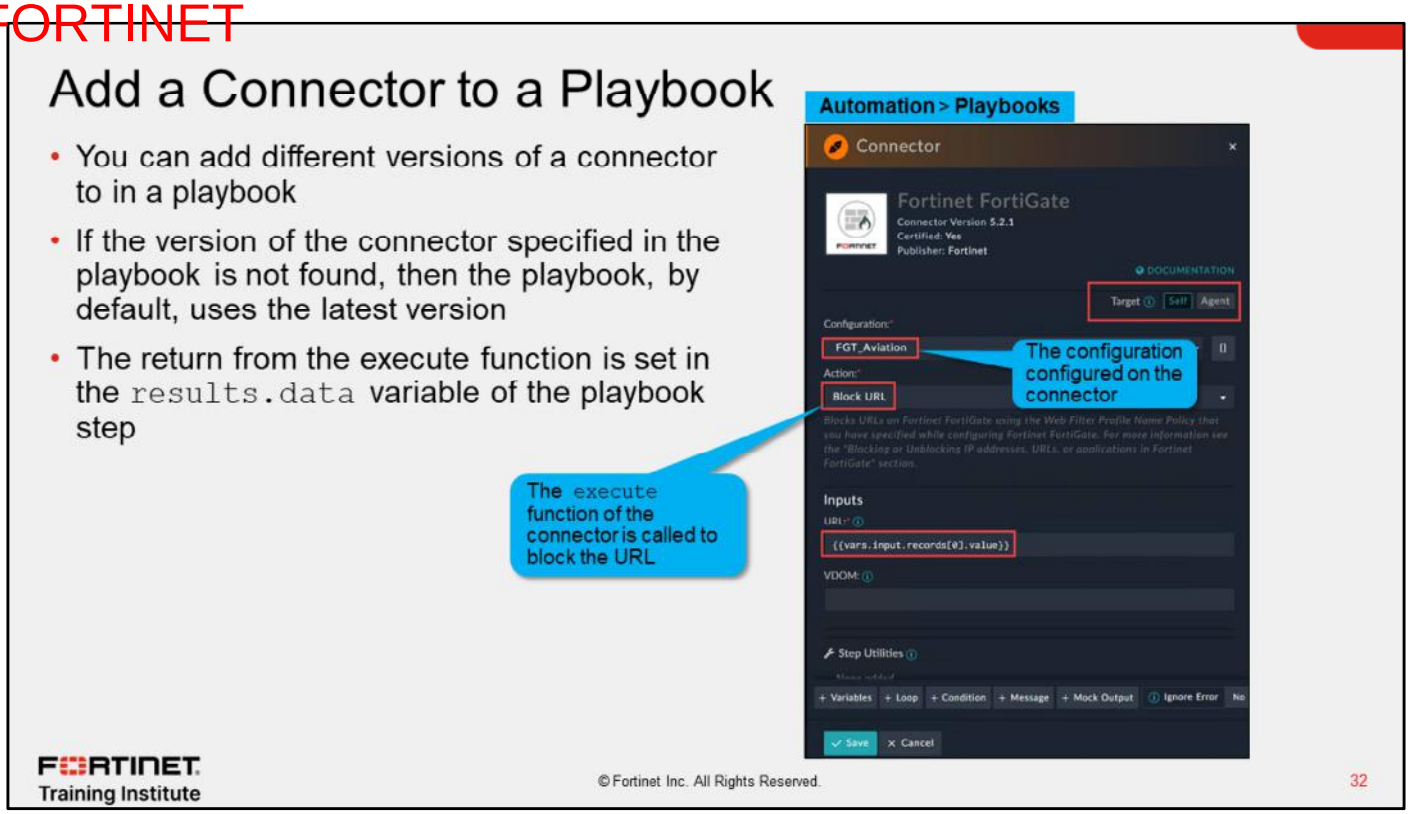


Use the **By Connector Names** tab to first choose a specific connector, and then choose the operation that you want that connector to perform. Or use the **By Actions** tab to first choose the action that you want to perform, and then choose the connector that you want to use to perform the selected action.

The **By Connector Names** tab displays all the connectors that are configured in your system. Use this tab if you want to use a specific connector to perform a particular action. Use the **Search Connectors** section to search for connectors by name.

An annotation can have multiple connectors configured to perform that action, and, if more than one connector can perform the same action, then the GUI displays a list of connectors when you click the name of the action.

DO NOT REPRINT
© FORTINET

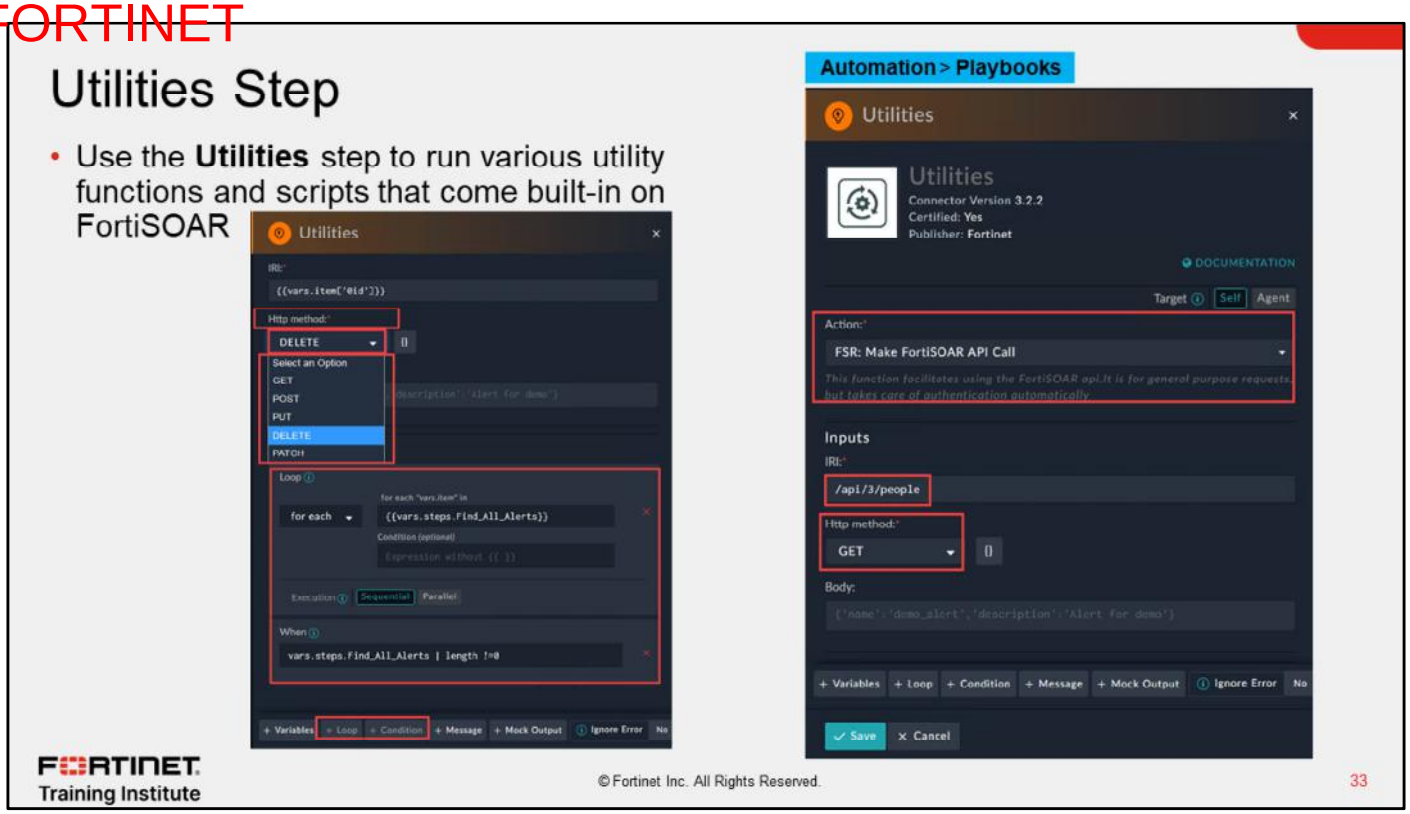


From version 6.4.1 and later, you must specify whether you want to run the action on the current FortiSOAR node, or remotely on the agent node, by clicking the **Self** or **Agent** buttons besides **Target**. By default, **Self** is selected, which means that the action runs on the current FortiSOAR node. Then you must select the configuration by clicking the **Configuration** drop-down list to select the action, because the FortiSOAR node can have multiple configurations. Configurations are based on the configuration names that you specify when you are configuring the connector.

You can install different versions of a connector, and, while adding a connector operation, you can specify a specific version of a connector within a playbook. If the version of the connector specified in the playbook is not found, then the playbook uses the latest version by default.

The return from the execute function is set in the `results.data` variable of the playbook step.

DO NOT REPRINT
© FORTINET

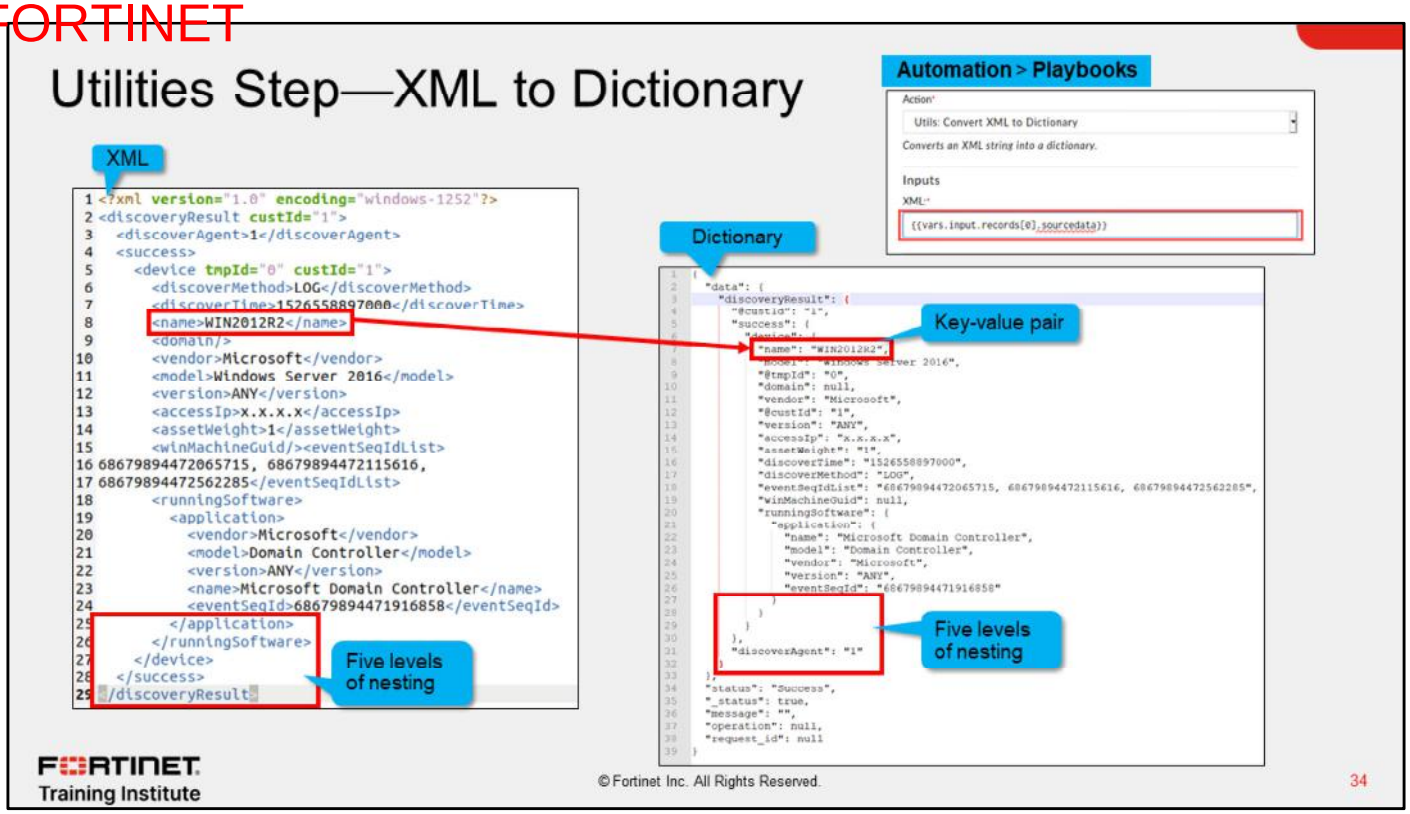


Use the **Utilities** connector for performing operations on FortiSOAR, such as performing a FortiSOAR search using the Query API, updating a FortiSOAR resource, and creating a FortiSOAR resource. This connector also contains other useful utilities, such as extracting email metadata, like indicators, from an email file; uploading a file to FortiSOAR and associating that file with an attachment, such as providing the file IRI in the output; converting file formats, such as XML to JSON or CEF to JSON, and zipping and password protecting a file. This connector is ready to use, and you do not need to configure this connector.

One of the utilities on FortiSOAR is **FSR: Make FortiSOAR API Call**. You can use this action to call an API within FortiSOAR and perform various actions. You can use HTTP methods such as **GET**, **POST**, **PUT**, **DELETE**, and **PATCH**. In the example shown on this slide, a **DELETE** HTTP method is used for every item from a previous step. The **for each** runs for every item in the array that is fetched by the step named **Find_All_Alerts** and then deletes them.

Another example is using the **GET** function of the HTTP method. This retrieves users from the IRI `/api/3/people`.

DO NOT REPRINT
© FORTINET

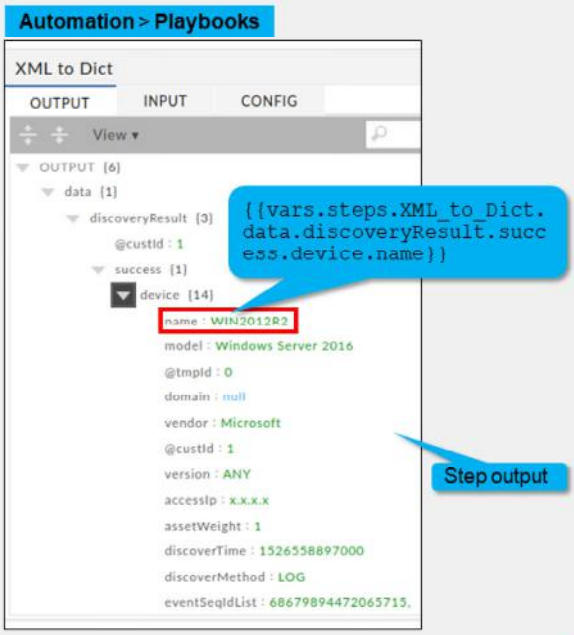


The XML to Dictionary is another useful utility that you can use to convert XML data to dictionary format. The dictionary data is in JSON style and preserves the nesting level structure in XML data. For example, if the XML data has five levels of nesting, the dictionary data also has five levels of nesting. The data within an XML tag is considered to be the value and the XML tag itself is considered the key. In dictionary format, it the value plus the key forms the key-value pair.

DO NOT REPRINT
© FORTINET

Utilities Step—XML to Dictionary Output

- You can access any key-value pair from the dictionary in the playbook environment
- For example, if you want to access the name key, then you must enter the Jinja `{{vars.steps.XML_to_Dict.data.discoveryResult.success.device.name}}`



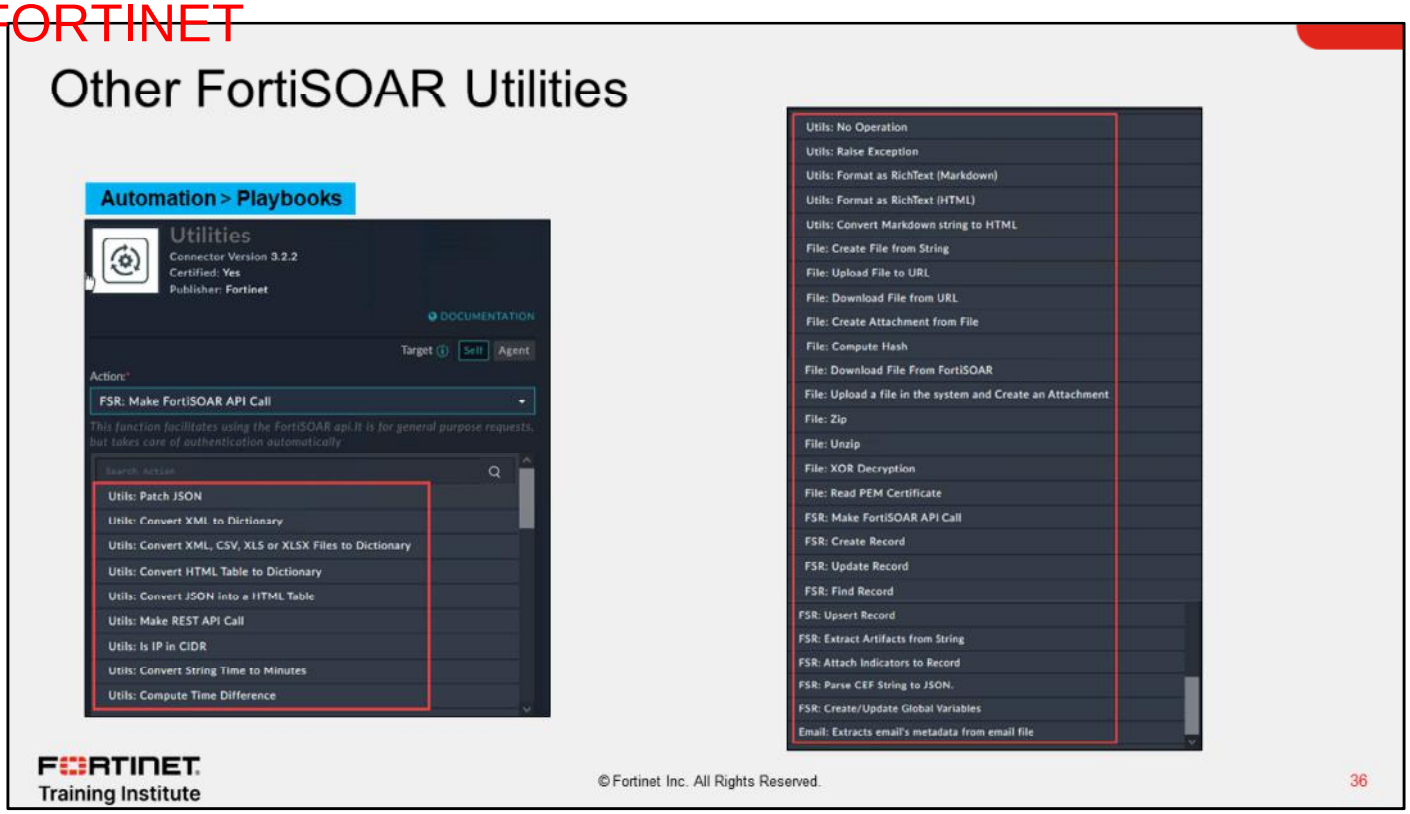
FORTINET Training Institute

© Fortinet Inc. All Rights Reserved.

35

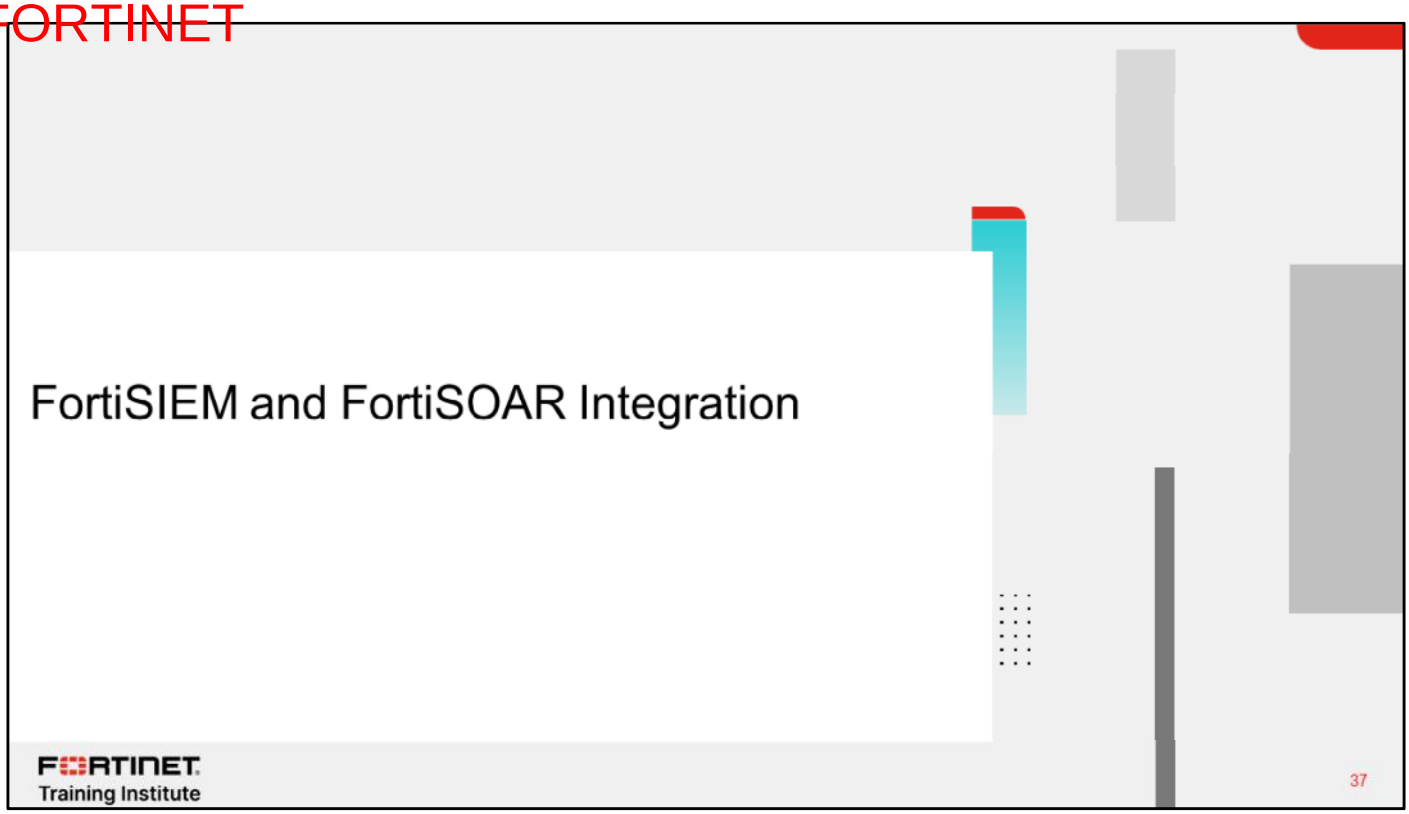
After you convert XML data to the dictionary, you can use that dictionary data in any remaining playbook steps.

DO NOT REPRINT
© FORTINET



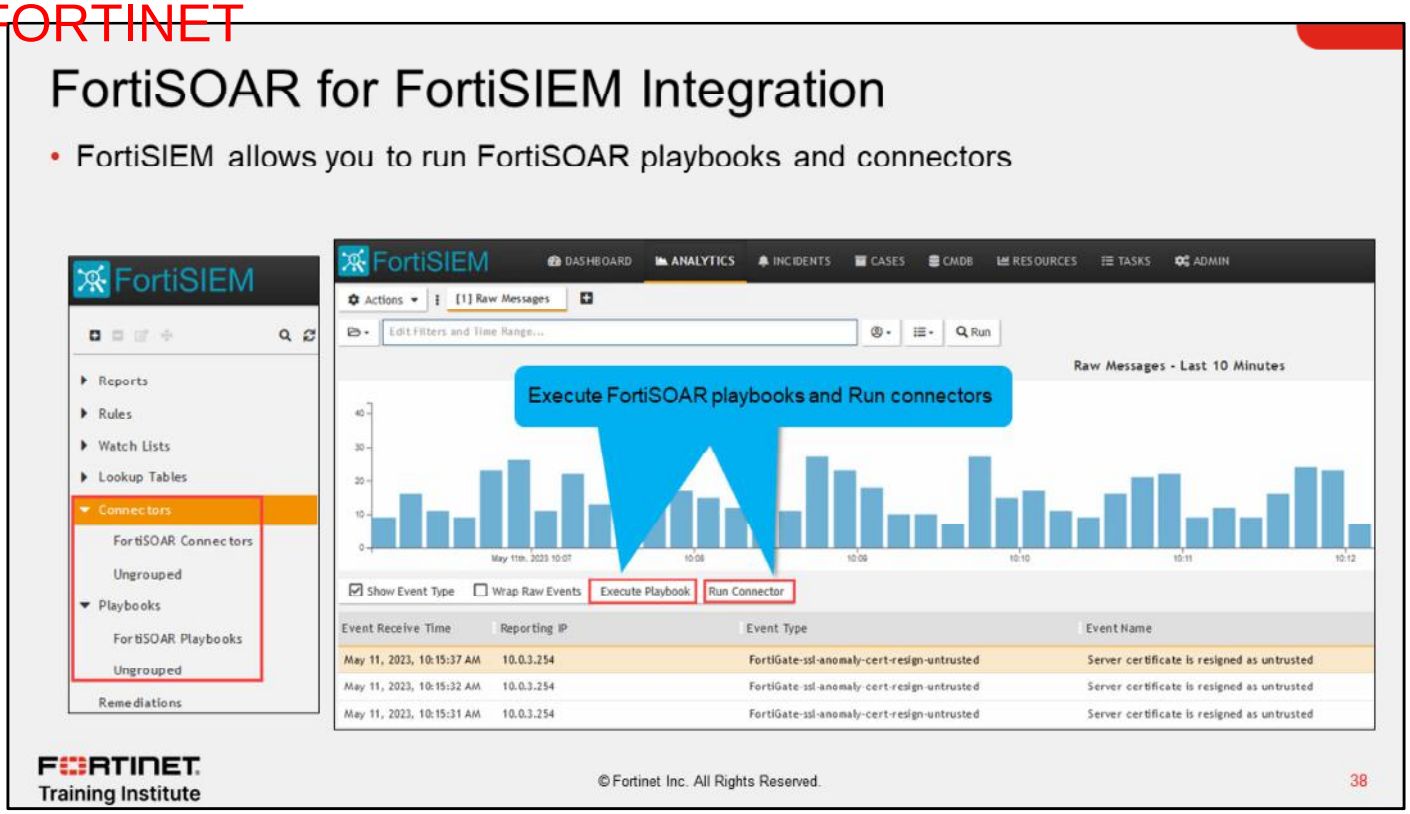
There are other FortiSOAR utilities that you can use in your playbook. This slide lists the other utilities that are available to you while designing a playbook.

DO NOT REPRINT
© FORTINET



In this section, you will learn about leveraging FortiSOAR integration in analytics and incidents on FortiSIEM using FortiSOAR playbooks and connectors.

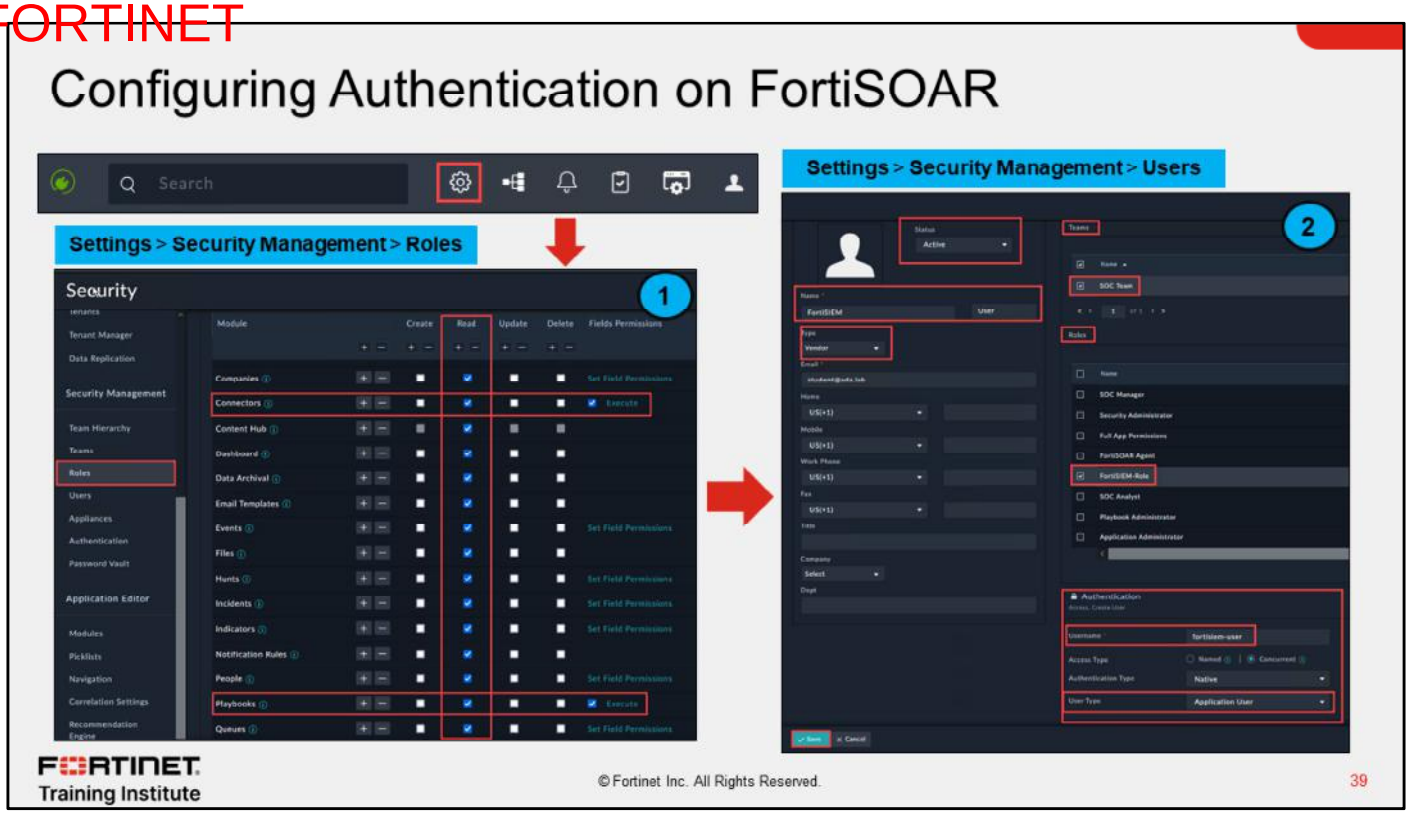
DO NOT REPRINT
© FORTINET



You can integrate with FortiSOAR and execute any connector actions or playbooks. This is very useful because an analyst can leverage the integrations made possible by having FortiSOAR to speed up decisions (by leveraging FortiSOAR through FortiSIEM for Incident Response) or use FortiSOAR to gather more information about a specific IOC. You configure FortiSIEM and FortiSOAR and then leverage that integration from the FortiSIEM incidents section.

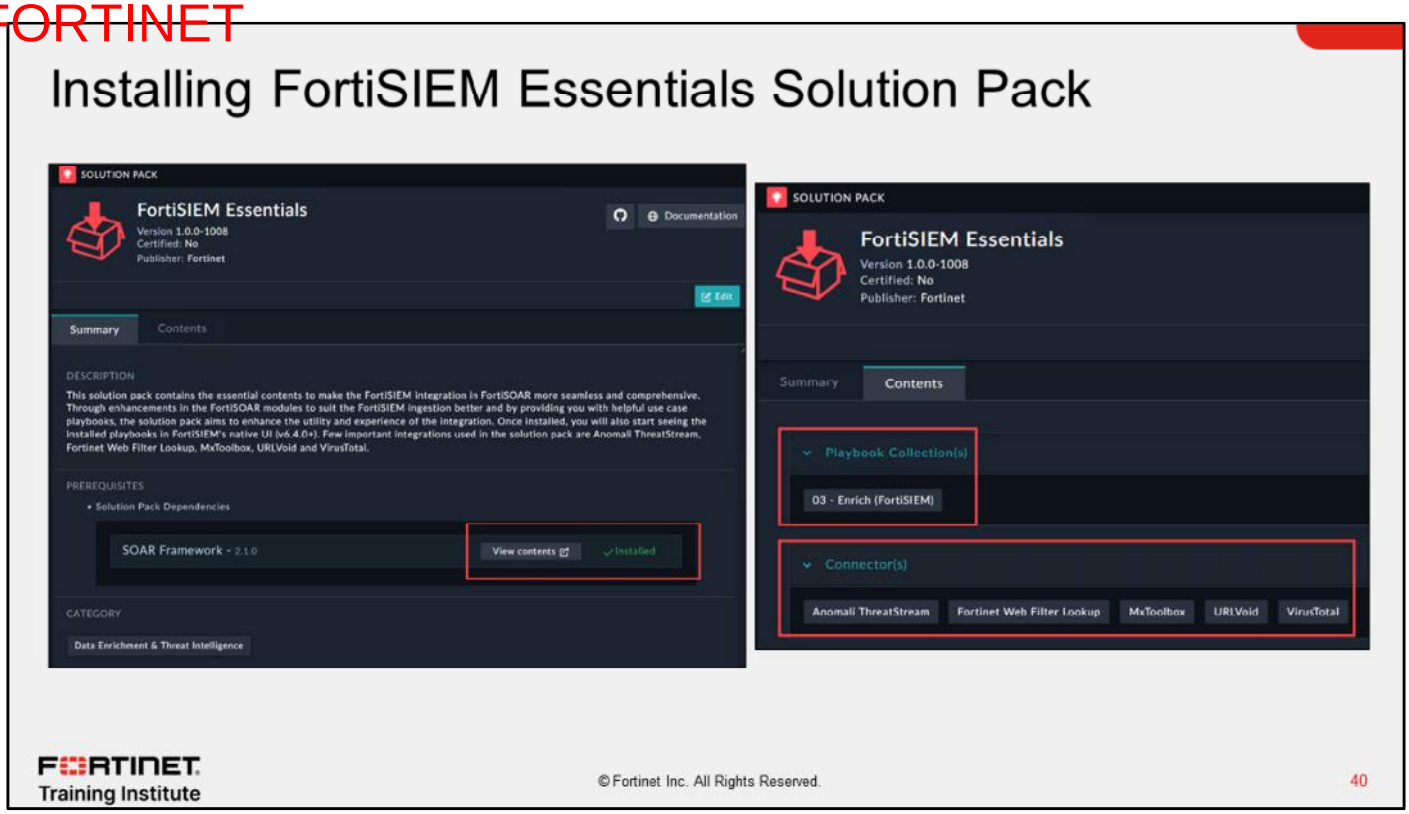
You can use a FortiGate or FortiGuard connector, or you can add other connectors if you want to add the API keys, for example Virustotal, Alienvault OTX, and so on.

DO NOT REPRINT
© FORTINET



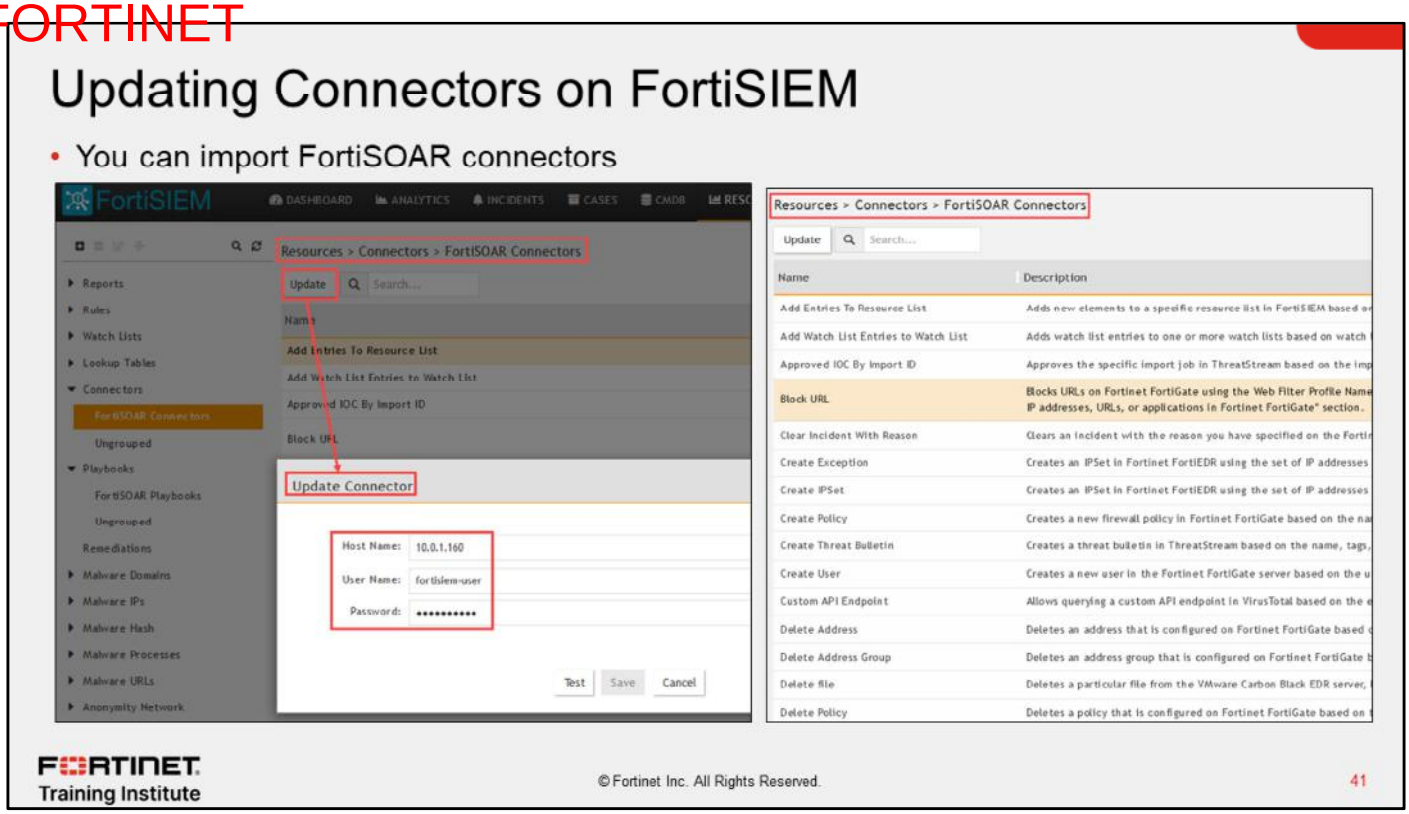
You must set up FortiSOAR with a role assigned to a user account so you can leverage the use of FortiSOAR playbooks and connectors from the FortiSIEM GUI. Use HTTP basic authentication for API for authentication. You will configure a role assigned to a user that FortiSIEM can use to read and run playbooks and connectors.

DO NOT REPRINT
© FORTINET



This solution pack contains the essential contents to make the FortiSIEM integration in FortiSOAR more seamless and comprehensive. Through enhancements in the FortiSOAR modules to suit the FortiSIEM ingestion better, and by providing you with helpful use case playbooks, the solution pack aims to enhance the utility and experience of the integration. Once installed, you will also see the installed playbooks in the FortiSIEM native UI. A few important integrations used in the solution pack are Anomali ThreatStream, Fortinet Web Filter Lookup, MxToolbox, URLVoid, and VirusTotal.

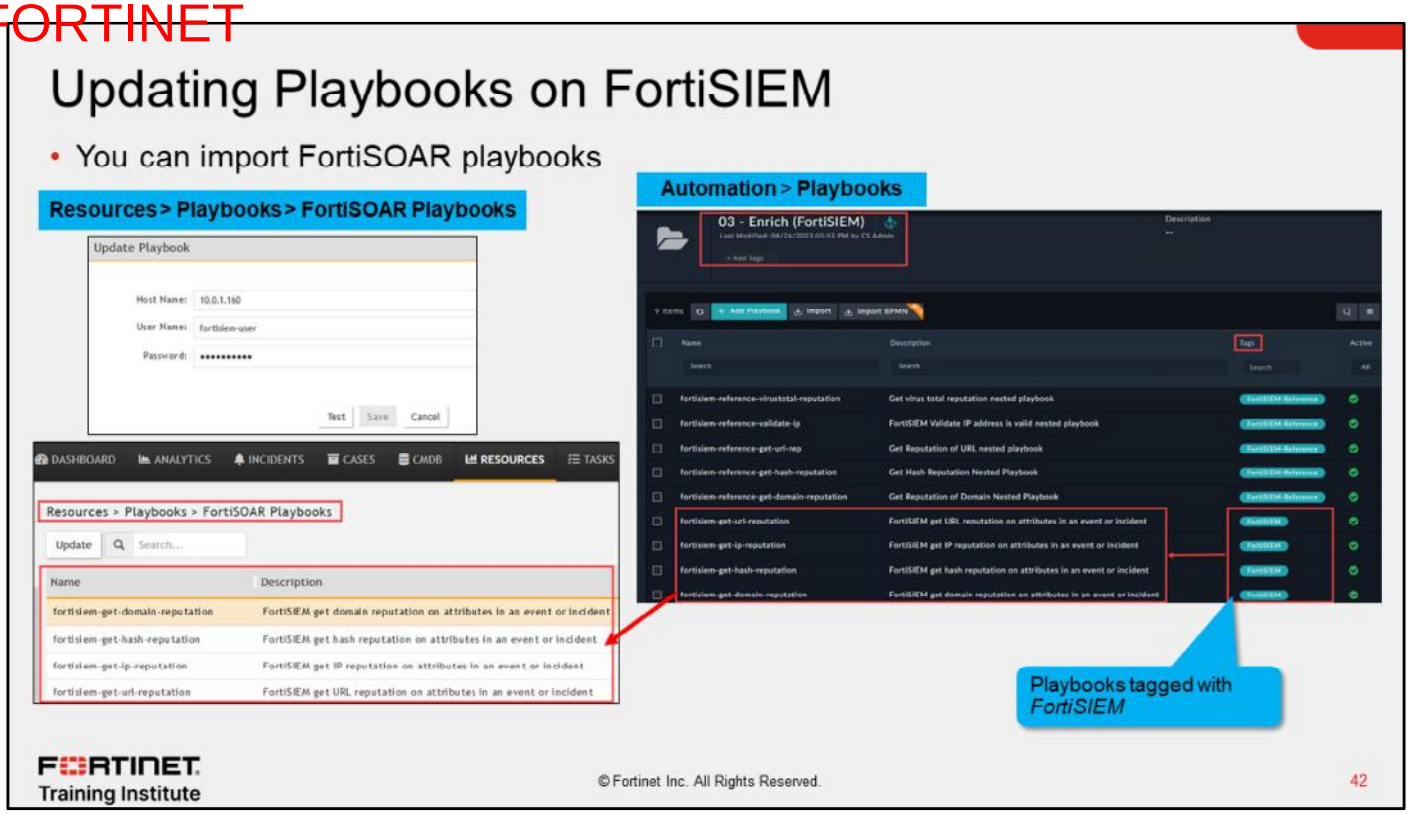
DO NOT REPRINT
© FORTINET



You can use connectors to send and retrieve data from various third-party sources. Using connectors, you can connect to external cybersecurity tools and perform various automated interactions using FortiSOAR playbooks. In this integration, the `fortisiem-user` has permissions to run connector actions, which means that any connector action is available to execute from FortiSIEM.

Note that you must configure connectors on FortiSOAR.

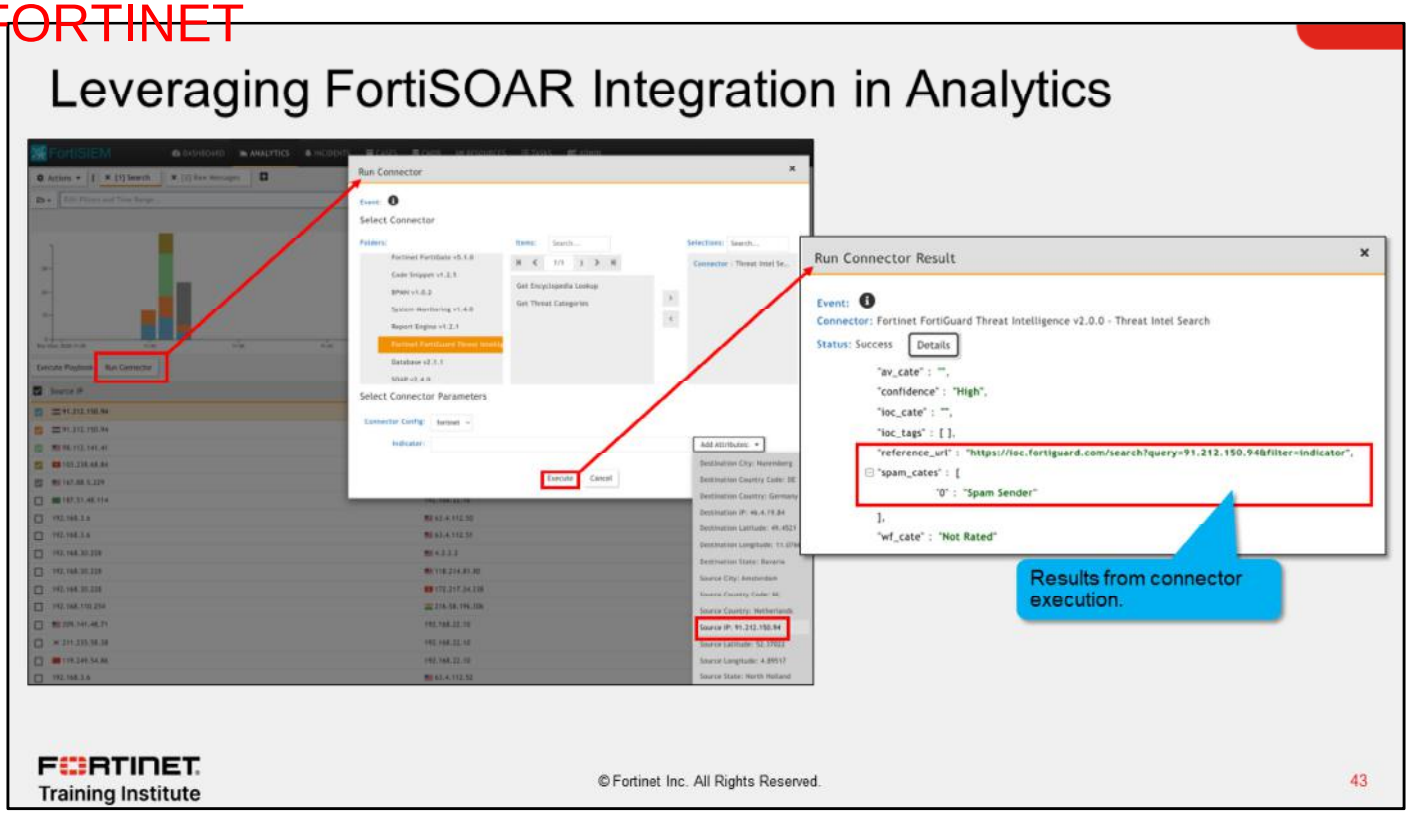
DO NOT REPRINT
© FORTINET



Playbooks on FortiSOAR allow you to automate your security processes across external systems, while respecting the business processes required for your organization to function. You can customize playbook templates to follow an organization's current procedures while leveraging the automation capabilities of FortiSOAR. FortiSIEM imports only playbooks that have the tag "FortiSIEM". You can add multiple tags to playbooks within FortiSOAR, but the integration looks specifically for playbooks tagged FortiSIEM.

Note that you must configure playbooks on FortiSOAR.

DO NOT REPRINT
© FORTINET



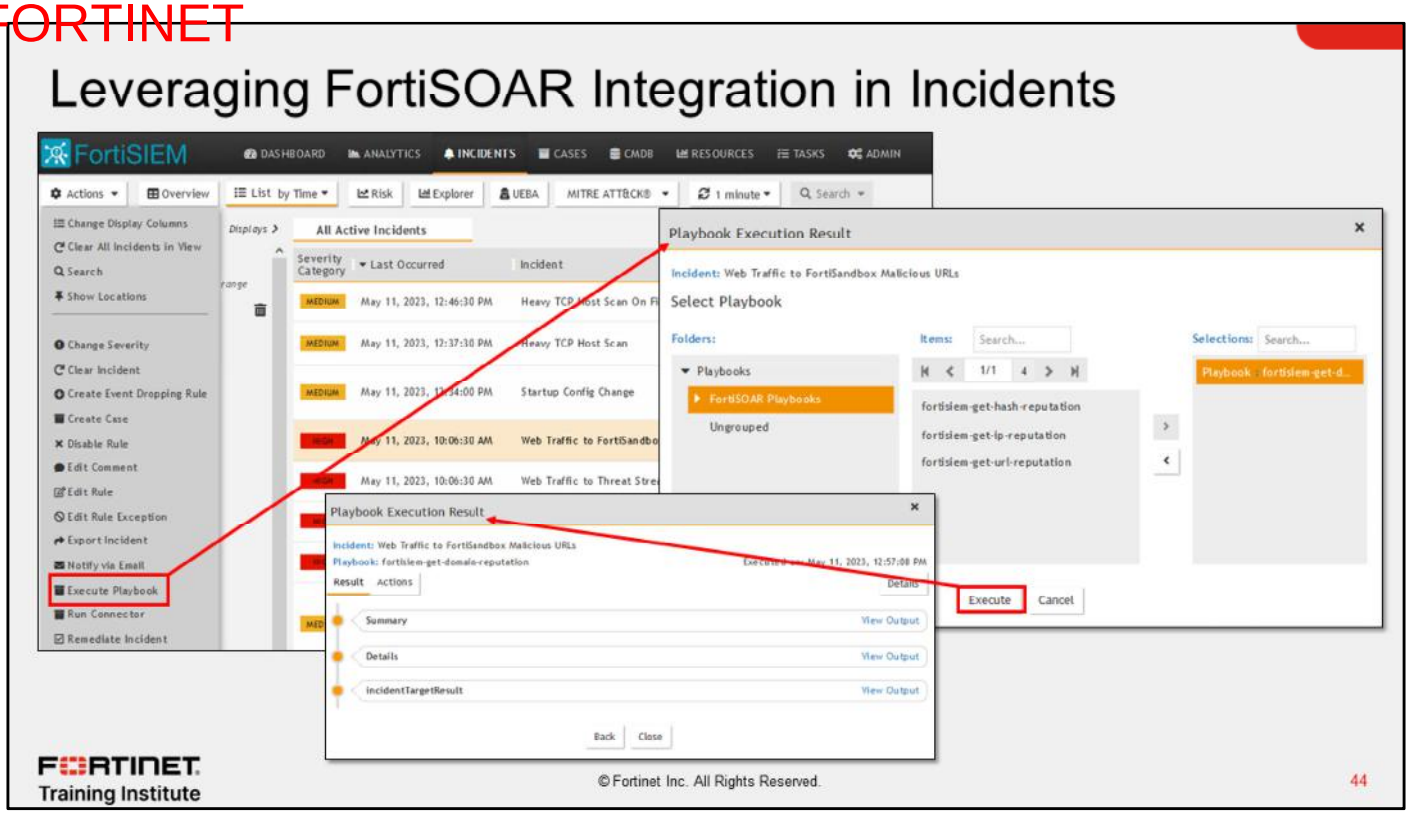
You can execute configured FortiSOAR connectors and playbooks in analytics FortiSIEM. Different use cases will leverage different connectors.

In the example shown on the slide, you want to leverage a FortiSOAR connector to lookup IP address information on a third-party threat service. Use IP address `91.212.150.94`, select the IP address from analytics search results, and then click **Run Connector**.

Select the FortiGuard connector and the **Threat Intel Search** action. In the connector configuration select the available connector configuration **fortinet** in this example, and select the attribute you want to leverage from the normalized log event, which is **Source IP** for this example). Click **Execute** to view the results.

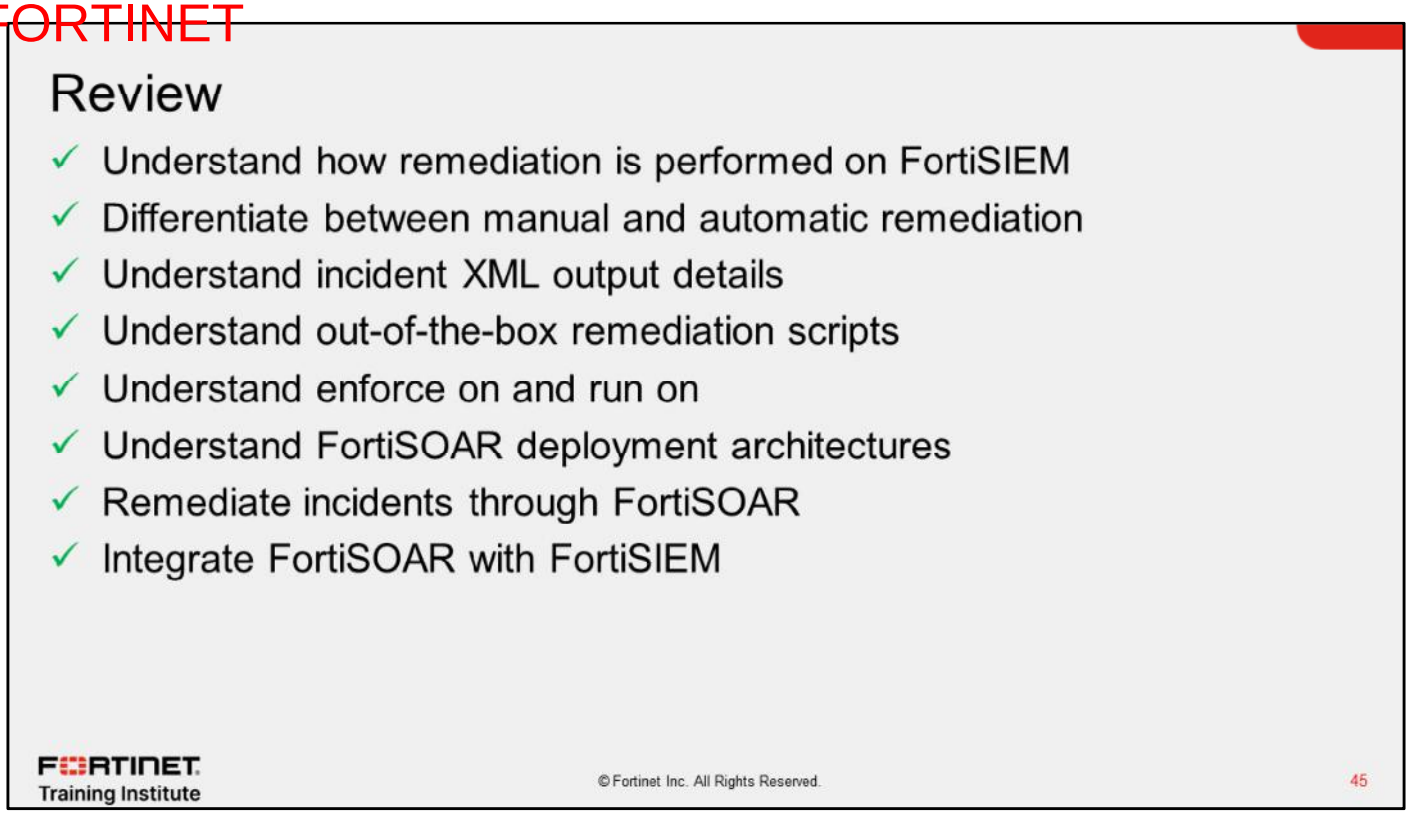
The results from FortiGuard connector for IP address `91.212.150.94`, confirmed as **Spam Sender**. You can use this information for data enrichment purposes on FortiSIEM.

DO NOT REPRINT
© FORTINET



Now you have seen how to leverage the FortiSOAR integration in analytics on FortiSIEM. To obtain more context about the incidents, you can execute playbooks and run connectors on an incident as well. The example shown on the slide shows the execution of a FortiSOAR integrated playbook on an incident on FortiSIEM.

DO NOT REPRINT
© FORTINET



This slide shows the objectives that you covered in this lesson.

By mastering the objectives covered in this lesson, you learned how to remediate incidents manually and automatically.

DO NOT REPRINT
© FORTINET



No part of this publication may be reproduced in any form or by any means or used to make any derivative such as translation, transformation, or adaptation without permission from Fortinet Inc., as stipulated by the United States Copyright Act of 1976.

Copyright© 2023 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., in the U.S. and other jurisdictions, and other Fortinet names herein may also be registered and/or common law trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance and other metrics contained herein were attained in internal lab tests under ideal conditions, and actual performance and other results may vary. Network variables, different network environments and other conditions may affect performance results. Nothing herein represents any binding commitment by Fortinet, and Fortinet disclaims all warranties, whether express or implied, except to the extent Fortinet enters a binding written contract, signed by Fortinet's General Counsel, with a purchaser that expressly warrants that the identified product will perform according to certain expressly-identified performance metrics and, in such event, only the specific performance metrics expressly identified in such binding written contract shall be binding on Fortinet. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. In no event does Fortinet make any commitment related to future deliverables, features, or development, and circumstances may change such that any forward-looking statements herein are not accurate. Fortinet disclaims in full any covenants, representations,and guarantees pursuant hereto, whether express or implied. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.