Cisco Secure Workload and Secure Firewall

March 2024

Contents

Abstract

In a world where application workloads are deployed anywhere at any time, across hybrid multicloud solutions, applying network security controls is no trivial task. The policy control toolset just keeps growing, with multiple enforcement points in the network to protect application workloads using different approaches such as host firewalls, network firewalls, Software-Defined Networking (SDN) controllers, or cloud-based in the form of security groups. Adding to the equation, different teams manage each policy control and often work in organizational siloes. Given this reality, it should be no surprise that these circumstances often lead to inconsistent islands of policy controls across the environment. With Cisco® Secure Workload, organizations can embrace the Zero Trust microsegmentation journey to harmonize different network security policy controls into a consistent, unified policy across hybrid multicloud environments in order to ultimately reduce the attack surface and contain lateral movement.

Target audience

This document provides a technical overview of the design principles, architecture, and use cases for Cisco Secure Workload and Cisco Secure Firewall integration.

The target audience for this document is network engineers, network security engineers, system engineers, security architects, and cloud architects.

Scope

This document covers multiple Cisco Secure Firewall architecture insertion options with their capabilities and related use cases.

## Cisco Secure Workload – Solution overview

Cisco Secure Workload is a holistic security solution designed to deliver in-depth application workload visibility and protection across hybrid multicloud environments. Secure Workload focuses on three main use cases:

- **Zero Trust Microsegmentation:** Using agent and agentless approaches, Secure Workload can discover workloads based on labels, automatically discover and suggest segmentation policies based on traffic flows, validate and test the policy without any operational impact, and enforce the dynamic policy on multiple enforcement points such as host-based firewalls, Data Processing Units (DPUs), network firewalls, load balancers, and built-in cloud security controls.

- **Vulnerability Detection and Protection:** Utilizing an agent, Secure Workload provides visibility into the application workload runtime, enabling the detection of vulnerable packages and vulnerable container images. It then leverages this information using vulnerability (Common Vulnerabilities and Exposures (CVE) attribute-based policies to quarantine workloads or perform virtual patching via Secure Firewall.

- **Behavioral Detection and Protection:** Secure Workload monitors running process for changes in behavior and a detailed process tree and process snapshot. It detects anomalous behavior using MITRE ATT&CK Tactics, Techniques, and Procedures (TTPs) or with custom forensic rules. By leveraging Secure Firewall's Rapid Threat Containment, protection of both agent and agentless workloads can be achieved.
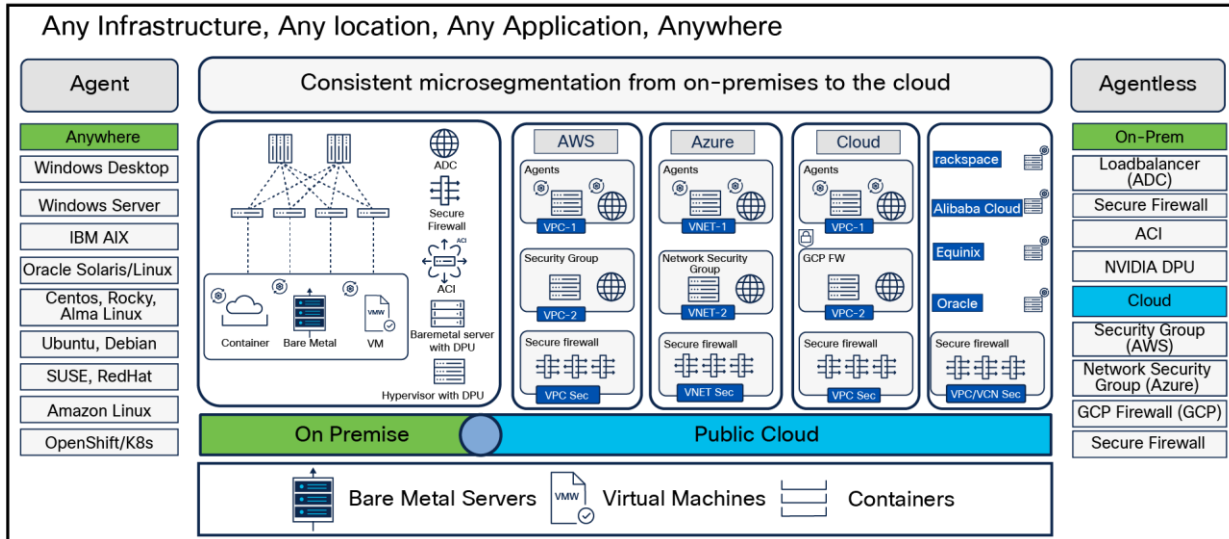
Figure 1.
Secure workload

Cisco Secure Workload and Secure Firewall – Microsegmentation use case
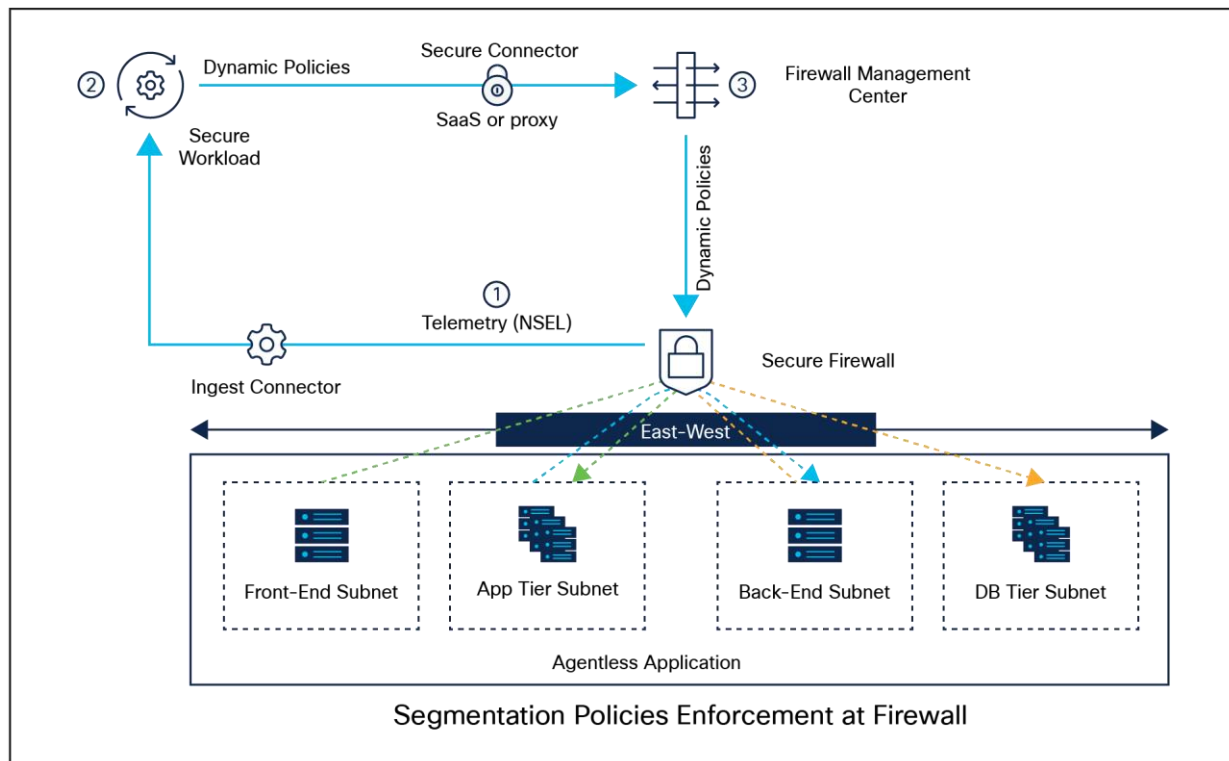
Figure 2.
Secure workload and Secure firewall high-Level architecture

Secure Workload and Secure Firewall provide a flexible means to implement microsegmentation for east-west traffic flows to protect application workloads where agent installation is not feasible. There are three main required capabilities to achieve this:

Visibility of agentless workloads

Secure Workload ingests telemetry from Secure Firewall via NetFlow Secure Event Logging (NSEL) and automatically discovers workloads by leveraging manual labels or external systems labels such as Configuration Management Databases (CMDBs) or IP Address Management (IPAMs). NSEL provides stateful IP flow tracking, including the flow bidirectionality.

- **Ingest connector:** NSEL events are streamed to the Secure Workload Ingest Connector for processing, and the flow data is then exported to Secure Workload. The Ingest Appliance can scale up to 45k fps per Secure Firewall Connector and up to 135k fps for an entire on-prem appliance.
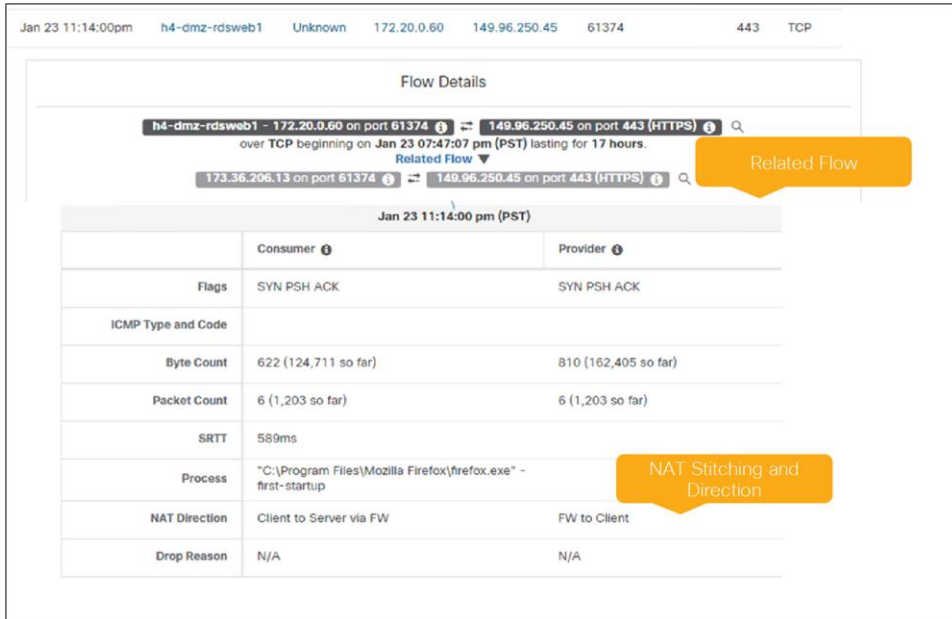

Figure 3.
Secure firewall connector

- **End-to-End Visibility:** Secure Workload is also capable stitching related flows (flow-stitching) to get end-to-end visibility even when Network Address Translation (NAT) is performed.
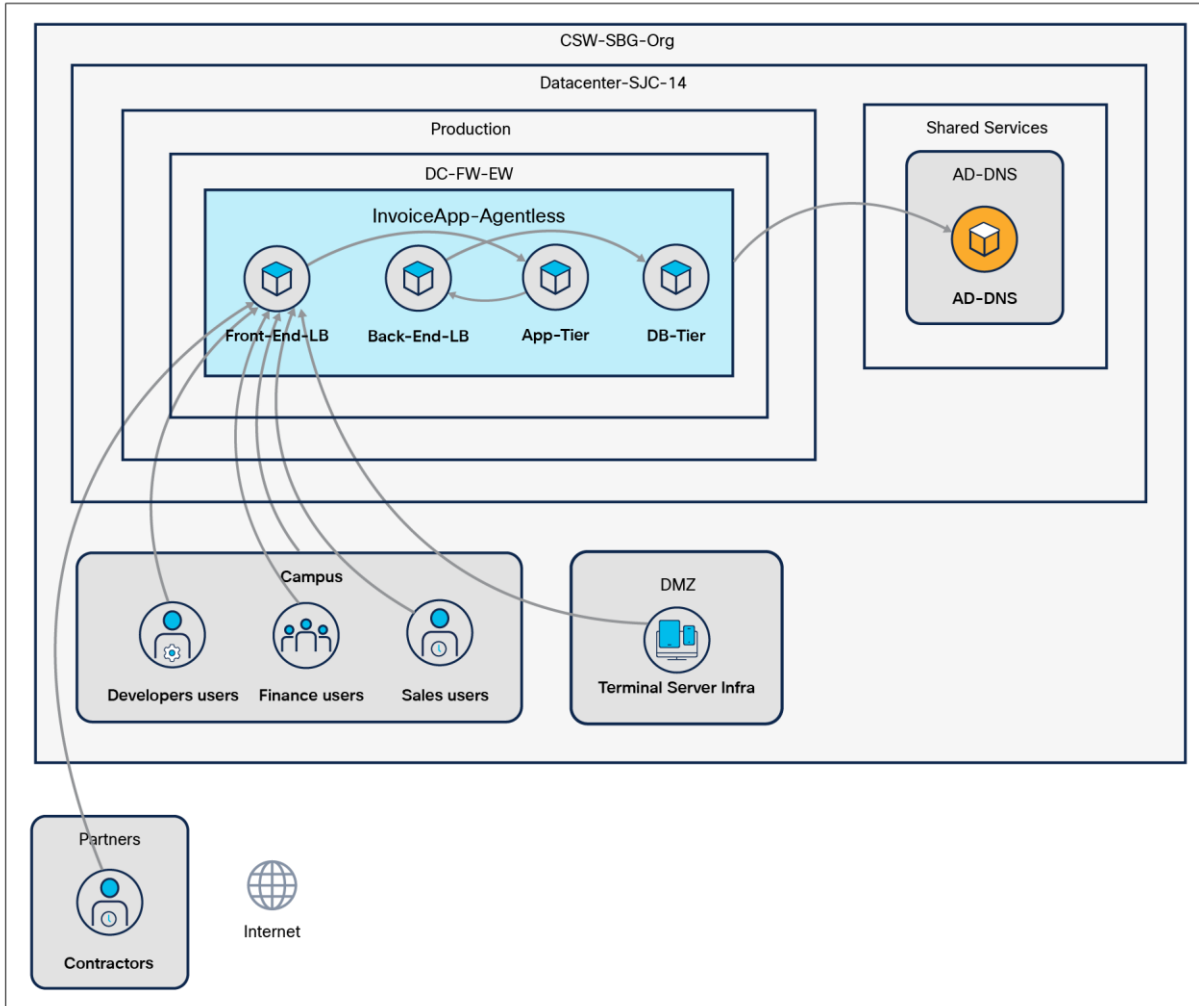
Figure 4.
Flow Stitching with Secure Firewall

Secure workload dynamic policy engine

With Secure Workload, the user can define policies manually or perform automatic policy discovery with Application Dependency Mapping (ADM). Once the policies are validated and enforced, they are pushed to Firewall Management Center (FMC). For the Secure Workload SaaS offering or if it is behind a proxy, Secure Connector for FMC is required.

- **Policy discovery and analysis:** By applying machine learning and behavioral algorithms on the ingested flow data, Secure Workload automatically discovers policies mapped to the application's dependencies. Once policies are discovered, they can be tested and validated against live traffic flow without impacting the application. When this is achieved the policy is ready for enforcement.

Figure 5.
Application dependencies discovered by secure workload
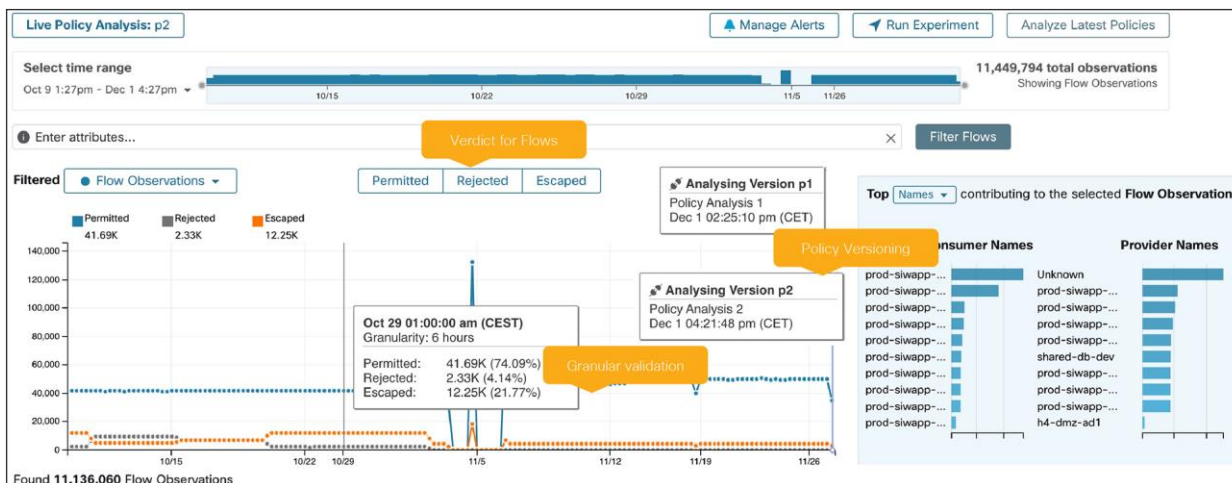
Figure 6.
Secure workload policy analysis toolkit
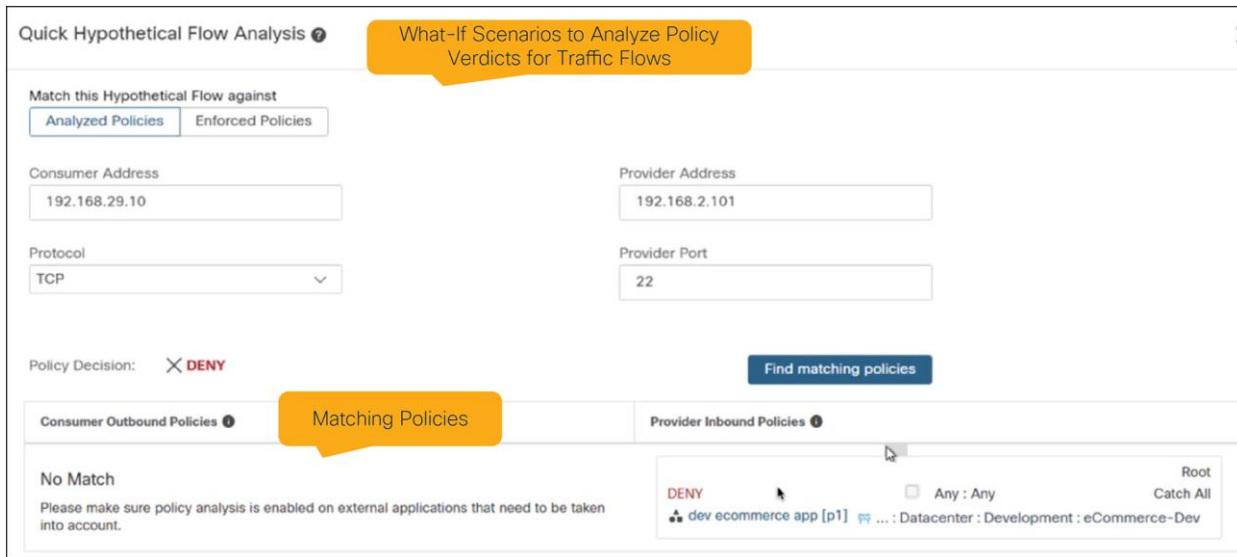
Figure 7.
Policy flow analysis

- **Firewall Management Center (FMC) Onboarding and Enforcement:** Before enforcing the policies, east-west firewalls are onboarded through the FMC Connector. The FMC Connector supports single domain and multi-domain deployments of Secure Firewall. The REST API user configured for Secure Workload must have administrative privileges.
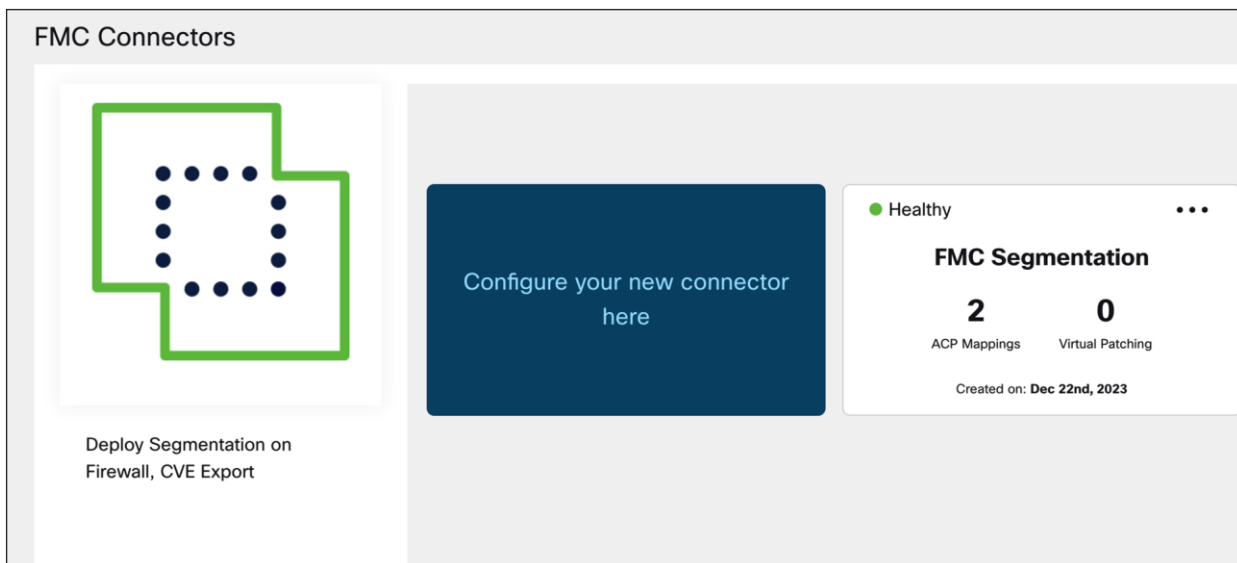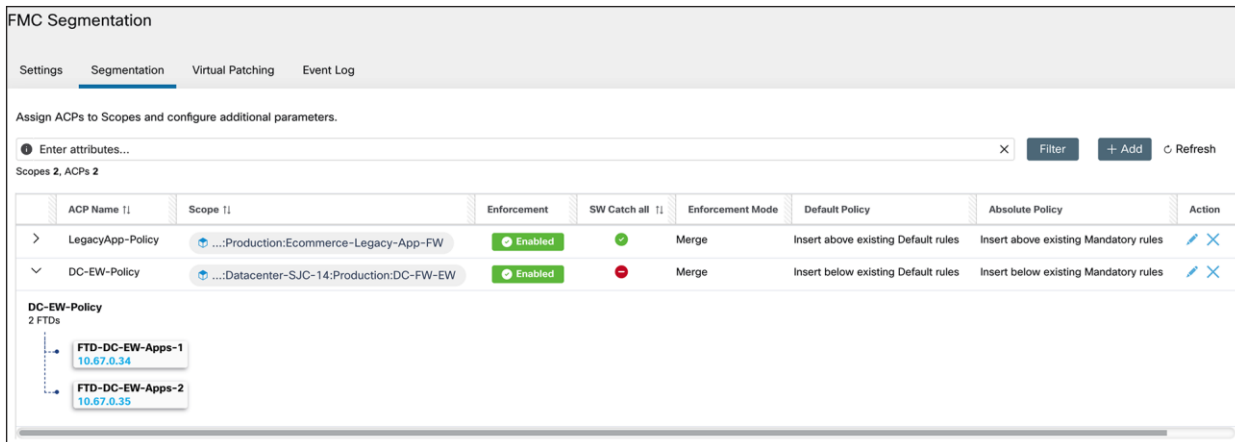

Figure 8.
FMC Connector

- [FMC Connector](#) leverages the concept of topology awareness to push only the rules necessary for enforcement to a specific Secure Firewall. Firewall onboarding happens on an Access Control Policy (ACP) basis by mapping an ACP to a Scope. Each ACP-to-Scope mapping capability can be modified depending on the application or segmentation requirements:
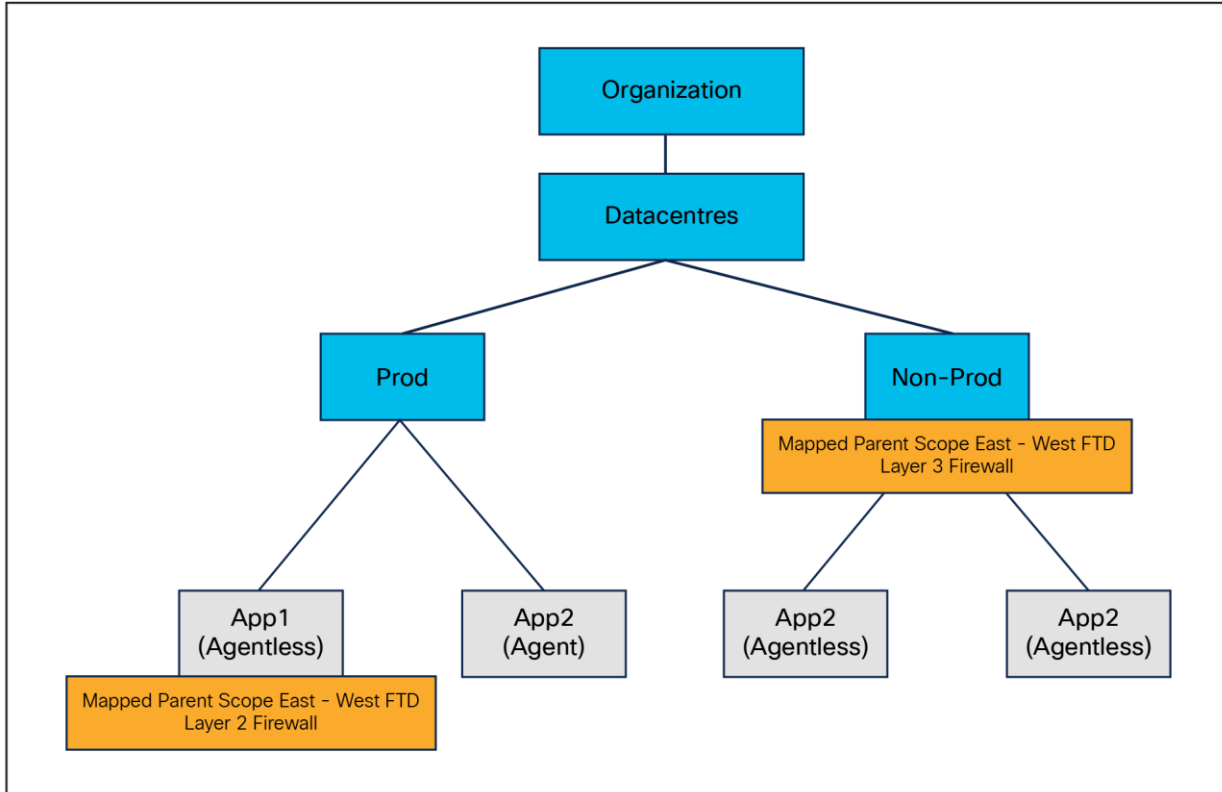
- **Enforcement mode:** With "Merge Mode," Secure Workload honors existing rules on FMC, allowing for policy dual-management. With "Override Mode," Secure Workload removes any existing rules on FMC, only allowing rules pushed by Secure Workload.

- **Rule ordering:** Secure Workload policies can be pushed either on top or bottom of existing FMC rules. "Default Policies" from Secure Workload will be pushed to the "Default Category" on FMC, whereas "Absolute Policies" will be pushed to the "Mandatory Category" on FMC.

- **Optional catch-all:** You can select whether to use Secure Workload "Catch-All" policies or leverage the existing Default action from FMC.

Figure 9.
FMC Connector segmentation use case

- There are two ways to perform the ACP-to-Scope mapping, depending on how many applications are being protected by the firewalls:

  - **Child scope / Single application:** Mapping a child/leaf scope is done when only one application is being protected by Secure Firewall. In this case, Secure Workload pushes only policies that belong to the child scope and any other parent/higher scope policy guardrail.

  - **Parent scope / Multiple applications:** Mapping a parent scope is done when multiple applications are being protected by Secure Firewall. In this case, Secure Workload pushes policies that belong to the mapped scoped but can also push policies from child/leaf scopes that are below the mapped scoped. This is done by virtue of the hierarchical policy model, inheriting the policies of child scopes. Parent policy guardrails will also be pushed.

Figure 10.
Scope tree topology

Figure 11.
ACP-to-scope mapping to parent scope
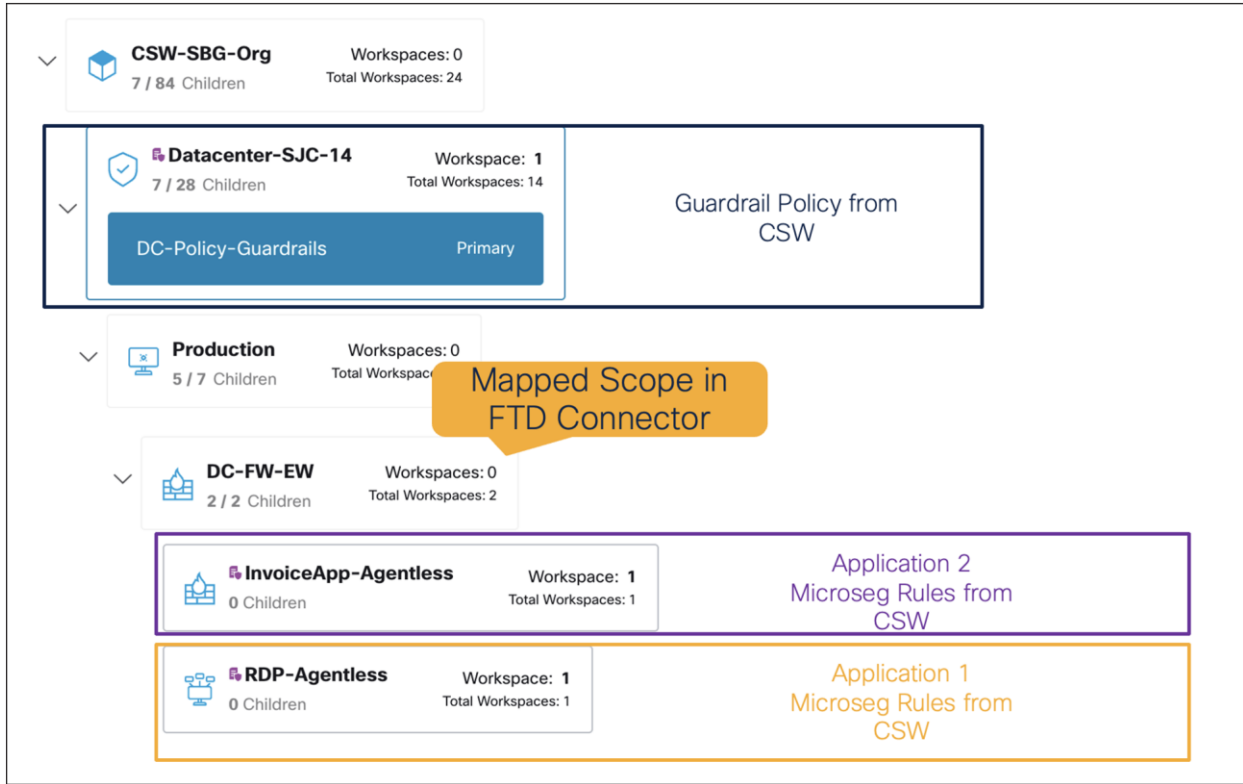
Figure 12.
Scope structure and mapped scope to onboard multiple applications
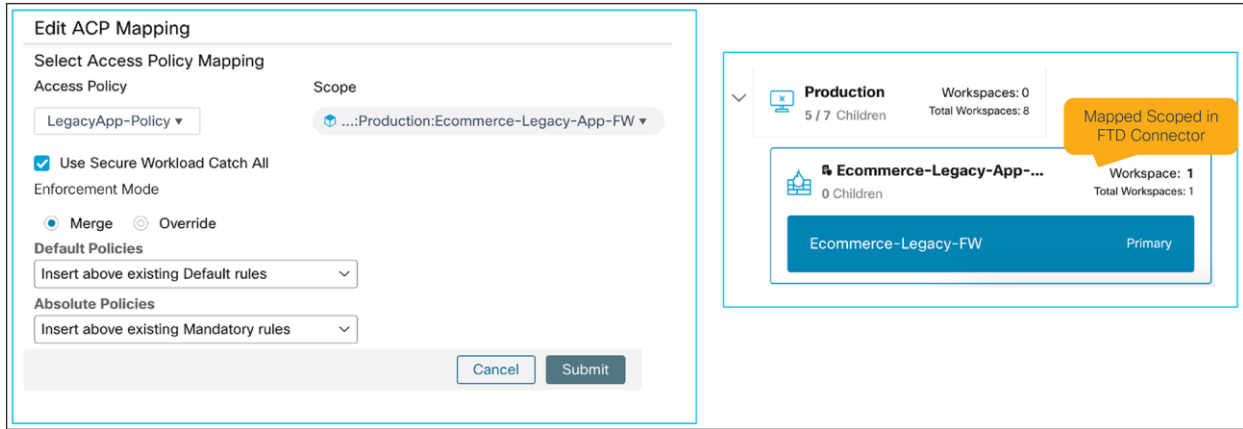
Figure 13.
ACP-Scope mapping for single application

Enforcement on FMC and Monitoring

Policies orchestrated from Secure Workload leverage FMC Dynamic Objects, so the policy is dynamic and doesn't require new policy deployments if an object changes.

- **Dynamic objects:** FMC will push the orchestrated dynamic policies from Secure Workload to the relevant Secure Firewalls in the environment.

Figure 14.
Multiple secure workload application policies pushed to FMC

- **Policy compliance:** The policy is constantly monitored to verify compliance. Alerts and reports can be generated for policy deviations to rapidly investigate and mitigate anomalies.
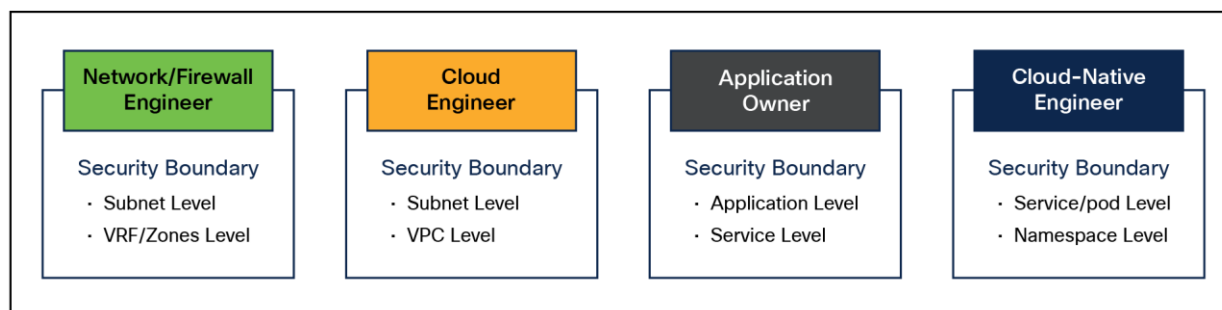

Figure 15.
Compliance alerts for rejected flows

Workload protection level definition

Before selecting the firewall insertion mode, defining the workload protection level based on the department or persona security/trust boundary is advised. The following outcomes are derived from this process:

- **Simplicity and abstraction:** Defining the persona security/trust boundary creates a bridge connecting the business requirements with the technical requirements. It also helps to abstract the intrinsic complexities of heterogenous environments, so the business outcome is more easily trackable.

- **Common language for different personas:** Microsegmentation can have different definitions based on the persona or department tasked with it. The definition must consider each department/persona security/trusted boundary, so it is understood by the whole organization.

- **Creates consistency:** By defining a persona-based security/trust boundary, the resulting segmentation controls on the network will be consistent, allowing each persona or department a deep understanding of the segmentation controls and their limitations, if any.

- **Prepares path for approach selection:** Depending on the persona security/trust boundary, an agent or agentless approach may be used. Agent-based is best when the constructs defined by the persona are not network-based, whereas agentless is typically a good fit for personas using network constructs as security/trust boundaries.

The image below shows an example of how a workload protection level definition looks like depending on the persona's security/trust boundary

Figure 16.
Example of workload protection level definition based on different persona's security/trust boundary

For the purpose of this document, the personas who manage and operate the Secure Firewall network security controls are the network/firewall engineers and/or cloud engineers, and the workload protection level security boundary in use is the subnet level.
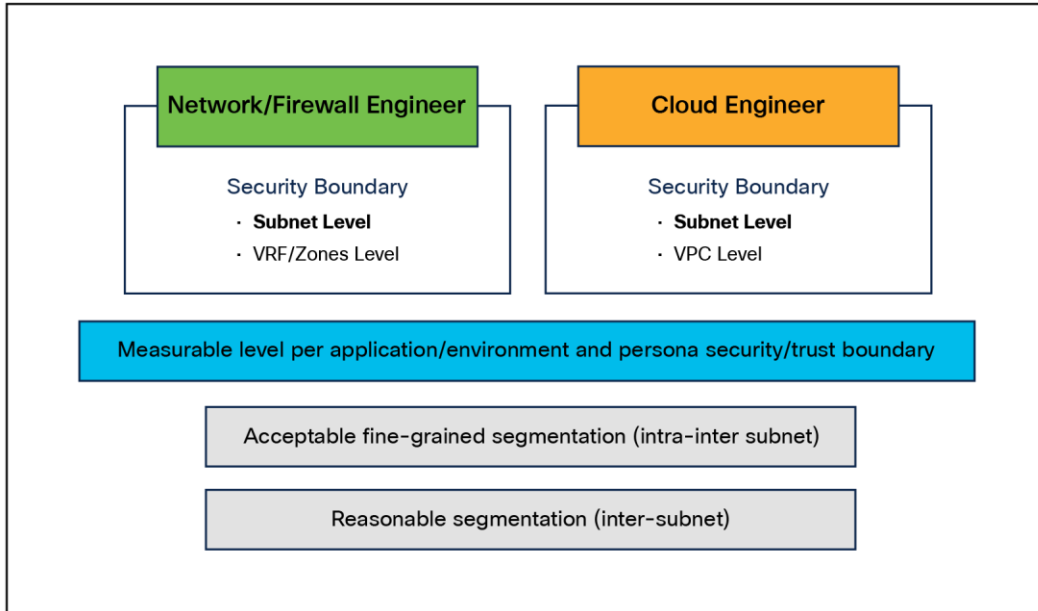
Figure 17.
Measurable workload protection level based

Cisco Secure Workload and Cisco Secure Firewall insertion options - on-premises

**Layer 2 Firewall (Transparent Mode) Insertion**
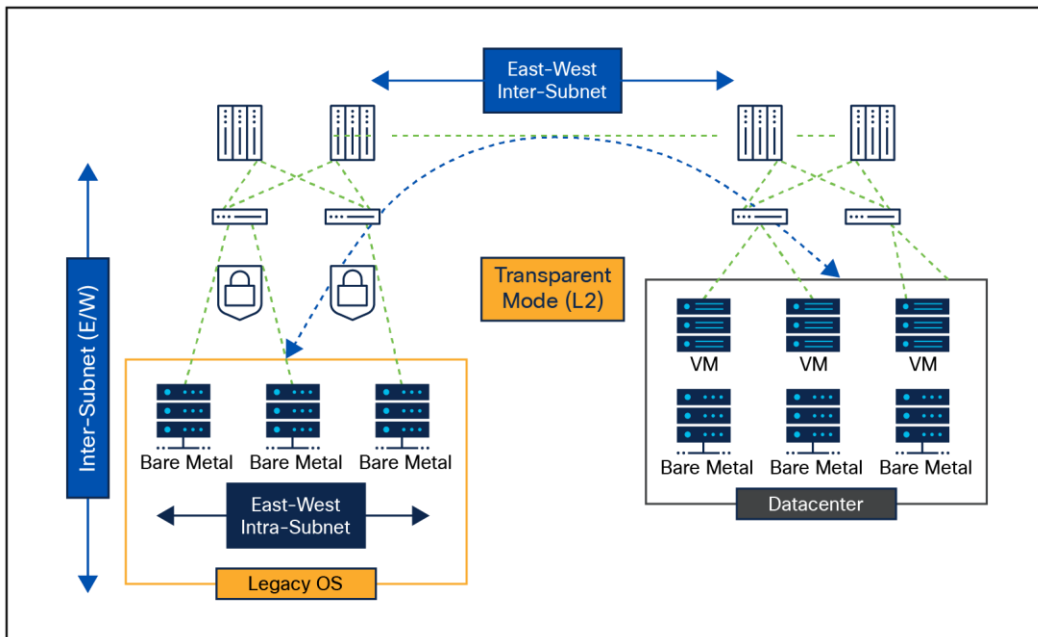
Figure 18.
Network microsegmentation for agentless workloads with layer 2 firewall

- Best fit for **localized workloads**
- Acceptable for fine-grained segmentation
  - Firewall as bump-in-wire on the data path

- Workloads that require fine-grained segmentation, but an agent cannot be installed, such as **legacy OS workloads**

- Protection at the network level

- Full flow visibility with NSEL

  - Intra- and inter-subnet flows

- Protection at the network level

  - Intra-subnet (App-App)

  - Inter-subnet (App-App and External-App)

- Allows policy dual management

  - Secure Workload-owned policies

  - FMC-owned policies

- Convenient for network and firewall engineers

Layer 3 Firewall (Routed Mode) Insertion
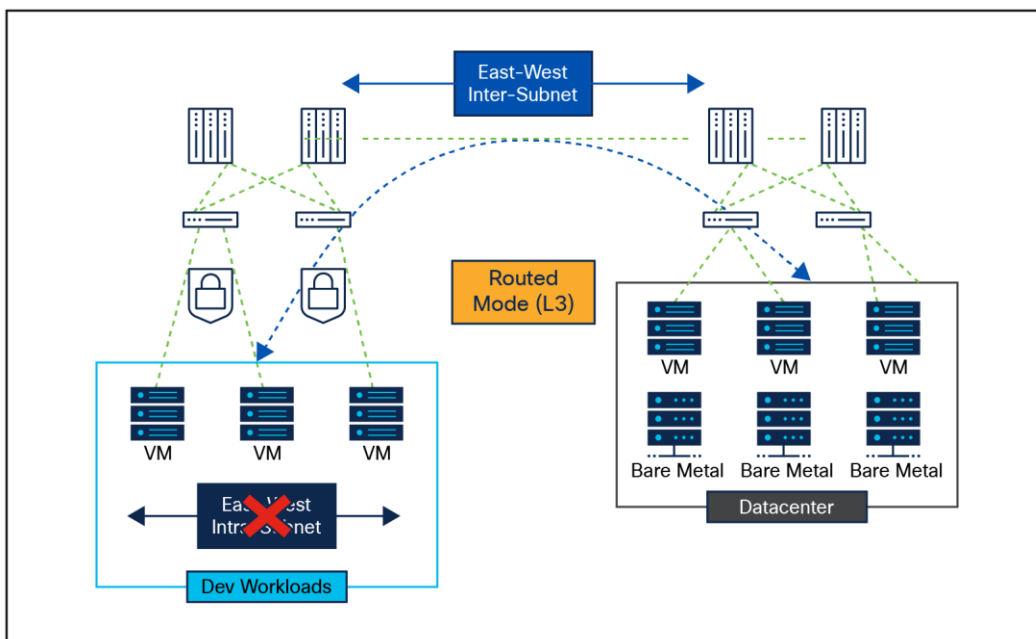
Figure 19.
Network Microsegmentation For Agentless Workloads With Layer 3 Firewall

- Excellent fit for **distributed workloads**

- Reasonable segmentation for workloads

  - Firewall as gateway

  - **Quick time-to-segment**. Good for segments that do not require fine-grained segmentation (such as non-production/development)

- Protection at the network level

- Inter-subnet only (App-App and External-App)
- Partial flow visibility with NSEL
  - Inter-subnet flows only
- Allows policy dual-management
  - Secure Workload-owned policies
  - FMC-owned policies
- Convenient for network and firewall engineers

Cisco Application Centric Infrastructure (Cisco ACI®) Insertion
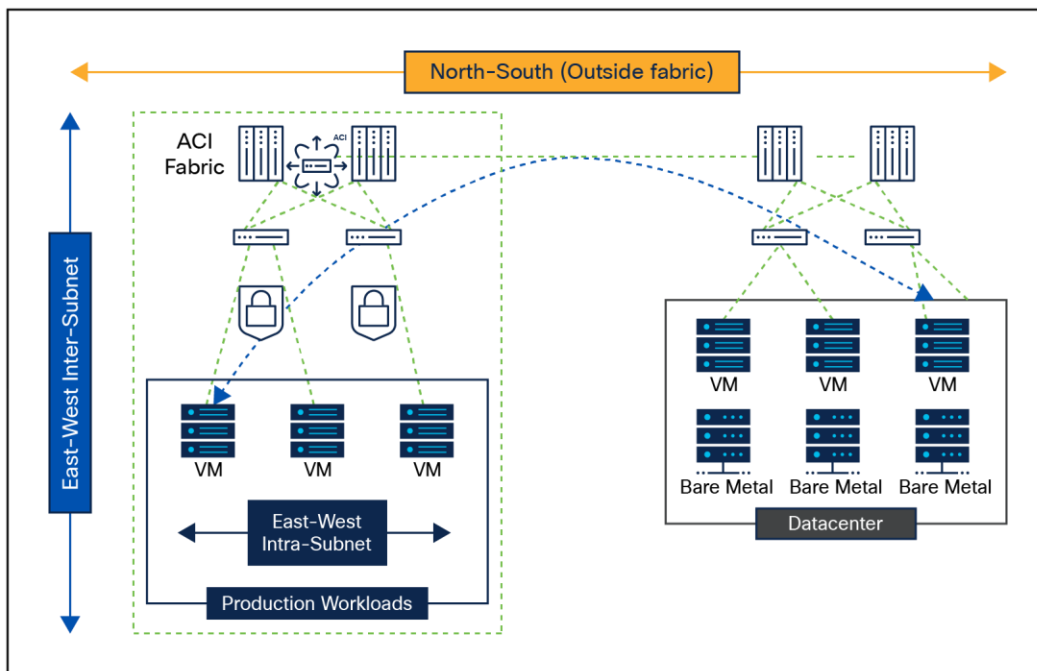
Figure 20.
Network microsegmentation for agentless workloads in ACI

**Service graph with policy-Based redirect**

- **No re-architecture**
  - Flexible and easy to configure
  - Firewall is selectively inserted in the path
- Supports both L3 and L2 firewall modes
  - Intra- and inter-subnet flow visibility (both)
  - Intra- and inter-subnet protection (both)
  - Preferred L3 mode
- Can do intra-ESG redirection

Service Graph Go-To/Go-Through Mode

- Firewall is in-path (Security over Connectivity)
    - Not very flexible and more complex
    - Typically used for North-South traffic
- Go-To
    - Inter-subnet visibility and protection
- Go-Through
    - Intra- and inter-subnet visibility protection

Service Graph Policy-Based Routing (PBR) and Firewall Insertion Protection
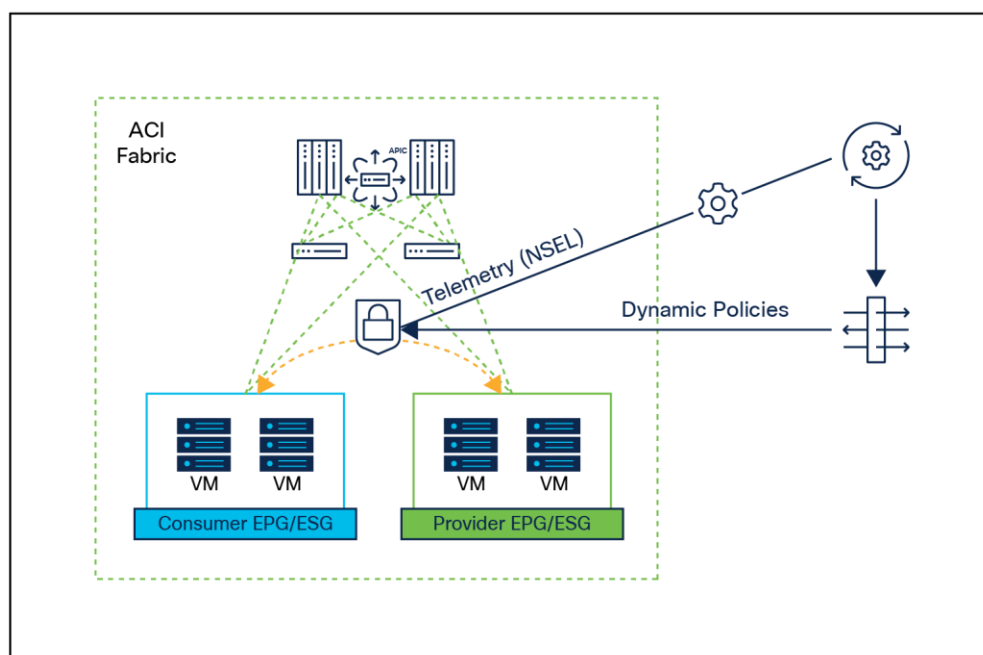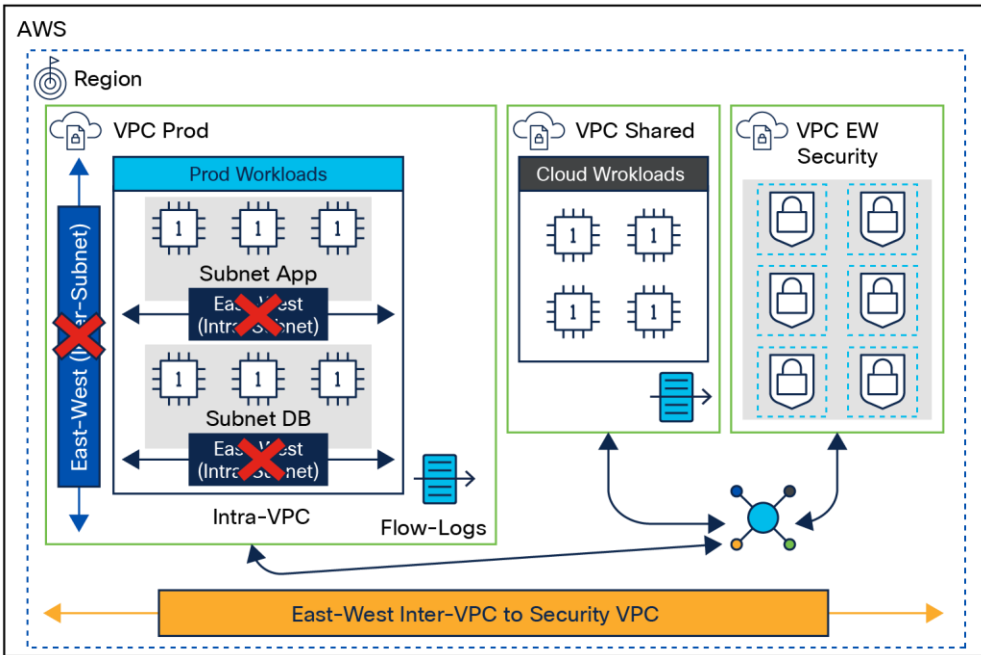

Figure 21.
Network microsegmentation for agentless workloads with service Graph PBR in ACI

- Flexible segmentation for workloads
    - Acceptable fine-grained
    - Reasonable
- Full visibility of flows with NSEL
    - Firewall inserted in data path with service graph
    - Intra and inter Endpoint Group (EPG)/Endpoint Security Group (ESG)
- Protection at network level
    - Intra EPG/ESG (intra-app)
    - Inter EPG/ESG (inter-app)

- Allows policy multi-management

  ◦ Secure Workload-owned policies

  ◦ FMC-owned policies

  ◦ ACI-owned policies

- Convenient for network (ACI) and firewall engineers

Cisco Secure Workload And Cisco Secure Firewall insertion options - Cloud

AWS Centralized East-West Insertion

Figure 22.
Network microsegmentation for cloud agentless workloads with centralized/Hub VPC secure firewall deployment on AWS

- Reasonable segmentation

  ◦ Full flow visibility with Virtual Private Cloud (VPC) flow logs and NSEL

  ◦ Intra- and inter-subnet flows

- Protection at the network level

  ◦ Inter-VPC / inter-subnet

    ◦ App-App and External-App

- FMC policy dual management

  ◦ East-West (Secure Workload +FMC)

  ◦ North-South – Ingress/Egress (FMC)

- Suitable for network/firewall engineers

## Amazon Web Services (AWS) Distributed East-West Insertion
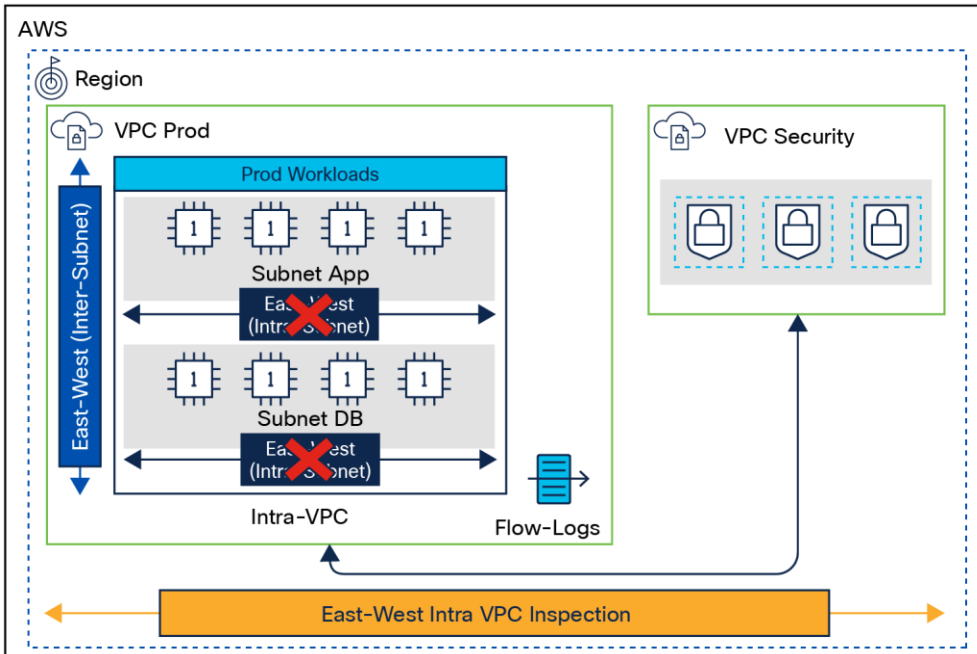
Figure 23.
Network microsegmentation for cloud agentless workloads with distributed VPC secure firewall deployment on AWS

- Reasonable segmentation
- Full flow visibility with VPC flow logs and NSEL
  - Intra- and inter-subnet flows
- Protection at the network level
  - Intra-VPC / inter-subnet
    - App-App and External-App
- FMC policy dual management
  - East-West (Secure Workload + FMC)
  - North-South – Ingress/Egress (FMC)
- Suitable for network/firewall engineers
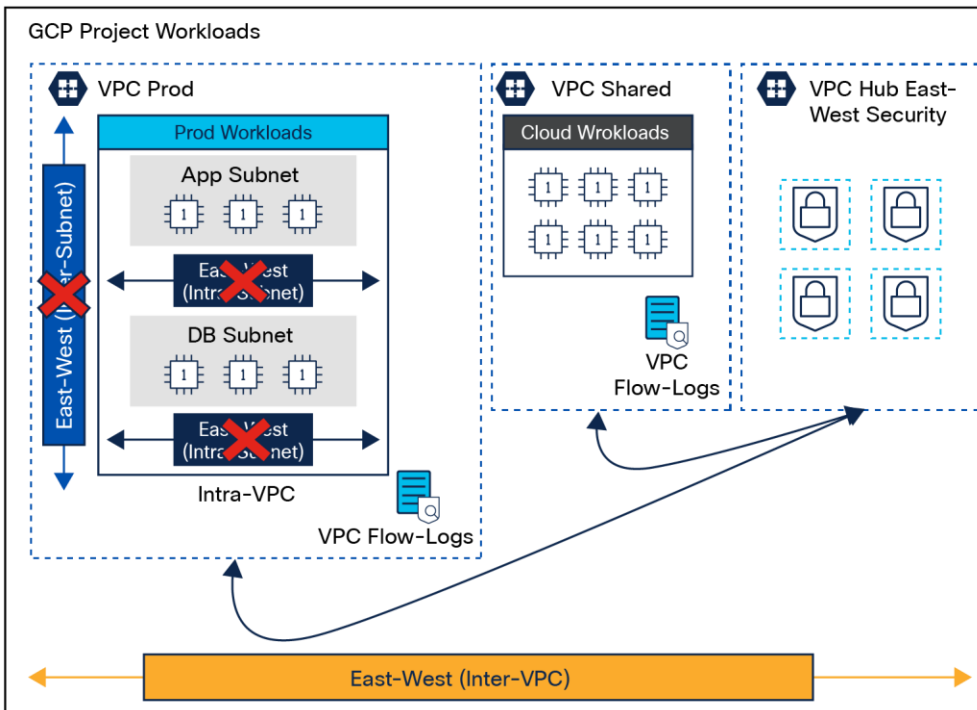
Azure Hub VNet East-West Insertion
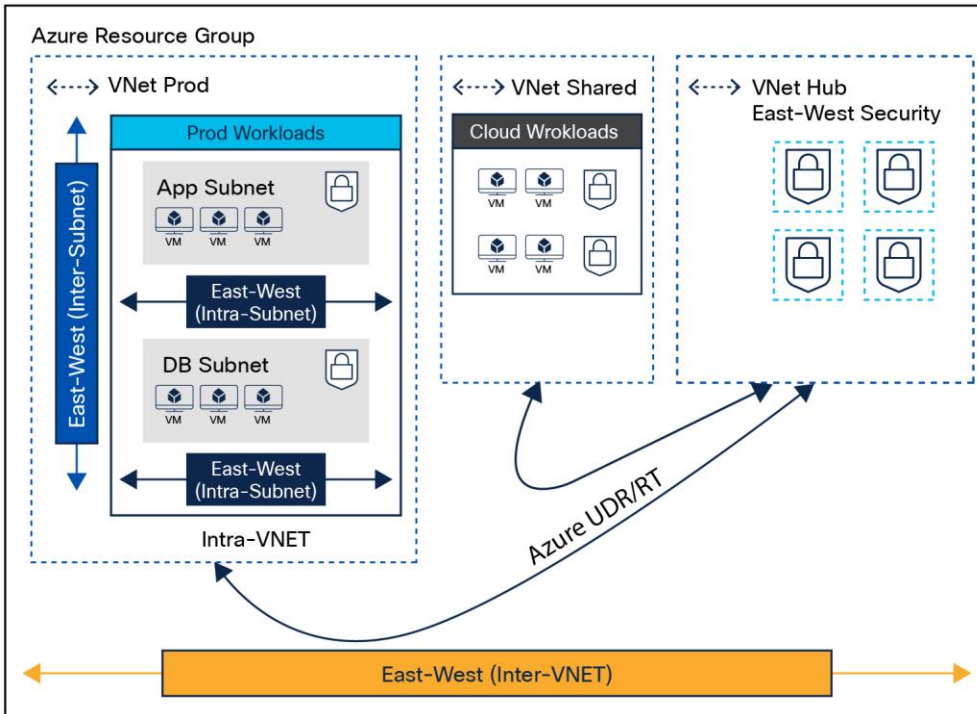
Figure 24.
Network microsegmentation for cloud agentless workloads with centralized/Hub VNet secure firewall deployment on Azure

- Acceptable for fine-grained segmentation
  - Azure User-Define Route (UDR)
- Full flow visibility with Network Security Group (NSG) flow logs and NSEL
  - Intra- and inter-subnet flows
- Protection at the network level
  - Intra-VNet
    - Intra-Subnet (App-App)
    - Inter-subnet (App-App)
  - Inter-VNet
    - Inter-subnet (App-App and External-App)
- FMC policy dual management
  - East-West (Secure Workload + FMC)
  - North-South – Ingress/Egress (FMC)
- Suitable for network/firewall engineers

Google Cloud Platform (GCP) Hub VPC East-West Insertion
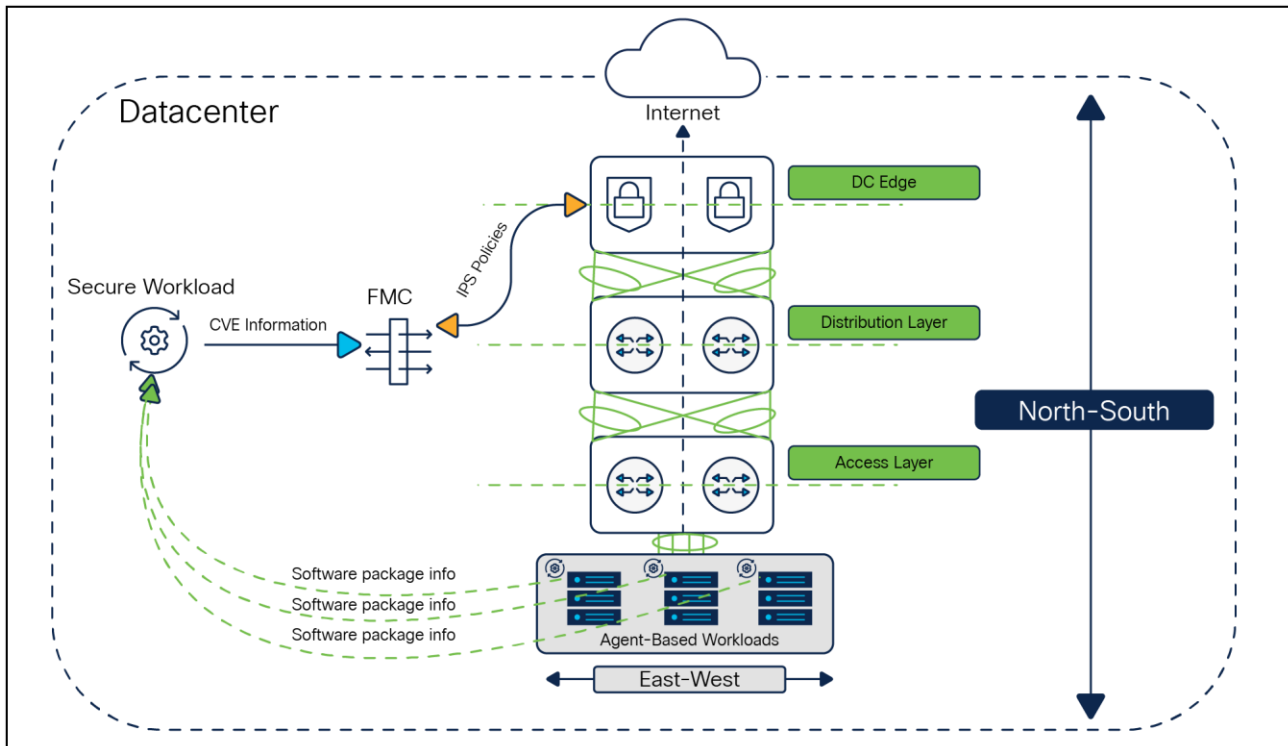

Figure 25.
Network microsegmentation for cloud agentless workloads with centralized/Hub VPC secure firewall deployment on GCP

- Reasonable segmentation
- Full flow visibility with VPC flow logs and NSEL
  - Intra- and inter-subnet flows
- Protection at the network level
  - Inter-VPC
    - Inter-subnet (App-App and External-App)
- FMC policy dual management
  - East-West (Secure Workload + FMC)
  - North-South – Ingress/Egress (FMC)
- Suitable for network/firewall engineers

Cisco Secure Workload and Cisco Secure Firewall – Virtual patch use case

Figure 26.
Secure workload and secure firewall virtual patch high-level architecture

Cisco Secure Workload delivers in-depth visibility of agent-based application workload runtime. Runtime information retrieved by the agent includes:

- Processes running, process snapshots, process tree, and process hash
- Software packages
- Software and kernel package vulnerabilities

Secure Workload can export the vulnerability information from workloads to FMC via the FMC connector. The FMC administrator can then run Cisco Recommended Rules to get fine-tuned Intrusion Prevention System (IPS) policies for applying a virtual patch in cases where there are important vulnerabilities in the environment that cannot be patched right away.
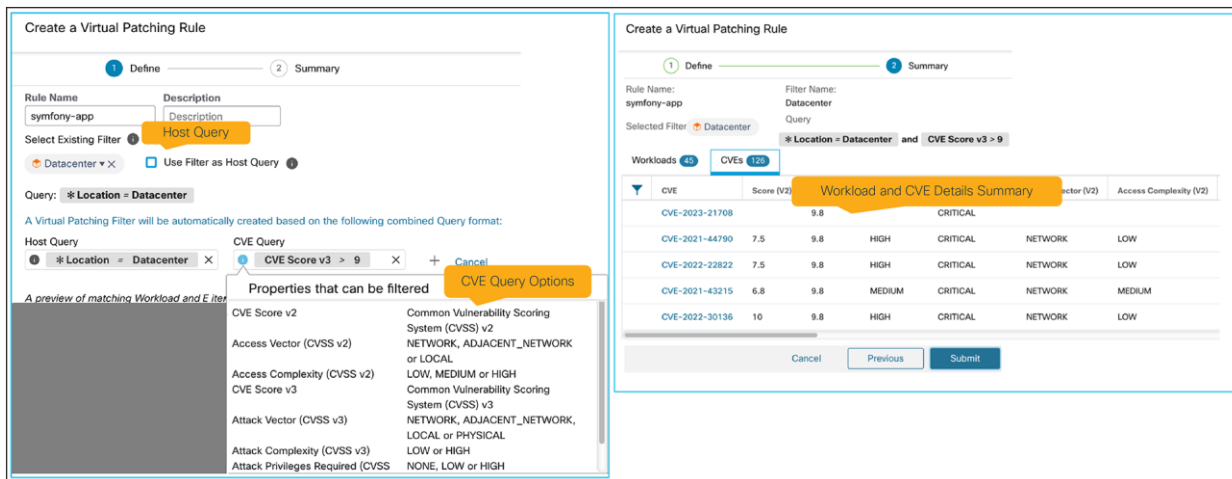
Vulnerabilities export

FMC connector manages both use cases (microsegmentation and virtual patching). The virtual patch use case has its own tab where all relevant virtual patch rules are created.
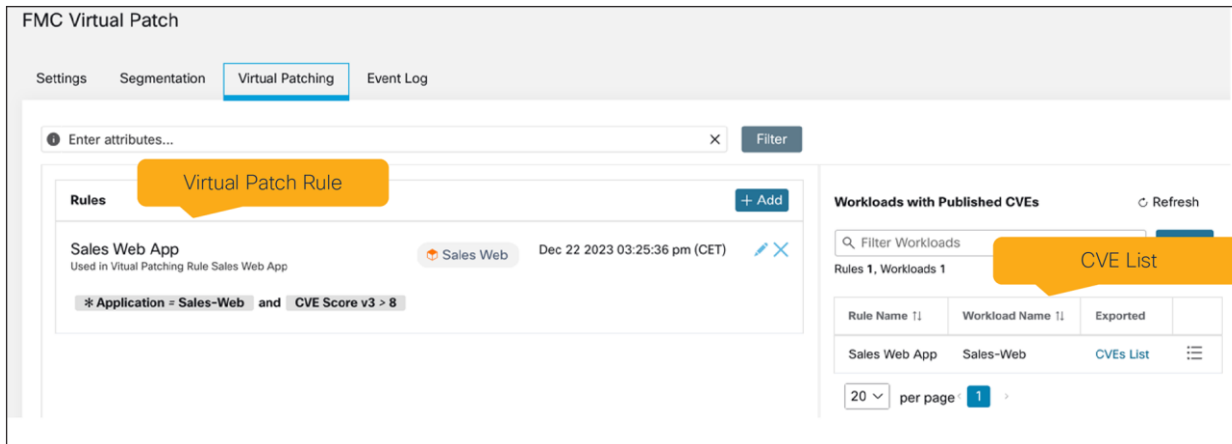
Figure 27.
FMC Connector virtual patch use case

To export vulnerability information from agent-based workloads, a virtual patch rule is required. The virtual patch rule consists of two elements:

- **Host query:** The host selector component to specify which workloads to export vulnerabilities from.

- **CVE query:** Query to specify the CVE selector for the workloads.

Figure 28.
Virtual patch rule definition

FMC Vulnerability import/visibility

Secure Workload will export vulnerability information to the Third-Party Vulnerabilities tab on FMC. This information is useful for FMC admins and SecOps operators, so they have visibility into the current hygiene of workloads.

For CVEs that have a Snort signature or signatures (SnortID) available, the CVE-to-SnortID mapping will be shown.

Figure 29.
Exported Vulnerabilities from Secure Workload to FMC

Cisco recommendations for fine-Tuned IPS policies

With the CVE intelligence exported from Secure Workload to FMC, automatic discovery of Snort signatures to mitigate the applicable CVEs on the workloads, can now be applied.

Cisco Recommended Rules must be run (previously known as Firepower Recommendations).

There are two main approaches to discover Snort signatures that are applicable to CVEs:

- **No rules active:** By selecting "no rules active," no preconfigured Snort signatures will be enabled on the IPS policy. Use this option if a fine-tuned IPS policy is required and only the Snort signatures that have a CVE mapped to them are required.

- **Selected base policy:** By selecting this option, a preconfigured "base policy" such as the Cisco default ones (e.g. Connectivity Over Security, Balanced Security and Connectivity, Maximum Detection) or a custom defined one. This option is useful when fine-tuned Snort signatures are not required but are preferrable to have coverage for both base policy Snort rules as well as those identified by the Cisco Recommended Rules.
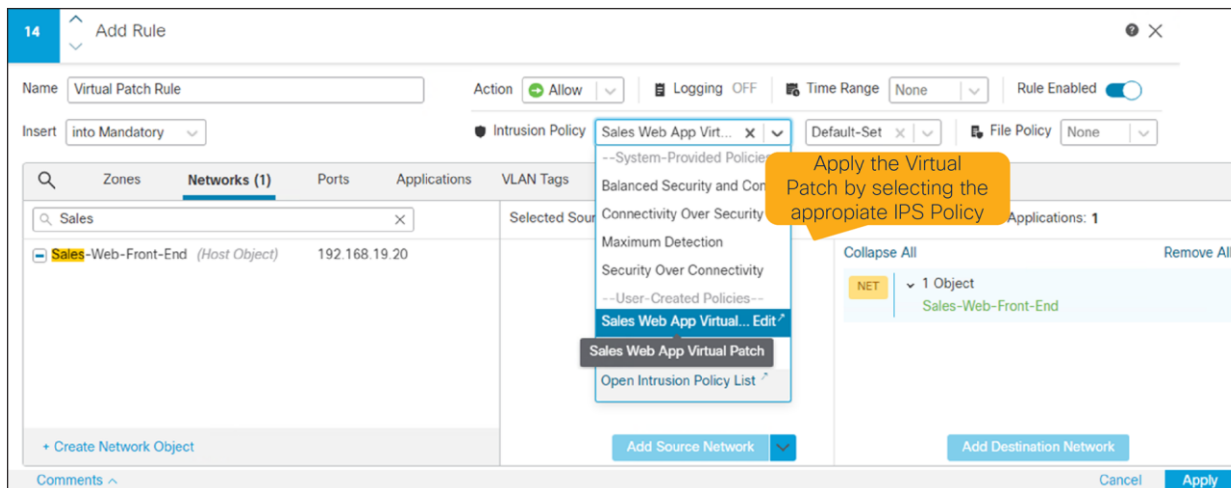

Figure 30.
Discovering snort signatures with Cisco recommended rules

Apply virtual patch

To apply the compensating control (virtual patch) for vulnerable workload traffic flows, it is required to modify or create an Access Control Rule in the Access Control Policies and to add the Intrusion Policy capability by selecting the virtual patch IPS policy.
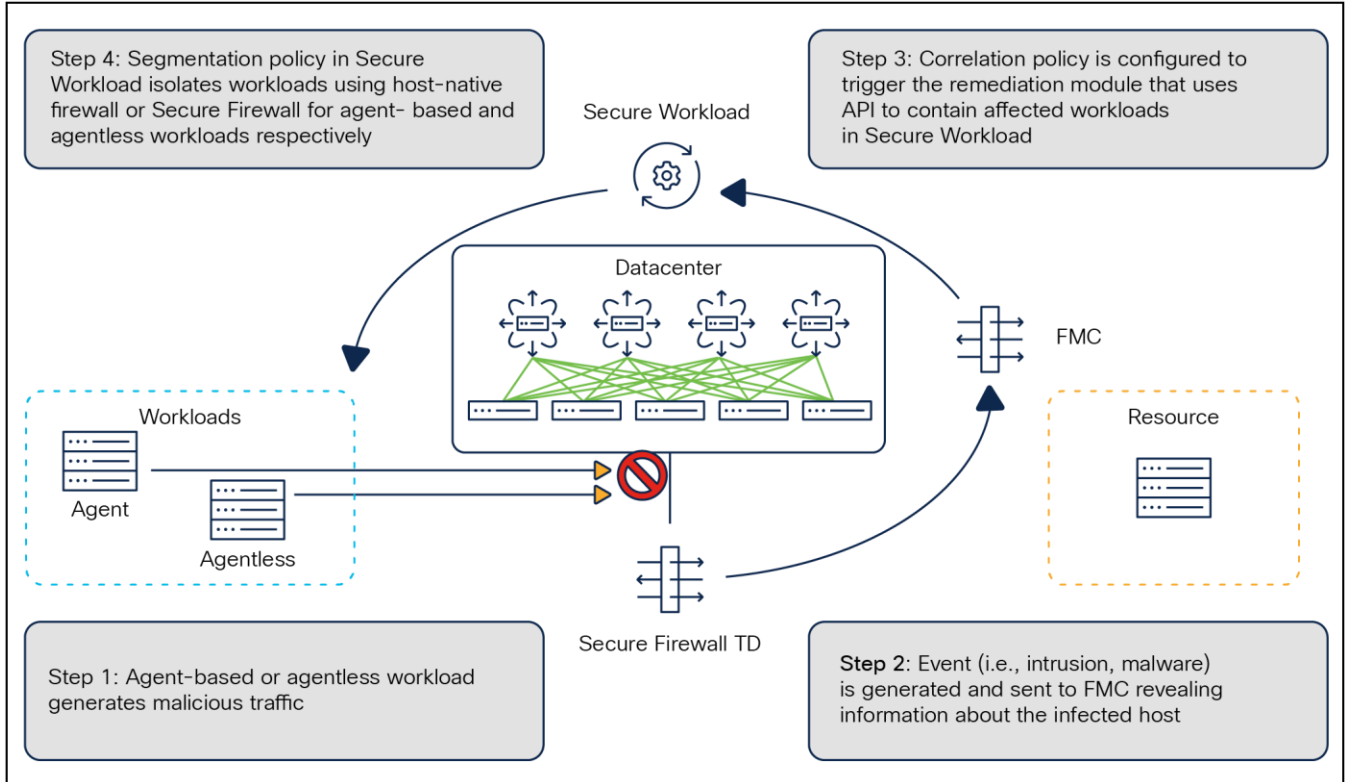
Figure 31.
Applying virtual patch to access control rule

Cisco Secure Workload and Cisco Secure Firewall – Rapid threat containment use case

The Rapid Threat Containment use case between Secure Workload and Secure Firewall enables network security and network operators to quickly identify and quarantine compromised workloads due to a detected anomalous behavior such as a malware event, intrusion event, or a correlation event.

This process consists of four steps:

- **Anomalous workload behavior:** An agent or agentless workload changes its behavior and generates anomalous or malicious traffic.

- **Secure firewall detection** Secure Firewall will detect the change in behavior and block the traffic flow if configured. An event is then sent to FMC with the workload's details.

- **FMC orchestration:** A preconfigured FMC correlation policy will track the change in behavior conditions and will orchestrate the response via API to Secure Workload.

- **Secure workload policy:** A predefined policy containing the quarantine attribute/label, which updates automatically via FMC, will propagate the policy through different enforcement points such as an agent within host-based firewalls or agentless with Secure Firewall to contain lateral movement and any malicious activity propagation.

Figure 32.
Rapid threat containment with secure workload and secure firewall

FMC Remediation module for secure workload

FMC Remediation Module for Secure Workload automates responses based on anomalous or malicious behaviors observed in network traffic flows or endpoints.

The package needs to be downloaded and uploaded to FMC, and requires some minimal configuration such as the Secure Workload Cluster IP, Root Scope, and the API Key for FMC. The only permission that is required is for "**User Data Upload**" in Secure Workload.

Figure 33.
FMC Remediation Module for Secure Workload

Correlation rules definition

Correlation Rules are a set of define events or signals that FMC must track in order to be triggered. Complex conditions can be built by leveraging the "AND" / "OR" operators for the rules.

Figure 34.
Correlation rules definitions

Correlation policy rules and response

After defining the correlation rules, the rules are grouped together into a Correlation Policy where each rule is assigned a priority and a response.
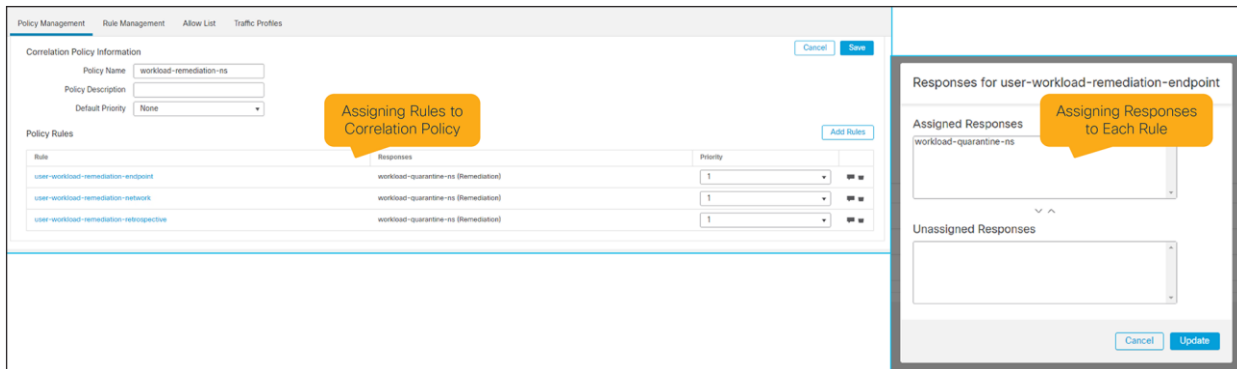
Figure 35.
Correlation policy and assigned responses
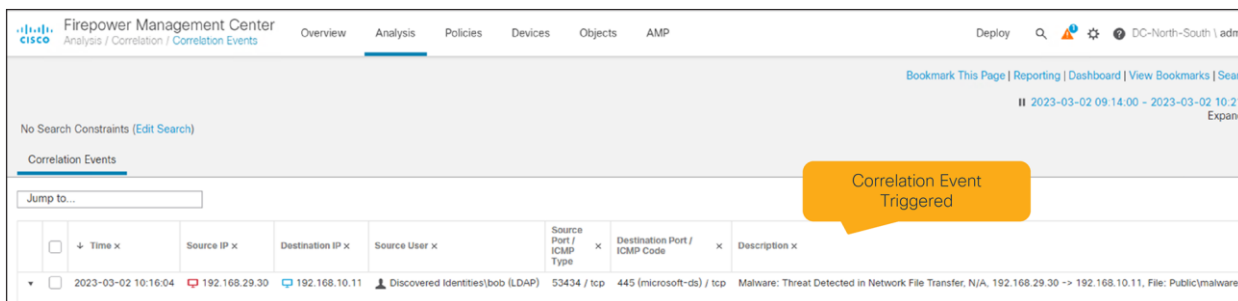
Remediation module and correlation policies events workflow

If a Correlation Rule condition is met, this will trigger the Correlation Policy

- **Correlation rule condition:** If a condition is met, this will trigger the Correlation Policy.

- **Correlation policy:** The Correlation Policy will initiate the response.

- **Remediation module:** The Remediation Module is instrumental to orchestrate the response to external systems, in this case Secure Workload. It will automatically send the affected workload or endpoint IPs to Secure Workload via API for later consumption in the segmentation policies.

Figure 36.
Correlation event triggered
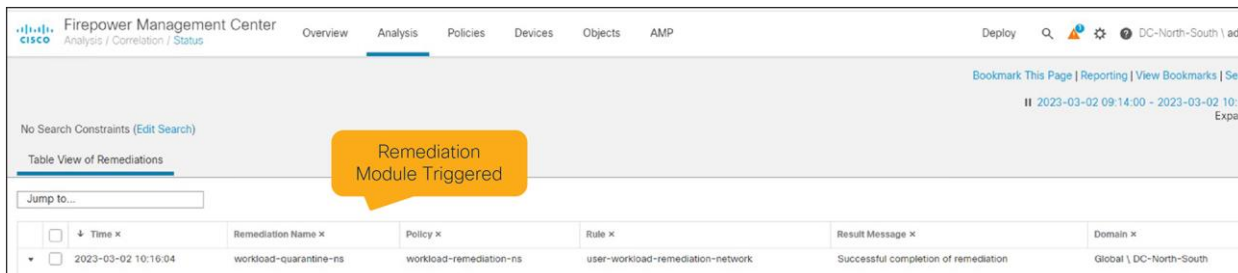
Figure 37.
Remediation module response triggered

Secure workload guardrail policy

Secure Workload uses human intent-based policy to define policy guardrails. These can easily be crafted with labels, which are used for context to automatically discover workloads and reduce the attack surface.

This integration can be used for the following use cases:

- **Quarantine workloads:** Block access to agent or agentless workloads with anomalous/malicious behaviors.

- **Deny access to compromised users/Endpoints:** Block access to compromised users or endpoints to applications on-premises or in multicloud environments.
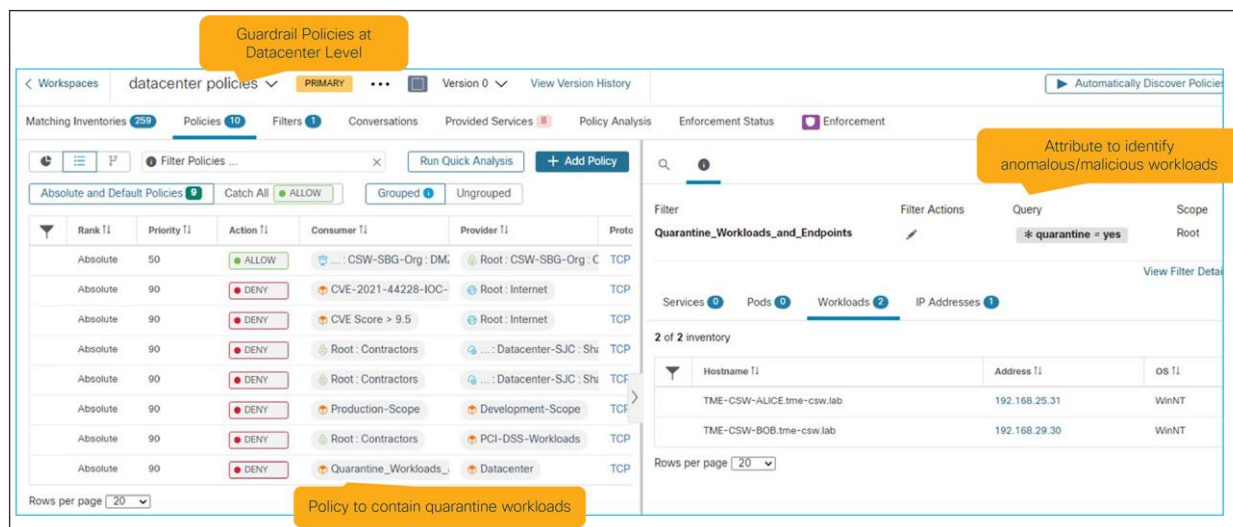
Figure 38.
Guardrail policies on secure workload

FAQs

- What happens if I have a mix of agent and agentless workloads behind the firewall mapped to the FMC Connector?

  ◦ The use case for Secure Workload and Secure Firewall integration is to protect agentless workloads. However, if there is a mix of agent and agentless workloads behind the firewall the solution will still work. The main difference is that rules from agent-based workloads (which are enforcing policies using the native host OS firewall) will be pushed to the FMC Access Control Policy (ACP) as well. This is a byproduct of Secure Workload's hierarchical policy model.

- Can I map more than one Scope to an ACP?

  ◦ No. Only one Scope can be mapped to one ACP.

- I want to enable Layer 7 capabilities and other FMC functionalities to the Secure Workload controlled rules. Can I do that?

  ◦ No. Secure Workload rules are orchestrated from Secure Workload and no modification should be done to them. While it is possible to add FMC functionalities, this is not advisable and not supported. At the time of this writing the use case is to provide east-west microsegmentation with L3/L4 policies.

- How can I have dual management of policies (Secure Workload owned and FMC owned)?

  ◦ Secure Workload has the capability to honor existing rules (merge) on FMC. Dual management is achieved by choosing to place Secure Workload rules either on top or bottom of existing ones. After

this is done, the rule ordering must be kept to preserve policy authoring from FMC and Secure Workload. If a rule is misplaced (out the intended order), Secure Workload will override the change and return the policies back to the original state.

- What happens if a rule owned by Secure Workload is modified?

  ◦ Secure Workload will override the change and return it to the original state.

- Which FMC versions are supported?

  ◦ To leverage dynamic objects any version above 7.0.1 is supported.

  ◦ To leverage Virtual Patch any version starting with 7.2 is supported.

  ◦ Current qualified releases by engineering includes 7.0.1 and 7.2.5