



WHITE PAPER

Micro-Segmentation and Security Orchestration for an Unassailable Cloud Defense

Contents

- The Cloud is Under Siege 3**
- Solution Considerations 3**
 - Visibility and Controls Configuration 3
 - Misconfiguration and Mismanagement of Security Controls 3
 - Planning for the Dynamic and Elastic Cloud 4
- Cloudvisory: The One-Stop Solution 5**
- Cloud-Native Visibility 6**
- Cloud-Native Control 8**
 - Limits to Legacy Micro-Segmentation Approaches 8
 - Intelligent Micro-Segmentation 8
 - White Listing 8
 - Organize and Orchestrate Security Policies 9
 - Auto-Discovery of Workload Context 9
- Policy Establishment and Management 10**
 - Point-and-Click Creation 10
 - Discovery of Policies 10
 - Automatic Orchestration 11
 - Rapid Policy Change Management 11
- Governance Through Monitoring 12**
 - Data Flow Violations 12
 - Automatic Handling of Violations 12
- Summary 13**

The Cloud is Under Siege

Many companies, from small businesses to multinational enterprises, modernize their computer environments with a shift to public and private clouds. Customer centricity, business agility, and reduction in costs are the main drivers for this infrastructure transformation. But cloud-based datacenters and architectures are under attack, just like conventional datacenters and infrastructures. Both environments are vulnerable to laterally moving threats.

Solution Considerations

Visibility and Controls Configuration

To limit the potential impact of threats, your organization needs to identify potential malware risks and quickly shut them down. This requires a reliable and consistent view into every aspect of your environment. For the cloud, this can be difficult, as 37% of respondents in one survey reported, “lack of visibility into infrastructure security,” as their “biggest cloud headache.”

Conventional network perimeter appliances and traditional network security strategies are less likely to be successful in cloud environments. Once beyond the firewall, there is little to stop attackers.

Every major cloud provider offers embedded controls that enable security decision-making before traffic reaches workloads. In contrast, legacy OS-based perimeter defenses rest inside the attack zone, which incurs a higher risk since security decisions are not made until data flow reaches the virtual machine. The granular, workload-level, white-listing approach of cloud-native controls must be explicitly configured before any data can flow in or out of a workload, instance or container. Configurations with least privilege access controls are critical to successful security.

Misconfiguration and Mismanagement of Security Controls

Ownership and operation of cloud-native security controls must be carefully considered. Initially, cloud provider DevOps teams might code security controls into their orchestration scripts, but scalability can become an issue:

- Security teams may have no visibility or understanding of what controls have been deployed, even with the cloud provider’s console.
- DevOps teams may use generic settings for complex controls which result in too much access and increased risk for the enterprise.
- Even though the cloud is an explicitly white-listed environment, DevOps teams can frequently misconfigure settings with too little access.
- Limited management and control over the security settings means DevOps needs to be continually involved in settings updates as workloads scale.
- Development and deployment slow because security teams have no easy way to adjust existing policies across multiple applications without complex coding and scripting.

“Advanced targeted threats bypass traditional perimeter and signature-based protection, creating a need for security controls to become automatable and adaptive.”

Source: Gartner

CLOUD NATIVE SECURITY CONTROLS:

“Gartner’s position on cloud security has been clear — public cloud services offered by the leading cloud providers are secure. Use IaaS provider’s native security controls.”

“Leveraging these controls allowed us to move from a secure network of workloads, to a network of secure workloads,”

— SVP of infrastructure.

“Through 2020, 80% of cloud breaches will be due to misconfiguration and mismanagement, not provider vulnerabilities.”

Source: Gartner

Planning for the Dynamic and Elastic Cloud

Security policy management can become complex, inaccurate and inefficient in any environment involving a cloud because workloads move, change, and scale often.

DevOps teams can attempt to learn the controls of different providers and write scripts to respond to changing requirements but this can result in deployment delays and create internal risk and audit concerns.

Strong security design requires a plan that includes:

- **Enhanced Visibility:** Visualization of entire infrastructure and its security controls to rapidly identify environmental risks and confirm correct policy deployment.
- **Security Orchestration and Automation:** Automatic provisioning and deprovisioning of specific controls to reduce misconfiguration issues and speed up operations.
- **Micro-Segmentation:** Automatic way to deliver exacting and granular policies for virtual machines and containers, based on workload function.
- **Monitoring and Governance:** Mechanism to track security status and identify potential risks and threats.
- **Buy Versus Build:** Reliable, commercially available solution to meet requirements.

The Details of a Modern Cloud Security Approach

In Gartner's June 2016 paper, "How to Make Cloud IaaS Workloads More Secure Than Your Own Data Center," they suggest:

1. Use IaaS provider's native security controls along with automation and DevOps practices.
2. Micro-segment by default: Move from a "secure network of workloads" to "network of secure workloads."
3. High levels of automation significantly reduce misconfiguration and mismanagement, thereby reducing attack-surface — greatly improving security.
4. Log everything and require persistent visibility — you cannot protect what you cannot see.

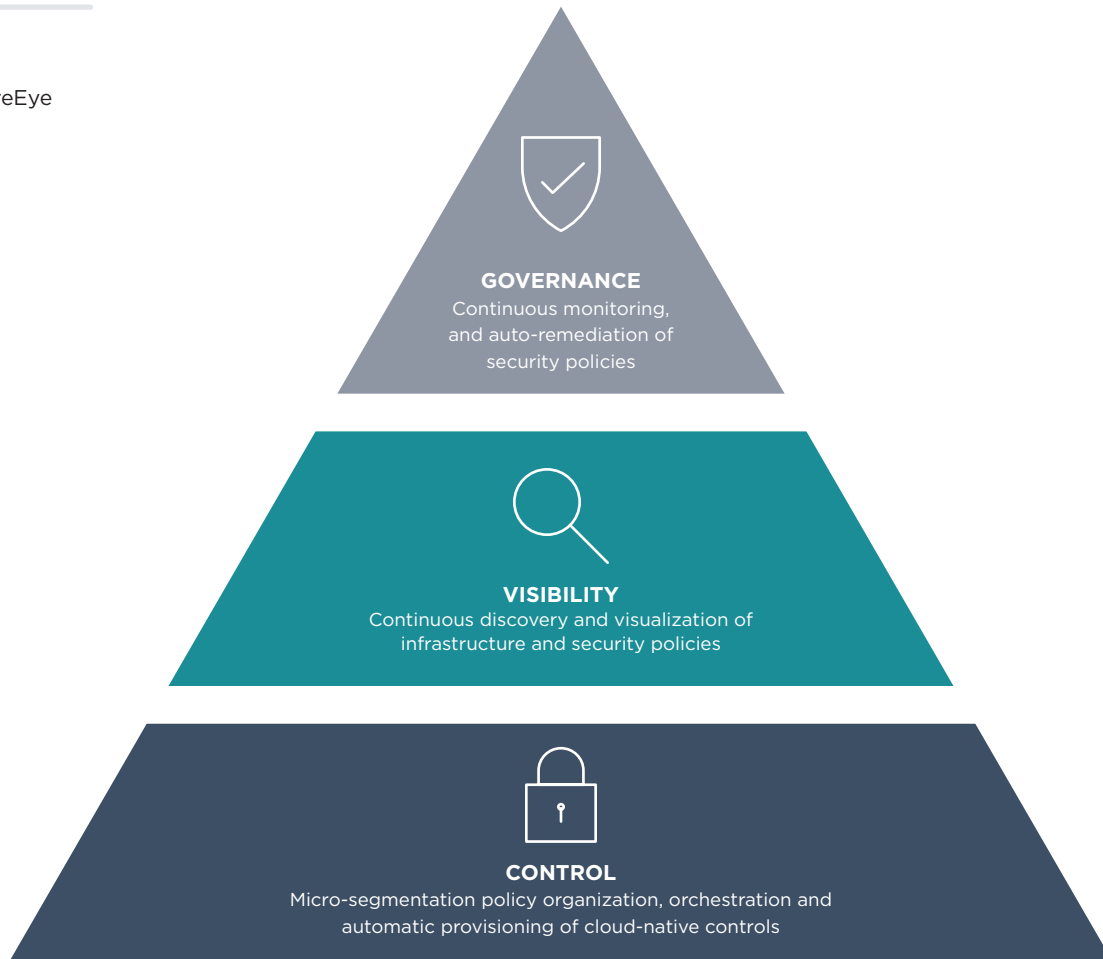
Properly leveraging cloud-native controls along with micro-segmentation, results in workloads that are better protected than in traditional data centers!

Cloudvisory: The One-Stop Solution

Powerful, centralized cloud security management and orchestration that makes use of the cloud-native provider controls is available through FireEye Cloudvisory.

Figure 1.

Elements of FireEye Cloudvisory.



Cloudvisory works across AWS, Azure, GCP, Kubernetes, OpenStack, and bare metal environments, enabling your organization to accelerate business, adapt to dynamic changes and reduce the risk of a security breach.

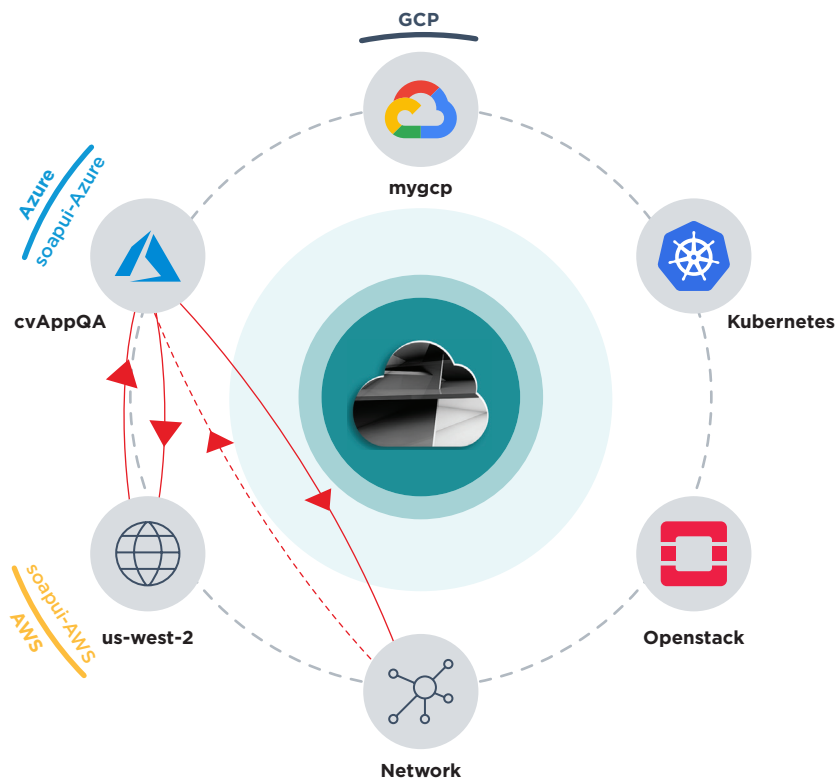
Cloud-Native Visibility

To manage, provision and remediate workload security, Cloudvisory continuously discovers and visually maps the cloud-native infrastructure and security objects of each cloud provider. For example, in AWS this includes Accounts, Regions, VPCs, Workloads, Security Groups and the intra workload network data flows. In OpenStack, this includes Accounts, Regions, Projects, Workloads, security groups, and data flows.

Data flows are mapped in real time, compared to deployed security controls and clearly differentiated as compliant or non-compliant to see the health of the underlying application environment.

Figure 2.

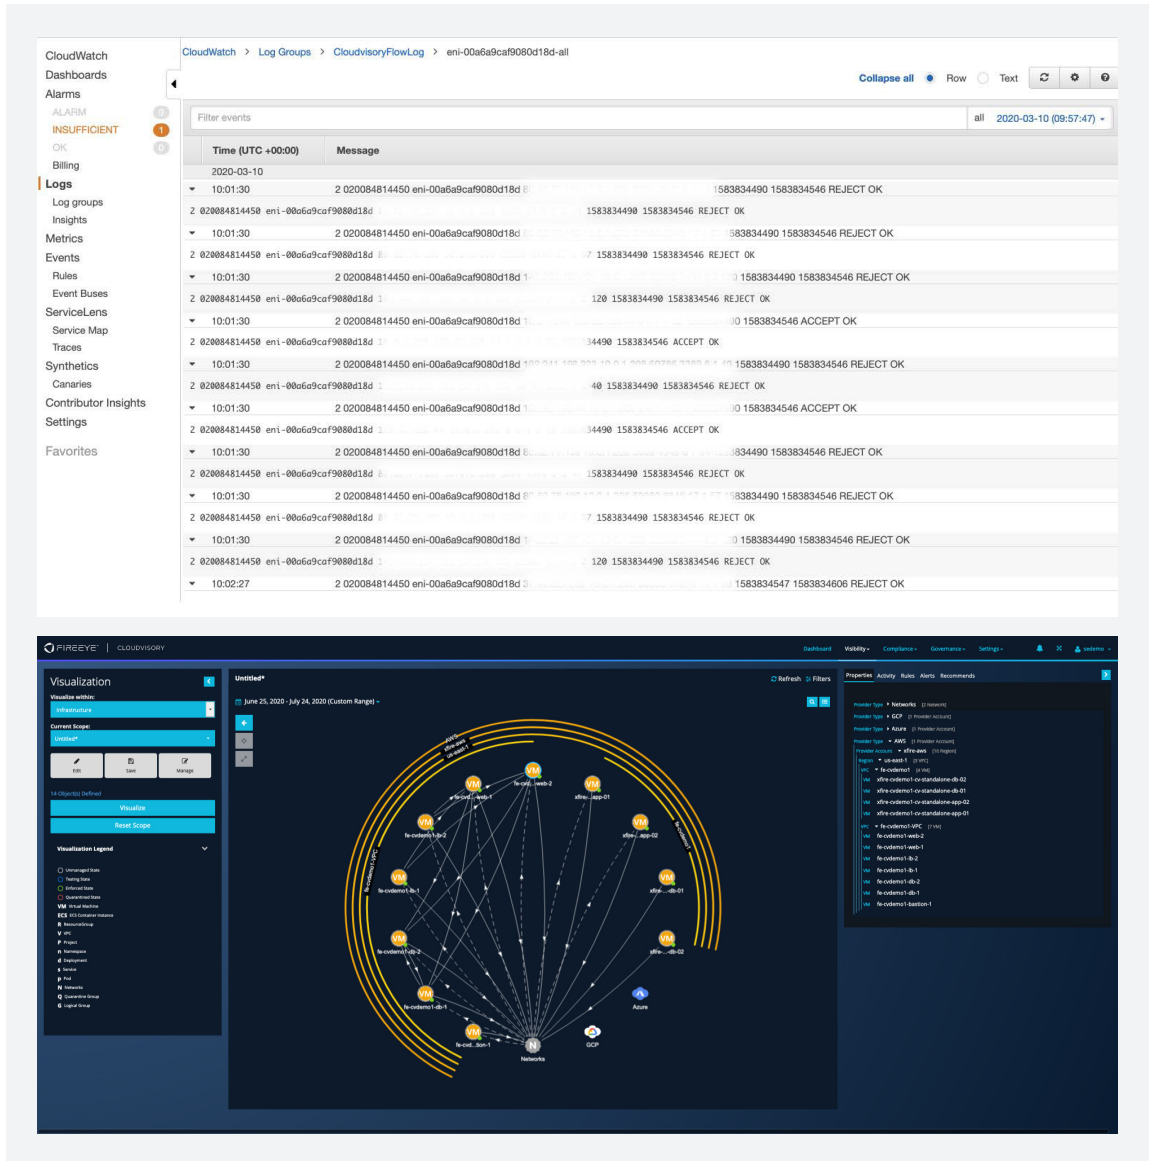
Cloudvisory visualization of a hybrid environment that includes AWS, Azure, GCP, OpenStack, Kubernetes and a traditional datacenter.



Cloudvisory can update the visual maps as assets move between environments (development to production) or between providers (datacenter to AWS), or if the environment itself changes. The DevOps teams can validate policy in a test mode, which does not block data flows.

Figure 3.

A typical list of VPC flows in AWS (above) can be visualized much more clearly in Cloudvisory (below). Non-compliant behavior, such as attacks, can be easily seen, alerted and quarantined.



Cloud-Native Control

To create, organize and manage security policies for environments with one or more clouds, The ideal solution:

- Delivers granular, intelligent micro-segmentation
- Simplifies policy creation using cloud-native security controls
- Uncovers misconfigured policy settings and helps correct them
- Organizes policy controls for consistent, repeatable and immutable security in dynamic environments
- Automates precise policy provisioning and de-provisioning across providers

Limits to Legacy Micro-Segmentation Approaches

Traditional micro-segmentation solutions tend to be very invasive and inflexible, and do not support native cloud security controls. Their reliance on operating system (OS) or inline firewalls positions the security governance point inside the attack zone where malware can compromise both workloads and security controls. Inline firewalls increase cloud configuration complexity and scalability issues.

A lack of native controls forces customers to manually configure all cloud provider security controls. These security vendors do not monitor cloud-native governance points, increasing organizational risk. Rogue or accidental changes can expose environments to hackers and interrupt active applications.

Intelligent Micro-Segmentation

It is entirely possible to move to a micro-segmentation of workloads, micro-services and containers. There are several reasons to aspire to and implement such a network of secure workloads:

- Current DevOps processes often provision incorrect or far too broad, security policies to workloads.
- Too much access greatly increases risk of malware and nation state actors.
- Malware that can easily move laterally in the environment will eventually find rich data targets.
- Rogue actors can be frustrated and blocked by micro-segmentation policies.

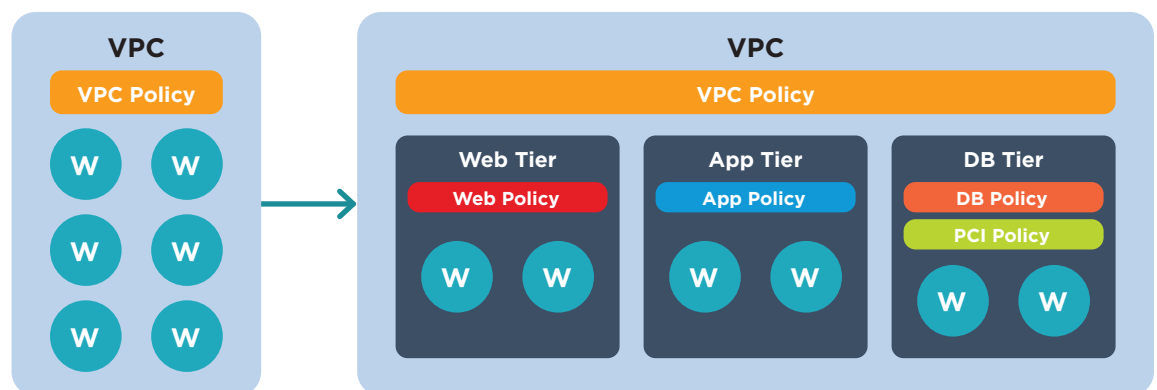
Managed correctly, micro-segmentation can harden security and end lateral hacker threats by fencing in smaller, discrete groups of workloads.

White Listing

In the cloud, workloads cannot receive any communications until cloud-native controls are configured. Whitelisting policies enable specific, granular ingress and egress as well as port and protocol rules that effectively provide a workload-level firewall (Fig. 4). This type of workload security frustrates hackers and halts the migration of malware within the environment. Therefore, Cloudvisory enables whitelisting policies by automating intelligence micro-segmentation through policy management and organizational infrastructure.

Figure 4.

Before and after whitelisting.



Organize and Orchestrate Security Policies

Power and scale for secure cloud operations relies heavily on organization and orchestration of micro-segmented policies.

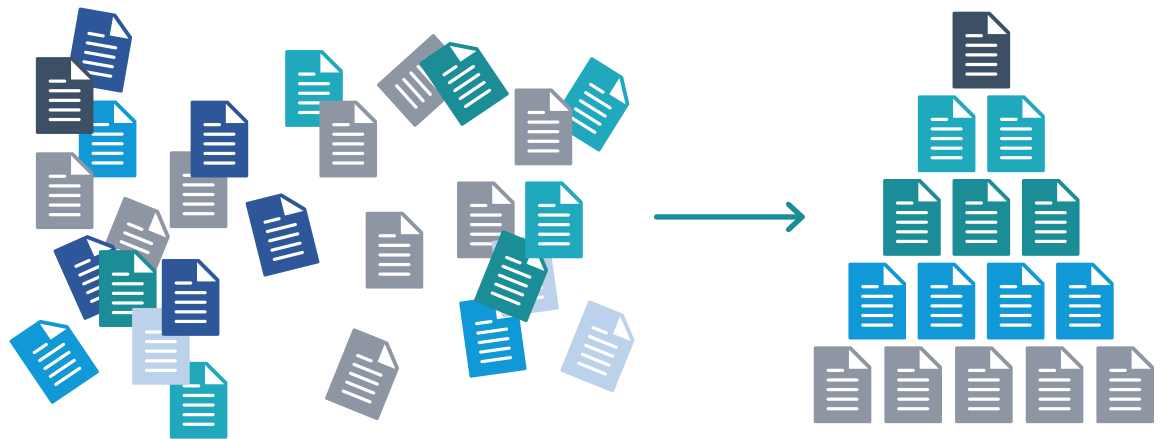
Cloudvisory can automatically respond to dynamic changes in the environment, adjusting cloud-native

security policies to keep workloads protected. To nimbly organize and manage policies, Cloudvisory enables:

- Auto-discovery and grouping of workloads based on context
- Automated policy creation, loosely coupled to workloads through logical groups

Figure 5.

Visualization of auto-discovery and grouping of workloads based on their context.



Auto-Discovery of Workload Context

Policy definitions are ultimately associated to workloads based on workload context, which is determined by logical group memberships and includes variables such as:

- Cloud provider
- Infrastructure membership
 - Account
 - Region
 - Resource group, virtual private cloud, project
- Application
 - Application tier
- Governance requirements such as
 - CIS, GDPR, HIPAA, NIST, PCI, OpenStack Security Checklist, and more
- Virtually any logical or ad hoc grouping required to manage policy

A continuous discovery capability allows Cloudvisory to discover a workload's context and automatically group workloads with common contexts.

Policy Establishment and Management

Cloudvisory enables two ways to create policies. Both methods are simpler and more accurate than writing code using commercially available orchestration tools.

Point-and-Click Creation

Cloudvisory users are not required to be experts in any particular cloud provider’s controls. By clicking through a sequence of guided screens, DevOps teams can create portable general policy rules.

Discovery of Policies

Developers are often unsure of the exact rules needed to control an application. As a result, too much access may be deployed, increasing risk to the environment. Cloudvisory is one of the fastest ways to establish least privilege policy control. It can discover the exact flows needed to run an application and use that information to create reusable and portable policy rules.

Figure 6. Cloudvisory Discovery Policy Creation Screen.

<input type="checkbox"/>	Source	Destination	Duration	Service
<input checked="" type="checkbox"/>	sangi_vm2	sangi_vm5	Unlimited	<input checked="" type="checkbox"/> UDP 67 UDP 67
<input checked="" type="checkbox"/>	sangi_vm2	sangi_vm3	Unlimited	<input checked="" type="checkbox"/> TCP 443 TCP 443
<input checked="" type="checkbox"/>	sangi_vm2	sangi_vm4	Unlimited	<input checked="" type="checkbox"/> TCP 3306 TCP 3306

Regardless of how policies are established, Cloudvisory translates policy definitions into the cloud-native controls for any cloud provider. Consider a PCI policy definition that determines which virtual instances can communicate with the PCI infrastructure back in the data center. Cloudvisory can translate this definition into the various providers controls, such as AWS Security Groups, Azure Network Security Groups and OpenStack Security Groups. Dynamic network parameters such as IP addresses of servers or moving an application to different servers are automatically managed by Cloudvisory. These capabilities free DevOps teams from having to become experts in each provider’s control set and reduce coding requirements.

Automatic Orchestration

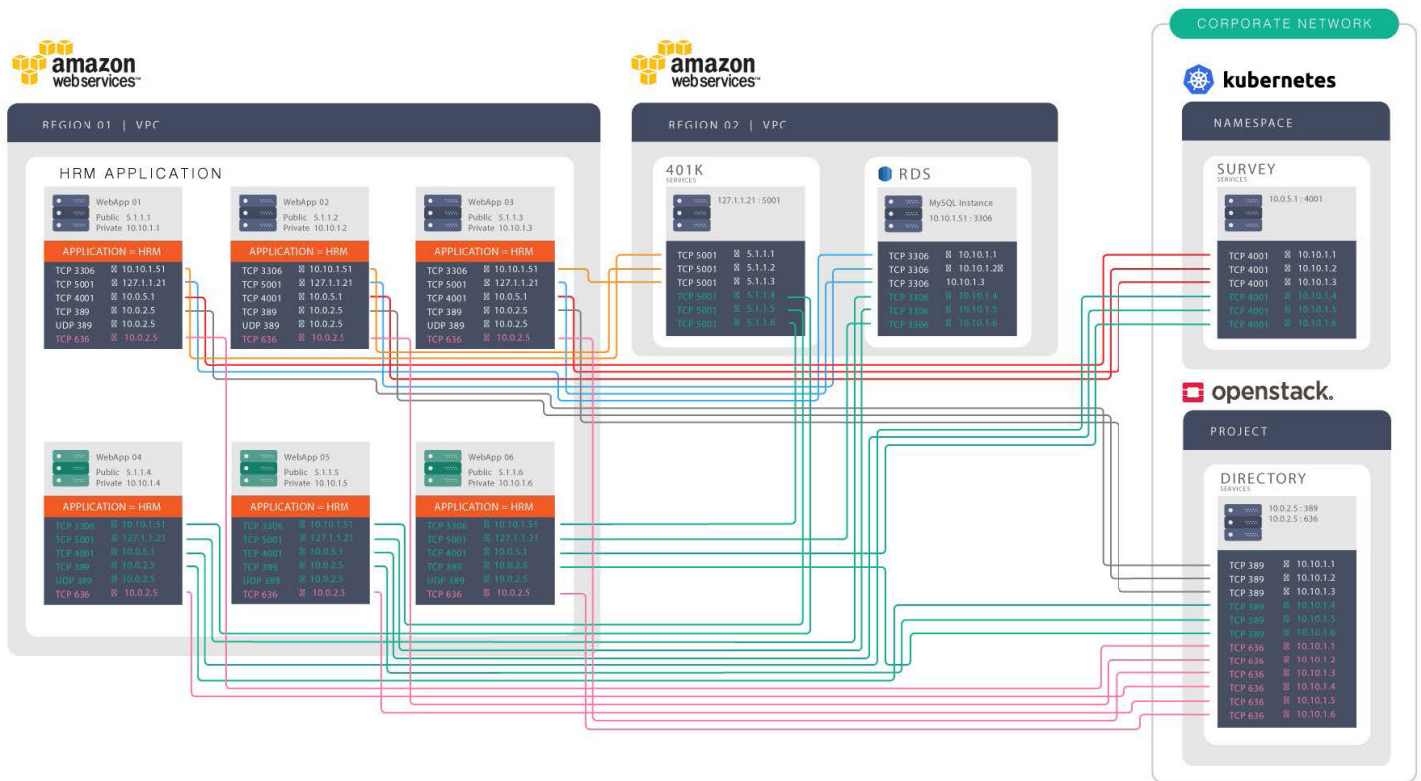
After Cloudvisory has established policy rulesets, they are loosely coupled to a workload’s context using an API or a point-and-click interface. Policies can then be automatically and correctly provisioned to workloads, even as new ones spin up, change roles, or move between environments, between providers or to other logical groups. In every case, Cloudvisory automatically identifies changes in context and updates the security policies in real time. No coding or complex scripting is required, which means greater security control and more precise policy management.

Rapid Policy Change Management

A complex hybrid cloud example (Fig. 7) that includes AWS, Azure, and OpenStack assets can help illustrate the power of Cloudvisory.

First, the HRM policy definition established by Cloudvisory for the application specifies the security ingress and egress rules for any workload containing the meta data ‘app=HRM’. As additional workloads spin up with this metadata, Cloudvisory calculates and provisions the correct cloud native controls for each service that is part of the integrated HRM application. Tracking all required connections, Cloudvisory ultimately provisions precise policies.

Figure 7. Complex hybrid cloud with assets in various HRM environments.



DevOps teams tasked to code or script these policies must keep track of immense complexity across all environments, public and private IP addresses, ports and protocols for ingress and egress. This takes considerable time and can lead to misconfigurations and non-functioning applications. Cloudvisory automates the entire process, removing the complexity associated with learning, organizing, provisioning, calculating, updating and managing cloud-native security policy. This not only takes less time and cuts costs, but also provides more consistent and accurate security.

Governance Through Monitoring

After policies are organized, established and provisioned, environments must be monitored for potential compromises and to verify immutable security.

Using the Cloudvisory UI or its APIs, business, operational and security teams can visualize data traffic at different levels of the infrastructure hierarchy. For example, the user can observe data traffic between virtual machines in the same hypervisor, different cloud providers, between cloud and traditional data center assets or in a single cloud account. Unlike traditional netflow collectors available in switches and routers, data flows are captured, stored and displayed with contextual information such as provider, infrastructure properties, application and user defined attributes. This context conveys information that helps understand application behavior, troubleshoot misconfigured policies, quickly analyze and triage security incidents and simplify data analytics processing.

After infrastructure discovery and policy establishment, Cloudvisory continuously monitors two critical areas:

- Data flows and byte counts of any managed workload
- Native policy governance points for each workload

Data Flow Violations

Data flows are compared against the actual deployed and compliant policies. Any data flows not matching allowed policies are immediately flagged, tagged as non-compliant, blocked and clearly indicated in the UI.

Alerts are configurable and an infected workload can be quarantined immediately to neutralize any potential threat.

For example, if a workload is not allowed to FTP and it is infected with malware that attempts to FTP out of the workload, Cloudvisory can detect the attempted data flow (and any corresponding increase in average byte count) and alert on this activity. The alert is displayed on the dashboard sent as email notifications to administrators

and integrated with SIEM solutions. The attempted communication is blocked. Cloudvisory can be configured to immediately quarantine the infected workload and suspend all current outbound security policies and block all network traffic to and from selected workloads. This prevents lateral movement of threats to other applications and resources and any further network infection until forensics teams can remediate the situation.

Alerts for non-compliant traffic due to misconfigured policies can cause application malfunctions. Administrators can quickly learn about any such policy rule gaps and resolve them using the point-and-click Cloudvisory interface.

Automatic Handling of Violations

Cloudvisory continuously checks the configuration of the established native security controls to make sure they remain in compliance. If an unauthorized change in the native security rules is detected, Cloudvisory generates alerts. Security policies can be configured to auto-rollback non-compliant changes. In this case, any unauthorized changes in native security controls are automatically reverted, bringing policy governance back to compliance. In such cases, it also generates an alert and audit log entry to register the non-compliant event.

For example, if an administrator using an Azure console accidentally removes security group rules for port 80 access, many cloud providers would not flag the action. However, the result would be a downed application, because it would be unable to contact the web interface of the application.

Cloudvisory would detect such accidental or malicious configuration changes and can immediately roll back the change to put the environment back into a compliant and working state. This reduces or eliminates downtime and automated triage and risk.

Summary

Every organization, of every size and level of complexity, is inexorably embracing public and private clouds. To ensure that their security stays abreast of their operational practices, they need to maintain visibility over all possible attack surfaces, configure security correctly even as the organization evolves and develop strong security designs that account for a dynamic and elastic cloud.

Elements of such plans include orchestration and automation for faster, error-free operation, micro-segmentation to implement workload-focused policies, monitoring and governance to quickly identify risks and threats, and selecting commercially available solutions to meet requirements.

FireEye Cloudvisory is a single solution that meets all critical multi-cloud security needs, from visibility and policy management to power and ease of use, for public, private and hybrid environments.

To learn more about FireEye, visit: www.FireEye.com

FireEye, Inc.

601 McCarthy Blvd. Milpitas, CA 95035
408.321.6300/877.FIREEYE (347.3393)
info@FireEye.com

©2020 FireEye, Inc. All rights reserved. FireEye is a registered trademark of FireEye, Inc. All other brands, products, or service names are or may be trademarks or service marks of their respective owners. CS-EXT-WP-US-EN-000312-01

About FireEye, Inc.

FireEye is the intelligence-led security company. Working as a seamless, scalable extension of customer security operations, FireEye offers a single platform that blends innovative security technologies, nation-state grade threat intelligence, and world-renowned Mandiant® consulting. With this approach, FireEye eliminates the complexity and burden of cyber security for organizations struggling to prepare for, prevent and respond to cyber attacks.

