

A SANS Product Overview

Securing Assets Using Micro-Segmentation: A SANS Review of Guardicore Centra™

Written by **Dave Shackleford**

June 2020

Sponsored by:

Guardicore

As security professionals, we are starting to rethink our approaches to network and workload security. Right now, our environments are becoming more complex, heterogeneous and interconnected with service providers, and the nature of application development is moving faster than ever before. We realize that we must address the following:

- **We must look at our entire environment as potentially untrusted or compromised, instead of thinking in terms of "outside-in" attack vectors.** Increasingly, the most damaging attack scenarios are almost entirely internal due to advanced malware and phishing exercises that compromise end users as a starting point for attacks.
- **We need a better understanding of application behavior at the endpoint and a better understanding of our application workflows.** Organizations need the capability to enforce policy at a more granular level and follow the workload, regardless of where that workload appears. This is almost always accomplished through host-based controls.
- **We must focus on trust relationships and system-to-system relationships within all parts of our environment.** At the same time, we also need a security methodology that can keep pace with a DevOps development and deployment model that brings efficiency, automation and speed to the enterprise.

©2020 SANS™ Institute

These are all worthwhile goals, but many of our traditional controls are not capable of accomplishing them. Compounding this is the advent of highly virtualized and converged workloads, as well as public cloud workloads that are highly dynamic in nature. Cloud workloads may change rapidly or exist only for very short time periods. Today, micro-segmentation technology can help prevent attackers from using unapproved connections to move laterally from a compromised application or system, regardless of environment. Essentially, micro-segmentation facilitates the creation of *affinity policies*, where systems have relationships and permitted applications and traffic, and any attempted communications are evaluated and compared against these policies to determine whether the actions should be permitted. This happens continuously. Effective micro-segmentation technology will also include some sort of analytics processing of attempted behaviors, adapting dynamically over time to changes in the workloads and application environments.

To implement a micro-segmentation model, security and operations teams need to focus on two key concepts. First, security must be integrated into the workloads themselves. By creating a layer of policy enforcement that travels with workloads wherever they go, organizations have a much stronger chance of protecting data, regardless of where the instance runs. This shifts security policy and access control back to the individual instances, as opposed to within the network itself, which is needed now due to modern cloud-oriented architecture (and rapid and dynamic deployments and workloads) that don't conform to traditional static network segmentation and controls. Traditional static network segmentation policies have also lacked in granularity—and more visibility into processes, identity and domain information is critical.

Second, the behavior of the applications and services running on each system needs to be much better understood, and the relationships among systems and applications need more intense scrutiny to facilitate a highly restricted, micro-segmented operations model without adversely impacting connectivity. Dynamic assets such as virtual instances (running on technology such as VMware internally or in AWS, Azure and others externally) and containers are difficult to position behind "fixed" network enforcement points. To resolve this problem, organizations can adopt a micro-segmentation strategy that allows traffic to flow only between approved systems and connections, regardless of their environment. Shifting security enforcement into the workloads themselves also increases speed and efficiency in rapidly changing environments.

SANS had the opportunity to review Guardicore's Centra™[1] platform, focusing on its micro-segmentation capabilities and unique differentiators in the zero trust market space. Because micro-segmentation is a broad topic, we started by focusing on the accessibility of the environment itself. The Guardicore team set up several test environments for the SANS team, with a variety of services, systems and applications running. In addition, there were pre-built attack scenarios that could help highlight key capabilities of the Guardicore platform.

[1] Centra™ is a trademark of Guardicore.

Product Overview

When we first accessed the platform, we explored the main dashboard to see what details were presented. The dashboard was laid out intelligently, with a breakdown of malicious behavior detection, unusual and failed network connections, events in the environment that Guardicore labeled as "incidents," and a range of other data (see Figure 1).

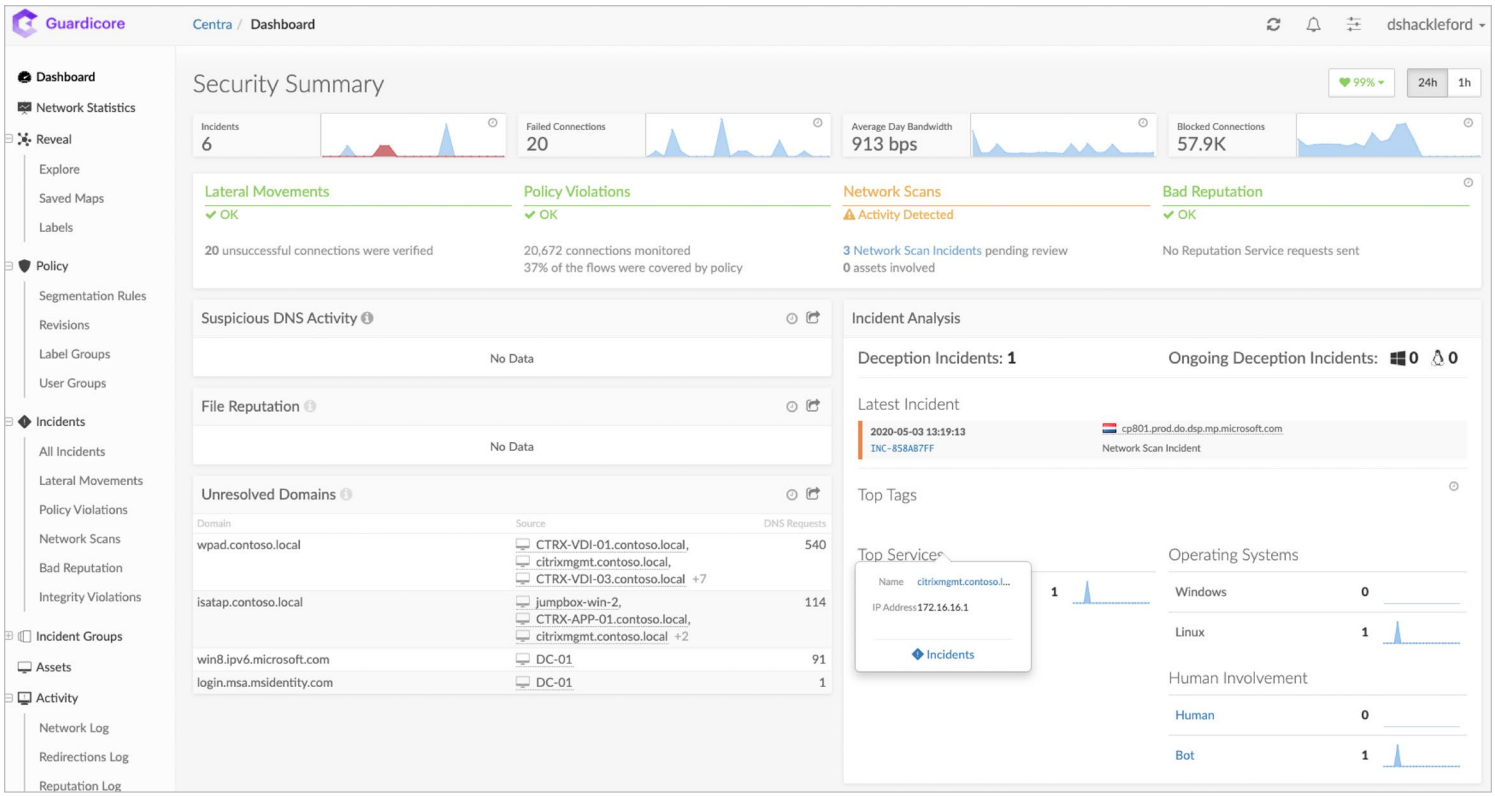


Figure 1. Centra Main Dashboard

We explored the console a bit more and found that the product excels in ease of use. The dashboard is easy to interpret, and the menu system on the left-hand side is broken down into categories such as:

- **Network Statistics—**Tracking and visualization of both internet traffic and east–west (internal connections)
- **Reveal—**Discovery and visualization detail on assets
- **Policy—**Segmentation rules and group definitions
- **Incidents—**Detected events and incident behaviors in the environment
- **Incident Groups—**Groups of related incidents and events
- **Assets—**A breakdown of all assets with associated risks and detected events
- **Activity—**Logs of network connections, web requests and more

The core Administration section is condensed and simple to navigate, too. In this section, security operations teams are able to configure detection policies, reputation analysis capabilities (covered later), Centra components and agents, integration with other solutions and threat intelligence, and users and permissions for the platform. Role-based access control of the platform is easy to set up, and multifactor authentication is supported.

Micro-Segmentation Focus and Asset Analysis

In addition to asset discovery, one of the most critical capabilities we look for in a micro-segmentation platform is support for a wide range of platforms and operating systems, because this can be a major limiting factor in rolling out a new access control model in many enterprises. The Guardicore platform supports numerous internal and cloud-based technologies, as we saw in the *Explore* section of the interface (representing the underlying asset discovery capability and reporting of the platform, shown in Figure 2) within the *Reveal* category. The product supports very flexible asset labeling, which makes customizable mappings and policy creation easy to adapt to individual use cases and environments. Organizations can enable dynamic labeling to designate asset criteria that Centra uses to automatically identify and properly label workloads over time, as well. This process removes the need to manually add, remove or update labels entirely.

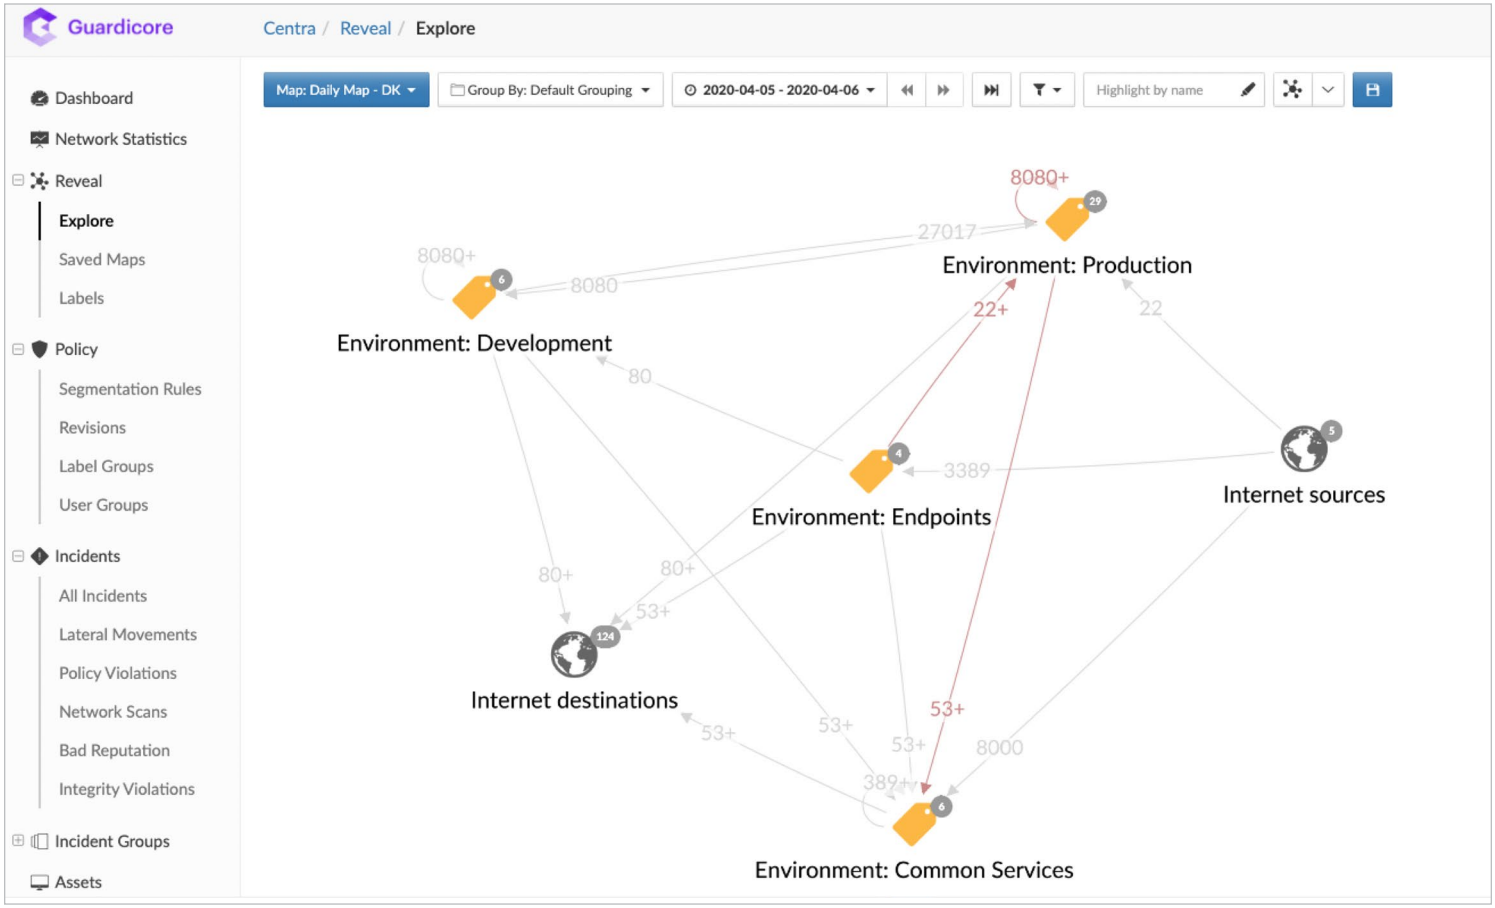


Figure 2. Guardicore Asset Exploration Dashboard

We quickly learned to modify the visualization of our resources in the test environment by changing the *Group by* menu in the top of this dashboard, selecting *Platform*, *App*, *Role*. This choice provided us with breakdowns of assets by location—with both real-time and historical perspectives (see Figure 3).

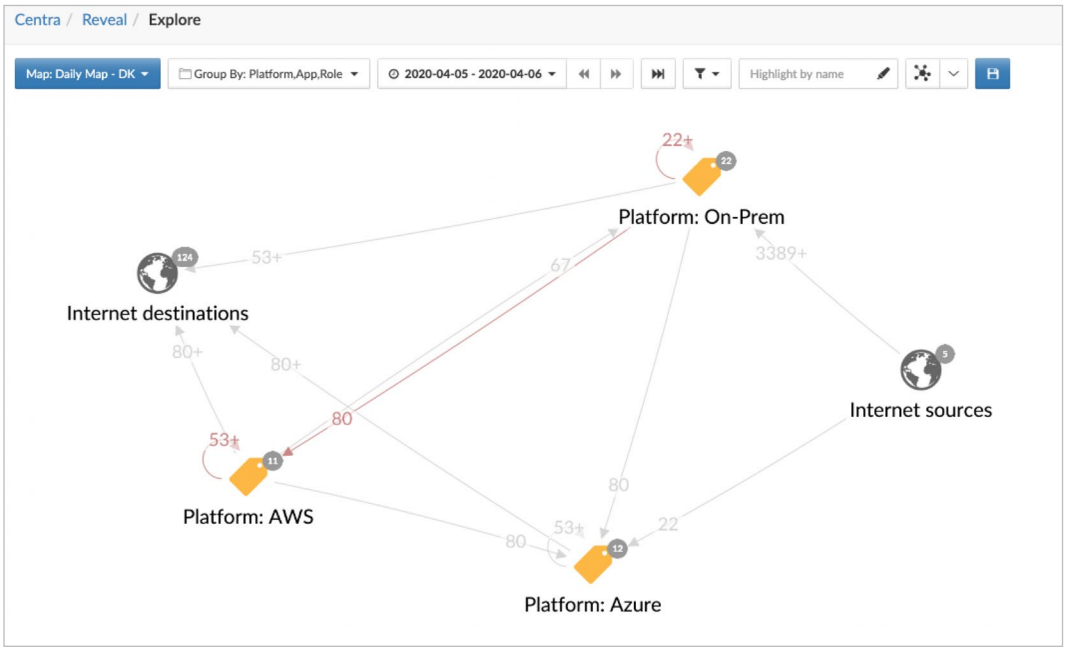


Figure 3. Asset Visualization by Environment

To meet the needs of a variety of organizations and use cases, it's imperative that a foundational access control platform be flexible in allowing customers to create a variety of views within their environments—and Guardicore's Reveal dashboard model meets this need. This is important when thinking through key micro-segmentation use cases; for example, security administrators and operations teams will likely need to see all activity in the environment, perhaps with emphasis on critical and sensitive assets. Incident responders will need to quickly select only certain assets or groups to drill into and uncover deeper asset and communication details, and auditors and compliance analysts may need to quickly view all resources of a specific type (PCI resources, for example). We dug into one of the represented groups by double-clicking the *Platform: AWS* group shown. The interface narrowed its focus to show only the applications running in the AWS environment. We right-clicked this icon and chose *Filter by this item* (see Figure 4).

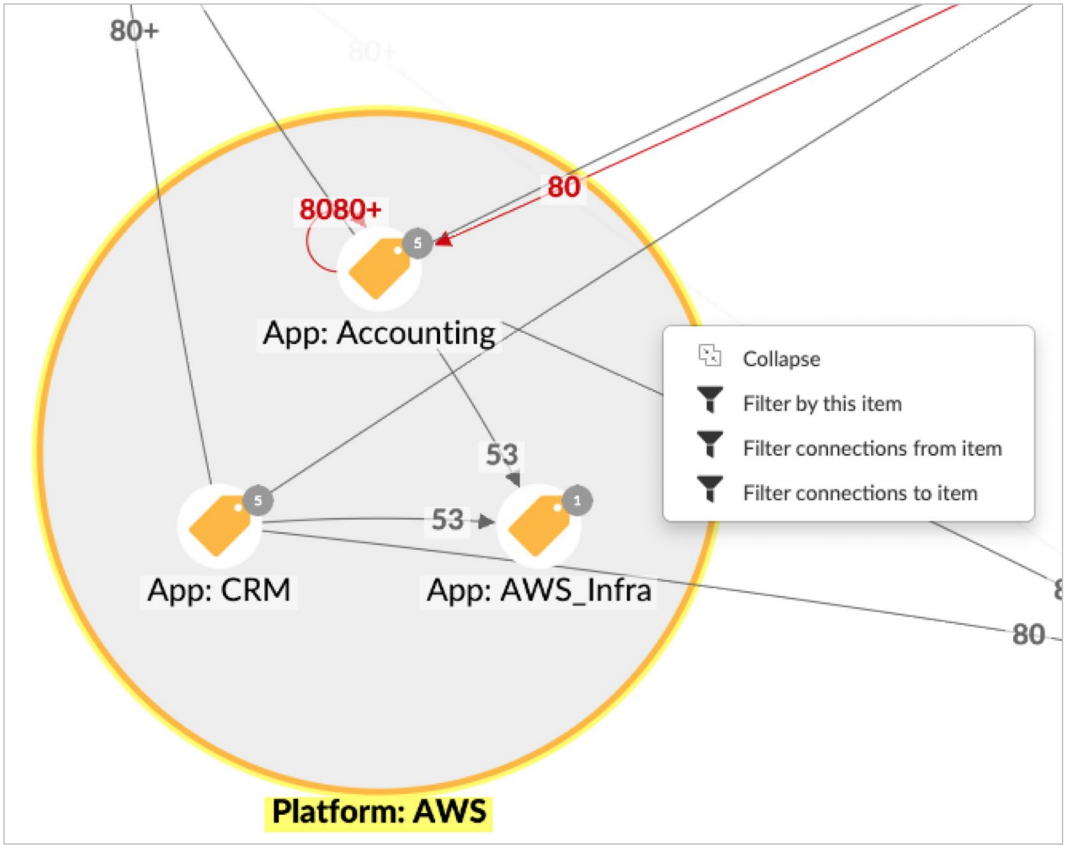


Figure 4. Filtering for a Specific Asset Environment

We expanded the *App: Accounting* assets in AWS to see distinct roles of assets within this application infrastructure—load balancers, web servers and databases (see Figure 5).

Security teams can easily drill into each of these distinct application assets. We highlighted the *Role: LoadBalancer* asset group and then selected the *Accounting-lb-1* server to get more details. Guardicore automatically shows analysts more detail about the selected asset, such as IP addresses, detected applications that are running on the asset and even cloud-specific information such as image information and cloud-native security metadata, including the AWS Security Groups (see Figure 6). This additional detail is immensely helpful for teams looking to enable junior or Tier 1 analysts to learn application behaviors and to be more successful and agile in identifying issues or specific asset information quickly.



Figure 5. Components of a Specific Application Environment

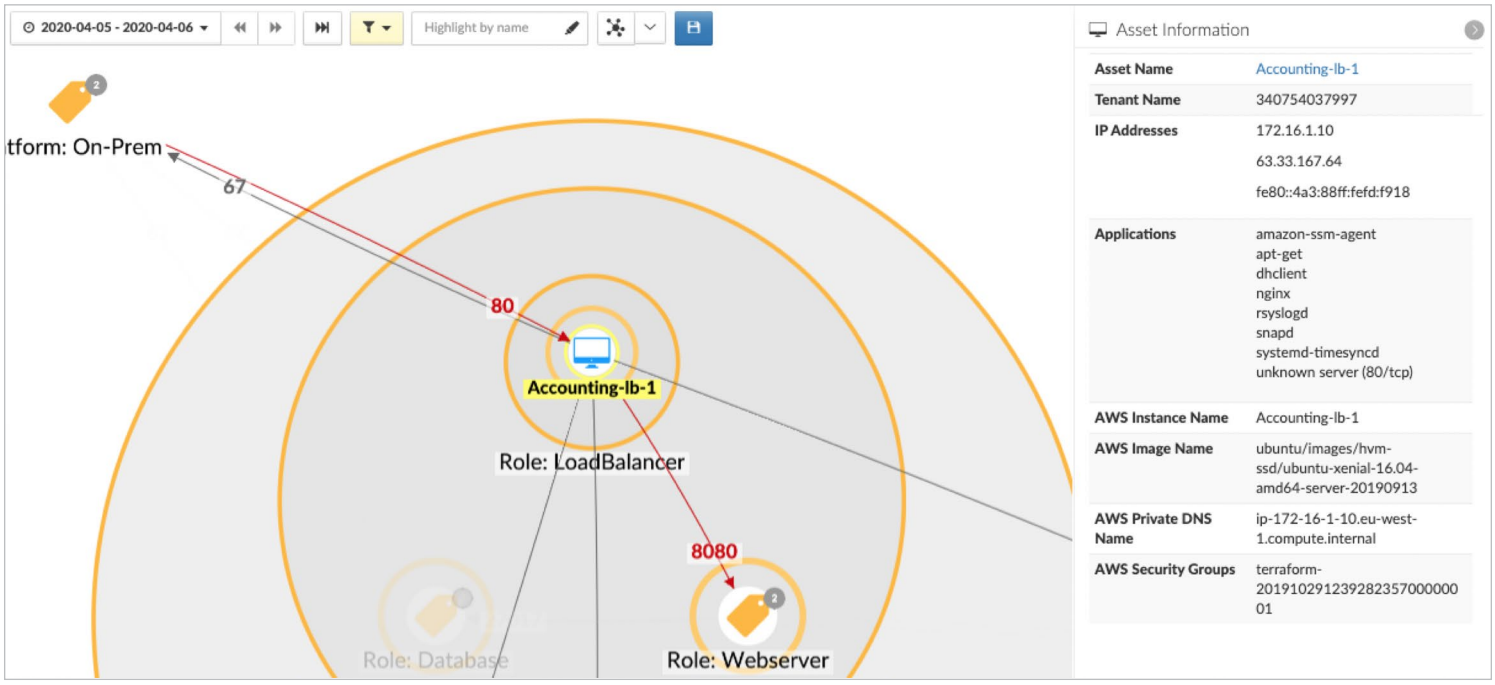


Figure 6. Asset Detail Breakdown with Cloud-Specific Metadata

By double-clicking on the asset, we can also see more detail about its detected running processes (see Figure 7).

Going even further, we clicked on the **`nginx`** process to see more detail about the process itself (see Figure 8). The capability to quickly and easily dig deeper into any assets to analyze (and potentially investigate) what is really running on the workload itself is a very important capability for any solution to offer. This capability shifts the use of a micro-segmentation platform from primarily an access control policy engine to being a useful analysis and investigation tool.

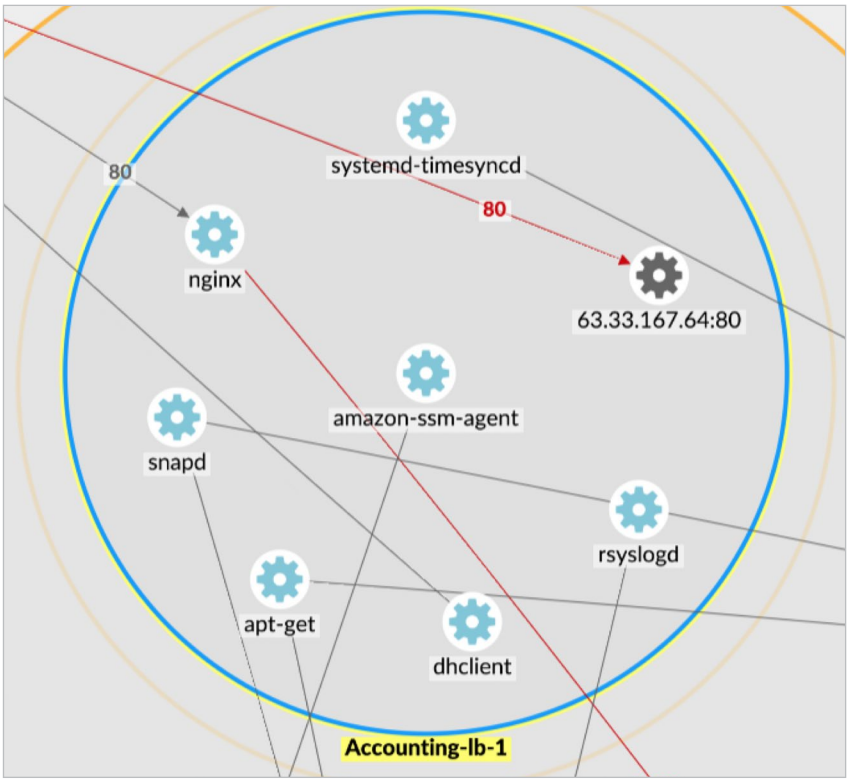


Figure 7. Detected Processes for an Asset

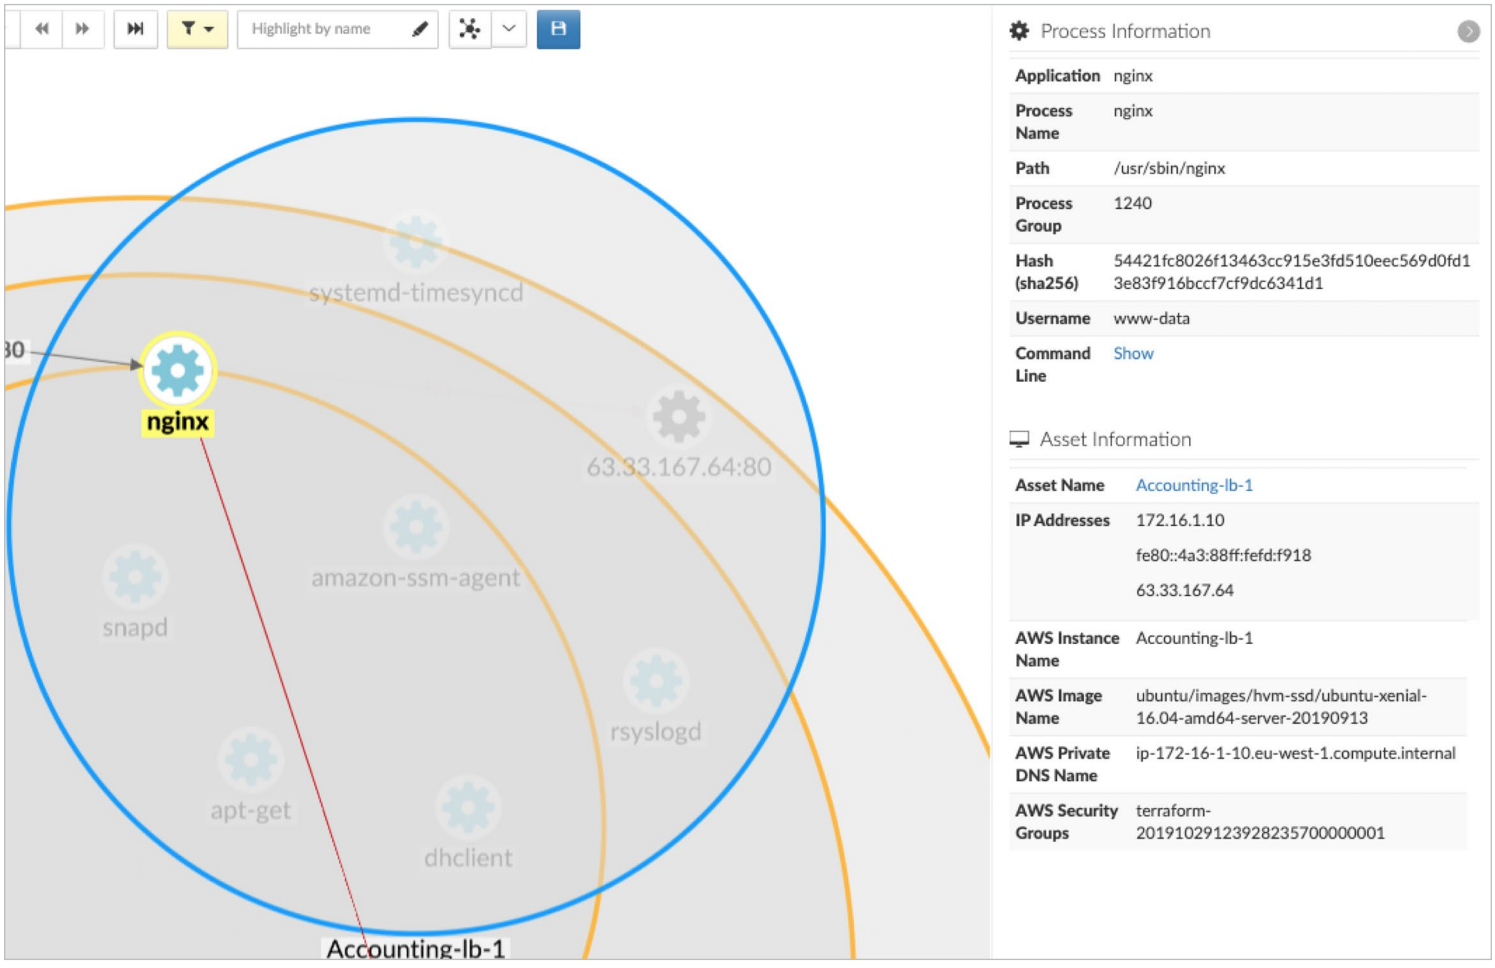


Figure 8. Detailed Process Information

Guardicore collects this data from a variety of sources:

- **Guardicore agent—**The Guardicore Centra agent is lightweight, with minimal overhead and utilization. Guardicore supports legacy and end-of-life operating systems such as Windows 7, 2008R2, Solaris, HP/UX, AIX and EoL Linux. It also has automated capabilities to add additional OS platforms within 48 hours of new kernels being released. Container environments require an agent on the container host. In serverless environments, Guardicore utilizes the cloud provider's API.
- **Read-only real-time query against the various platform APIs—**Centra can collect metadata from platform APIs, if available.
- **Guardicore Collector—**For network information, the Guardicore Collector can be implemented as a physical or virtual SPAN or TAP.
- **Guardicore REST API—**Guardicore supports integration with additional feeds from CMDBs and other enterprise data sources.

Guardicore Centra also collects data natively from environmental APIs (for example, AWS, Azure, Google Cloud Platform [GCP], Oracle and vSphere, among many other instances' metadata) that can aid in the allocation of dynamic labels. This makes mapping and policy creation easy, as well as making segmentation dynamic and auto-scalable. Labels can be dynamically automatable (or designated in playbooks and templates), which can eliminate the need for manual moves, adds and changes.

Centra Enforcement and Policies

One of the features that Guardicore Centra offers is consistent enforcement among Linux, Unix and Windows systems through a lightweight firewall agent. This brings about a number of enhancements over micro-segmentation tools that rely on native OS firewalls, such as the Windows firewall and Linux IPTables. First, this agent has been optimized for speed and efficiency, which likely reduces latency of policy evaluation and enforcement. Second, there are no requirements for administrator or root privilege use within the OS, which helps to sustain a least-privilege model of local behavior. Finally, this firewall is very capable and advanced across all major OS platforms, offering more features than the traditional Windows firewall or port/address specifications in many Linux firewall models. The engine also supports legacy and end-of-life OSes as well—which is critical for many enterprises today.

A robust micro-segmentation platform should also offer operations teams a wide variety of policy models. Centra includes both whitelisting and blacklisting policies, which could not be easier to create. In our test environment, we created a policy to detect and prevent `FTP` from being run. First, we navigated to the *Policy* menu and selected *Segmentation Rules*. This dashboard displays a list of all your policies in one place (see Figure 9 on the next page).

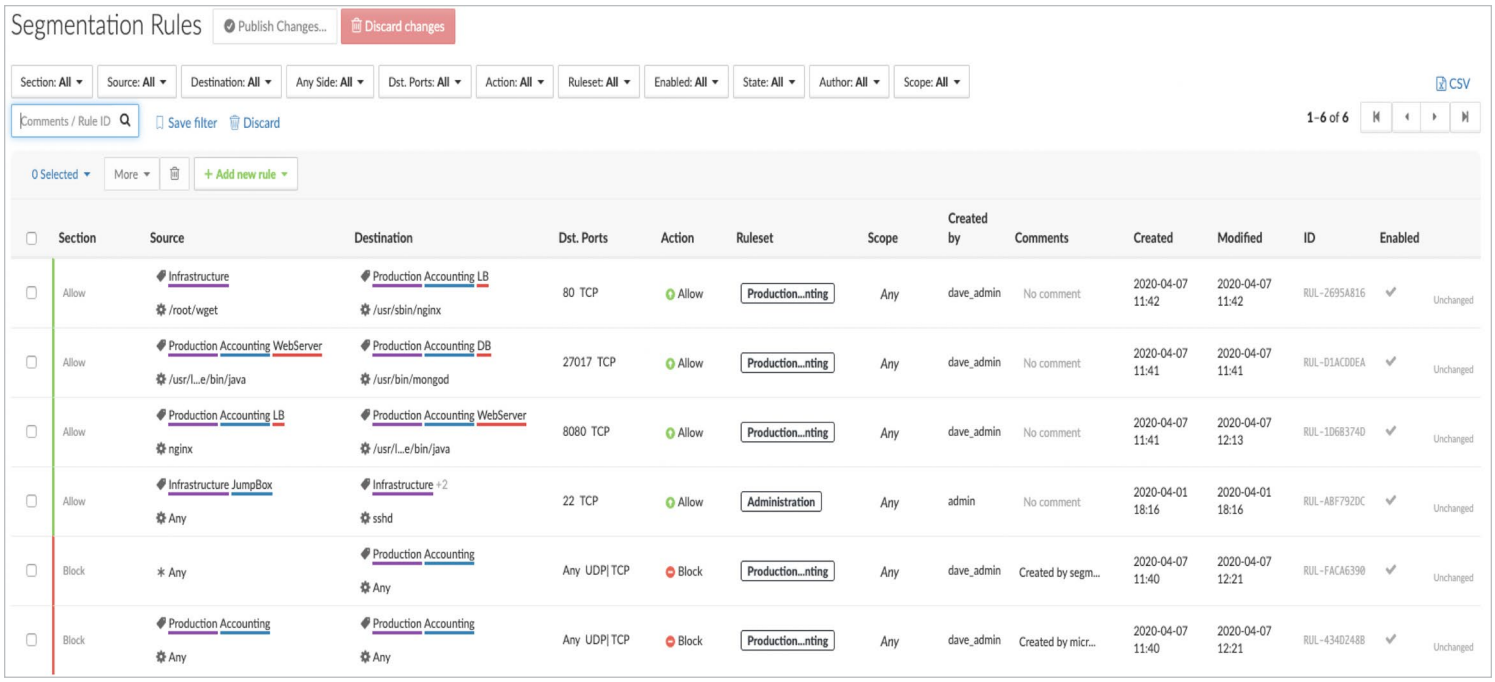


Figure 9. Segmentation Rules

We clicked on the *+ Add new rule* button and chose to create an Override Block Rule that would take precedence in the environment and "always apply." We chose a source of *Any* and then chose a destination of **`FTP`**. Rather than selecting a protocol or ports, Centra makes it simple to create a process-focused rule with the daemon/service name of **`ftpd`**, as shown in Figure 10. Note that blocking by process is more effective and efficient than trying to determine each individual port or range of ports in use for any given application flow.

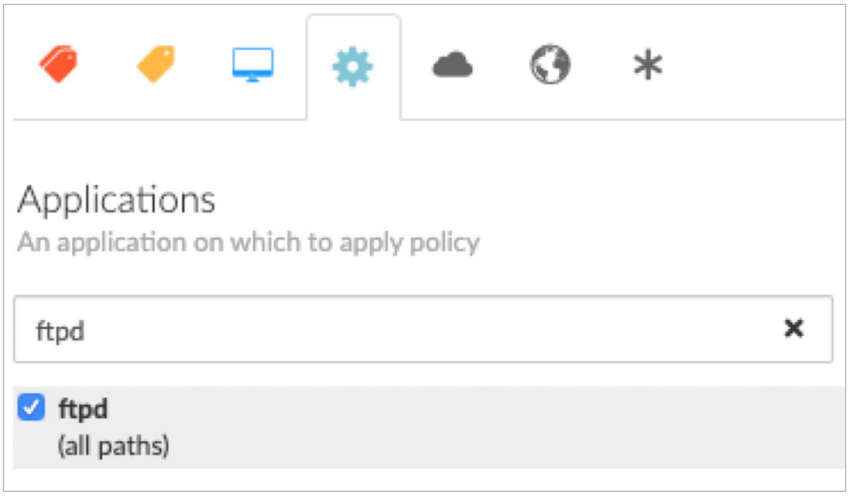


*Figure 10. Adding **`ftpd`** as a Rule Destination*

We then applied a label to the policy to assign it to an asset environment. For our purposes, we applied this rule to the Production environment. We then saved and published the policy (see Figure 11).



Figure 11. A New FTP Blacklist Rule

The entire process to create this rule was very intuitive and took less than one minute. The next type of policy we wanted to explore is at the heart of micro-segmentation deployments. After discovery has been enabled and assets have been identified, micro-segmentation platforms need to map out an application's dependencies and apply segmentation easily between an application and the world (ring fencing) and between the tiers of an application (micro-segmentation). To start, we visited the *Reveal* menu and explored Centra's data filter options. Centra has a variety of options for inclusion and exclusion, such as applications, assets, labels and label groups, as well as numerous conditional filters, such as connections to and from assets, addresses and protocols, and many more (see Figure 12).

For our first rule, we wanted to limit the Reveal visualization to include only the assets labeled *Application*, which we were able to accomplish easily with a data filter. We then wanted to look at only the Accounting application in the Production environment. To create both ring fence and micro-segmentation policies around this application and asset group, we clicked the *Edit Policy* button that starts up the Guardicore Centra Policy Wizard. Once the wizard opened, we highlighted the app visualization and chose *Create App Segmentation Policy*, as shown in Figure 13.

The policy editor dynamically created a policy with specific assets, local services and application components, and network communications/ports. We then followed the same process to create a micro-segmentation policy, which appeared in the policy pane as shown in Figure 14 on the next page.

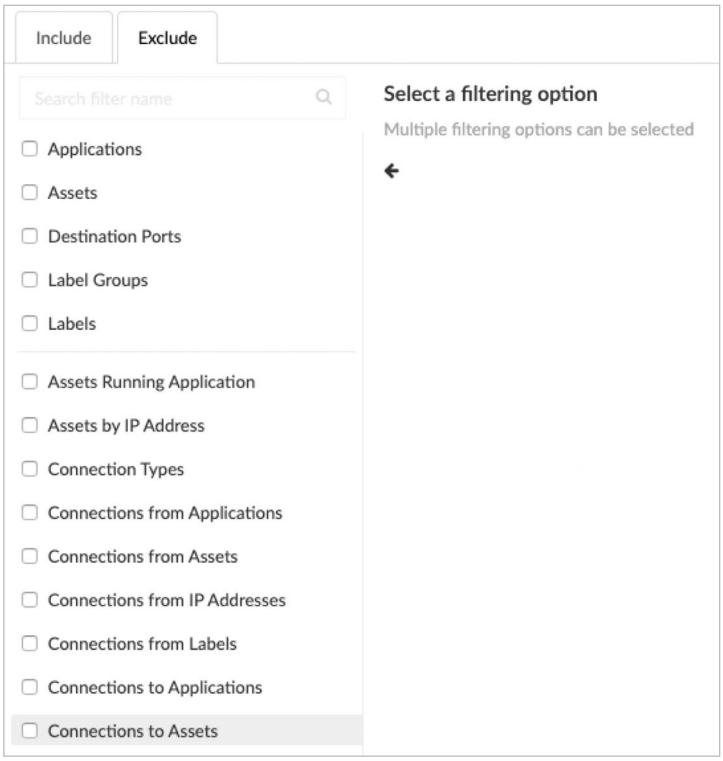


Figure 12. A Sample of Centra Data Filter Options

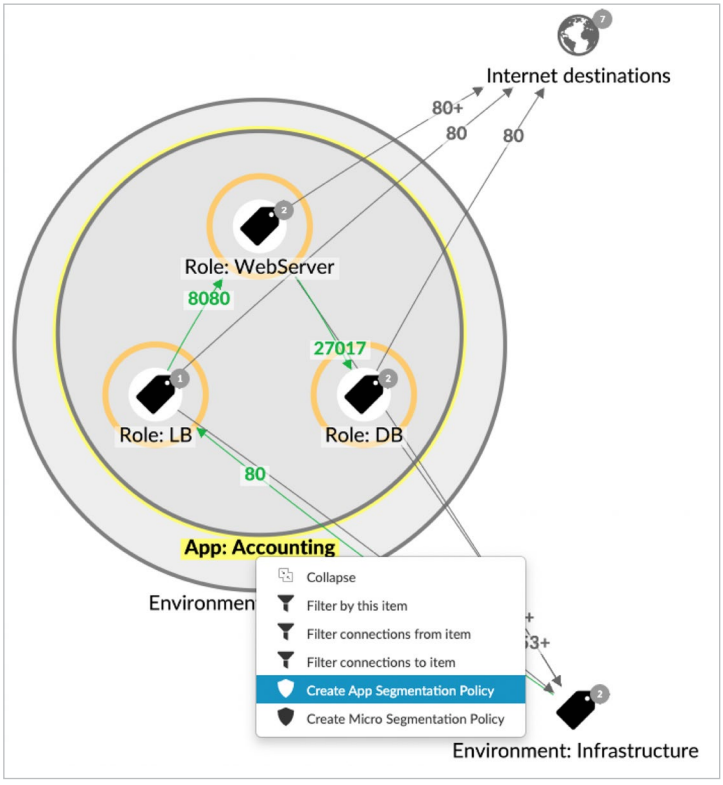


Figure 13. Dynamically Creating a Ring Fence Policy

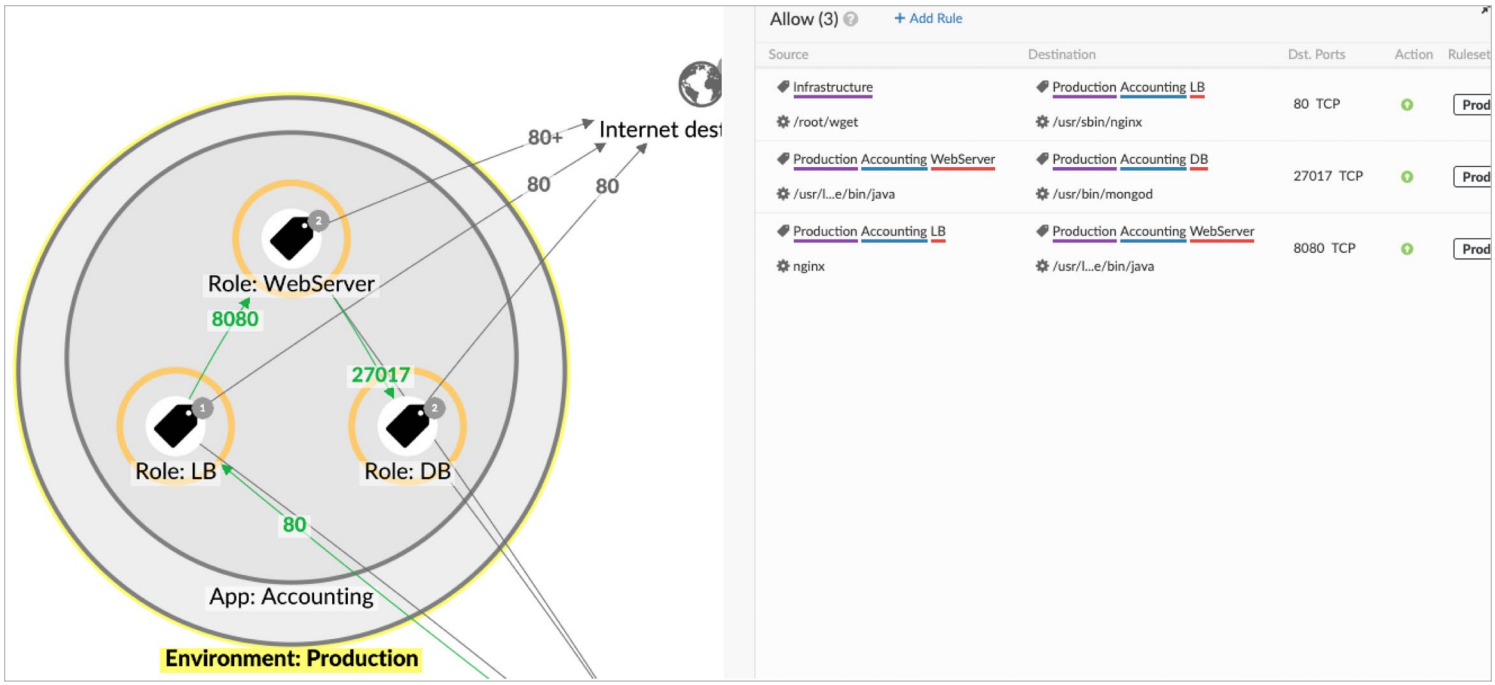


Figure 14. Dynamic Accounting Application Micro-Segmentation

We also modified several alerting policies to create blocking policies, which took less than one minute (shown in Figure 15).

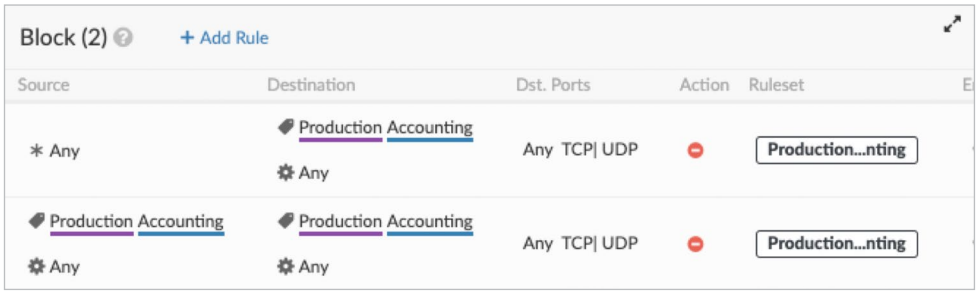


Figure 15. Alert Policies Changed to Block Actions

With these policies all now in place, we wanted to see any policy violations that might have been in existence within the current environment. We selected the *Incidents* menu and then navigated to the *Policy Violations* section. This revealed a variety of details about events and incidents in the environment. The one we focused on was the load balancer's **`nginx`** process communicating with the Accounting web service. But another service called **`attk`** was also communicating and, with our new policies, is now blocked (see Figure 16 on the next page). Guardicore's dynamic labeling capability can automatically label assets that meet specific criteria designated ahead of time to associate assets with specific groups, applications, and so on.

This exercise illustrated how useful the Centra platform can be in tracking down unusual behaviors and selectively isolating traffic and local application processes in policy creation.

The third type of policy we explored was focused on users and identity. In the *Segmentation Rules* section of the *Policy* category, we chose the *User Identity* types of rules to review. The team at Guardicore had a set of sample rules for us to dig into, created around the concept of IT administrator jump boxes for controlling access within

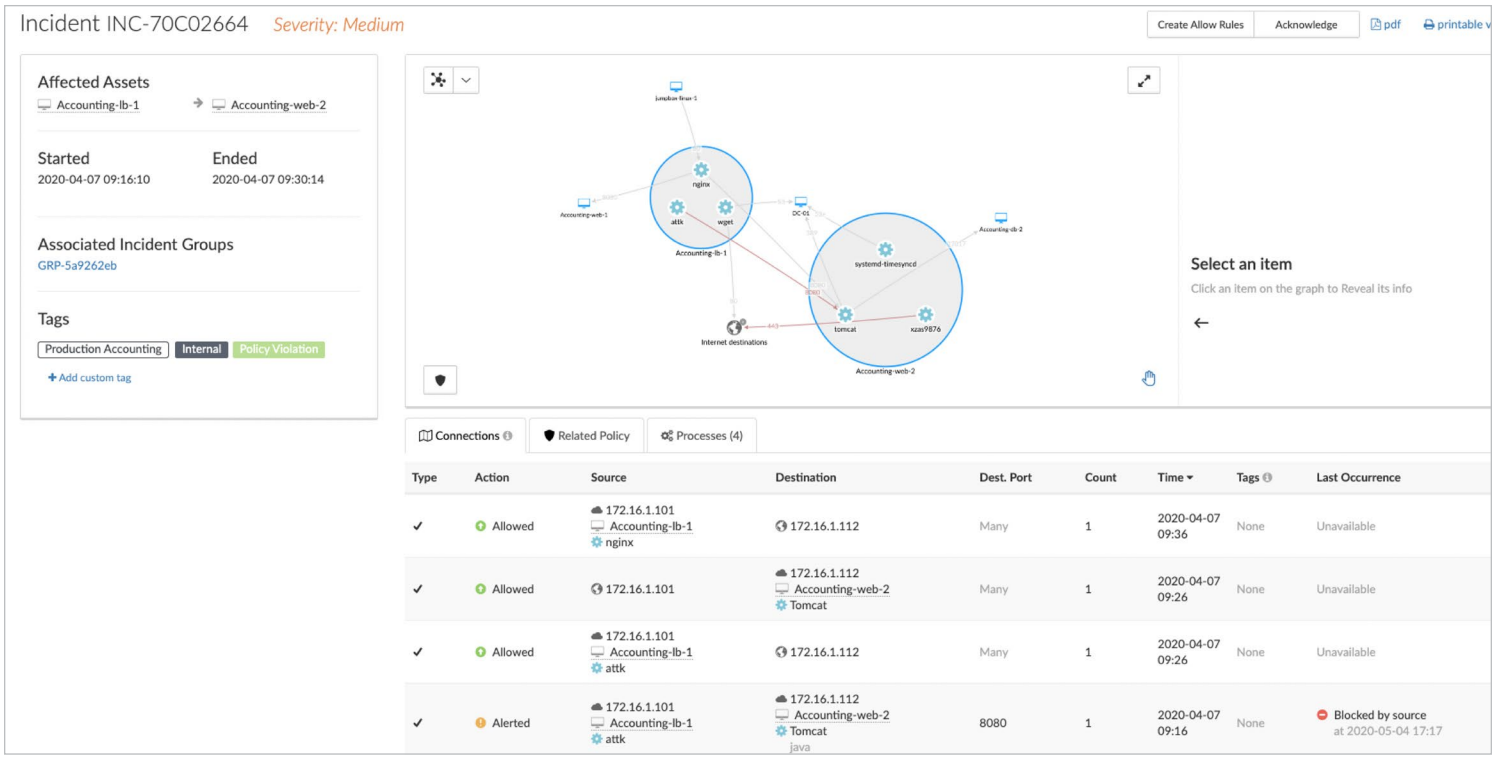


Figure 16. Policy Violation Details

the environment. The ability to tie the same series of jump boxes to different users and leverage different role-based segmentation rules is something that many organizations would likely embrace. Instead of having to manually segment using VLANs, security groups or firewalls, the grouping can be done with a micro-segmentation policy. This also means that various groups can use the same jump boxes, saving time, money and resources. There are three rules defined in Centra for this example (see Figure 17):

- The first rule allows server administrators to use Windows Jumpboxes in the Production and Common Service environments for **SSH** and **RDP**.
- The second rule specifically allows OrgPortal users access to web and **SSH** services within the OrgPortal application from the jump boxes.
- The third rule specifically allows Ecomm users access to the web and **SSH** services within the Ecomm app from the jump boxes.

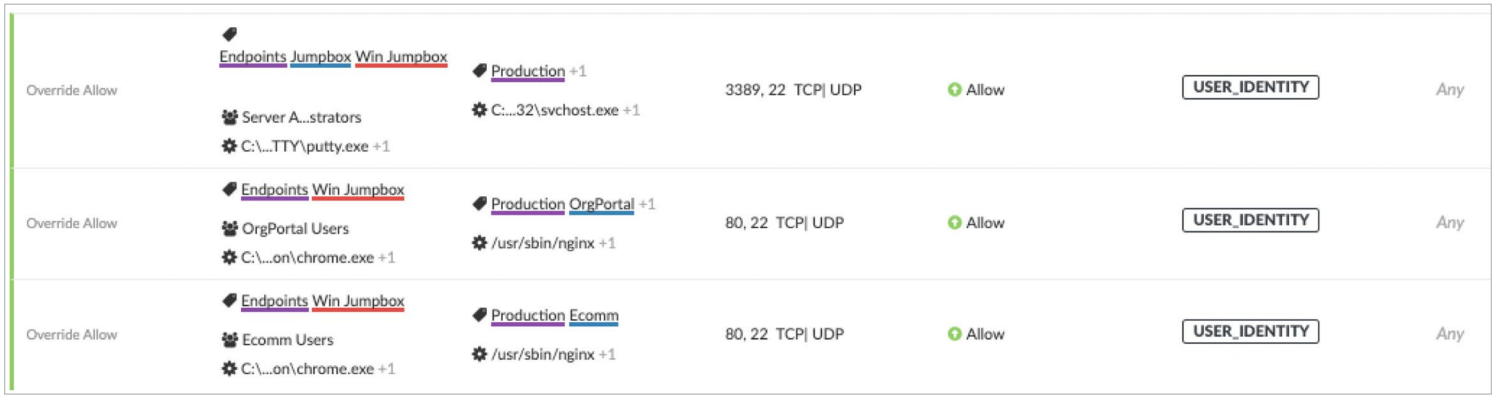


Figure 17. User-Oriented Policy Rules

The ease of understanding and mapping these users to policy rules demonstrates how flexible micro-segmentation can be, going beyond applications and components and helping drive real-world use cases in a more simplified way.

Our final example of policy rules focuses on the use of fully qualified domain names (FQDN) instead of IP addresses. Today, with the onset of cloud services and much more dynamic IT workflows, enterprises using auto-scaling and dynamic DevOps updates to containers and workloads are finding traditional network address rules to be difficult (if not impossible) to maintain. Anti-malware agent updates, patch and package distribution, kernel changes, **`GitHub`** and online code repositories may all use a variety of network addresses over time. Again, the team at Guardicore created a simple set of rules to review and analyze. In the first rule, assets in the Production environment can get to **`*.ubuntu.com`**, **`*.snapcraft.io`** and **`*.snapcraftcontent.com`** for web, secure web and **`ntp`** updates. All other internet access is denied (see Figure 18).



Figure 18. FQDN Rules

Altogether, we found the entire Centra policy engine to be highly intuitive and easy to configure.

Detection and Analysis

The capability to leverage micro-segmentation technology to better monitor and respond across all environments is another key feature that enterprise teams should evaluate. Beyond visibility and segmentation, Guardicore Centra also provides critical breach detection and response features, which we also evaluated.

The first capability we reviewed was the reputation services of local workload activity and behaviors that Guardicore offers. Guardicore has its own global sensor base that is aggregated into a reputation service that focuses on lateral movement through data centers and public cloud environments. With the agents deployed on your workloads, Centra evaluates every IP address, domain and process it sees and then can report on whether these are known to be trusted, untrusted or unknown. In the *Activity* section, we reviewed the *Reputation Log* to see what Centra found in our test environment. A snapshot of all malicious behaviors and sites is shown in Figure 19 on the next page. We filtered all the events to see only those items whose verdict was considered to be malicious.

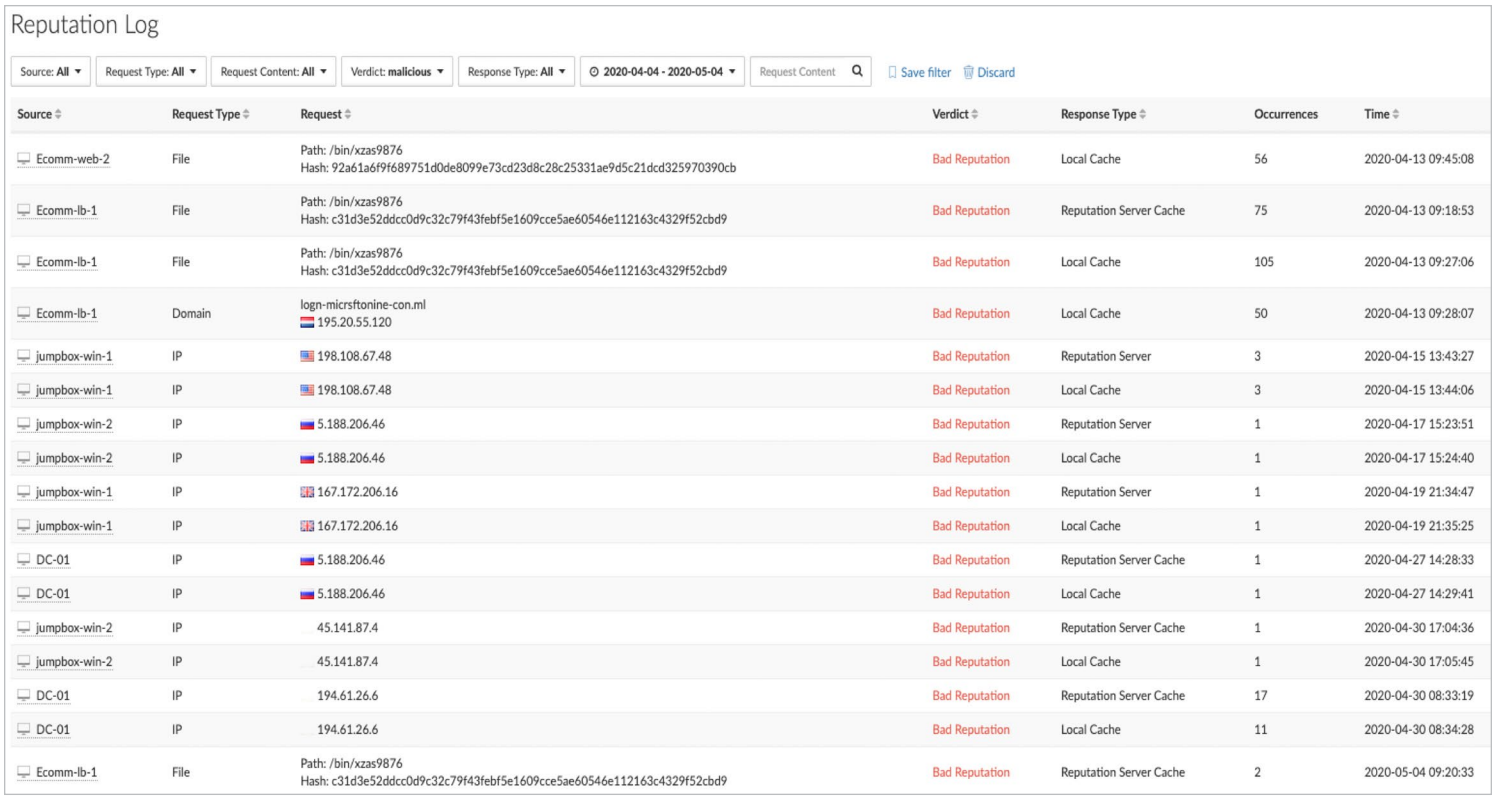


Figure 19. Reputation Events in Centra

Similarly, we can go even further to see what reputation events were involved in actual incidents detected within the Guardicore test environment. We selected the *Incidents* menu item and navigated to the *Bad Reputation* section. We drilled into a high severity incident alert to find out what was happening. This incident involved a malicious process (**`xzas9876`**) on the Ecomm application load balancer reaching out to a fake Microsoft update site that is a known credential-harvesting domain (see Figure 20).

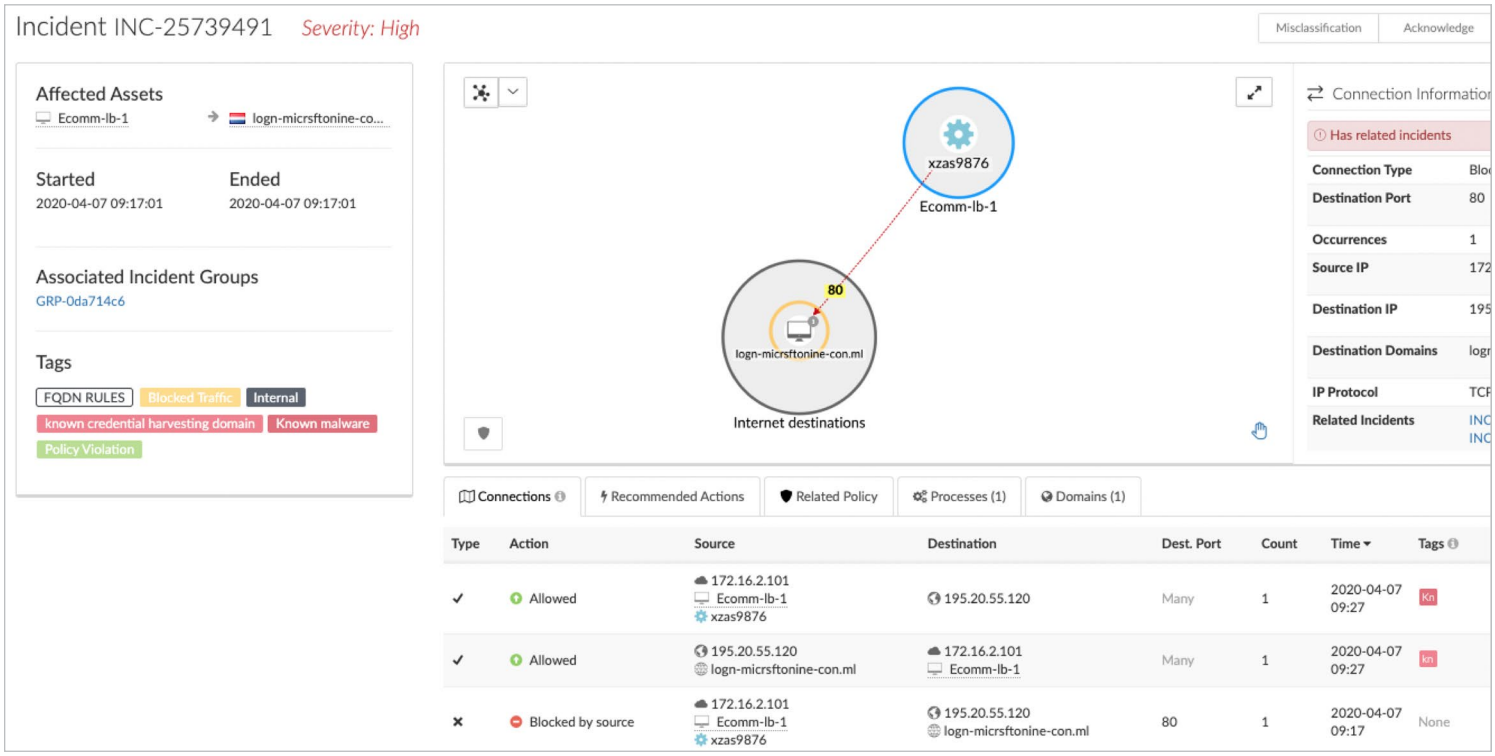
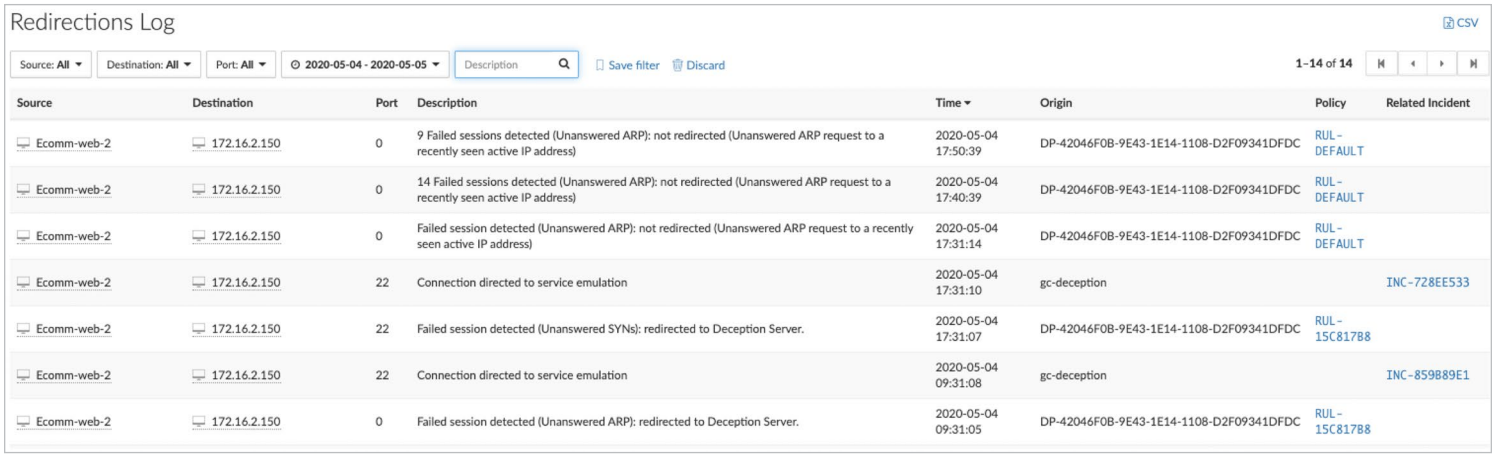


Figure 20. Reputation Incident Details

The amount of detail in this incident description was impressive. When we highlighted the connection in the visualization, additional connection information was displayed, along with related incidents and the action taken by Guardicore—in this case, blocking the connection. The detected malicious process information included the local path on the OS and the hash value to further enable threat hunting activities.

The next unique detection and response feature of Guardicore's Centra platform we reviewed was its dynamic deception capabilities. Guardicore Centra includes a dynamic honeypot function that monitors for suspicious connections that may indicate illicit lateral movement attempts in the environment and then redirects attackers to realistic honeypot servers for analysis. As Centra tracks legitimate assets in the environment, one of the benefits is reduction or prevention of false positives, capturing only traffic that is obviously suspicious. Suspicious connection attempts are redirected with a seemingly legitimate **TCP** three-way handshake that then leads to monitoring on the honeypots. We visited the *Activity* section again and chose *Redirections Log* this time around. Here, we can see all of the failed connections and whether they were redirected to the honeypot or not. The redirection behavior can be tuned, and the default is set to be stealthy (see Figure 21).


Redirections Log

CSV

Source: All | Destination: All | Port: All | 2020-05-04 - 2020-05-05 | Description | Save filter | Discard | 1–14 of 14

Source	Destination	Port	Description	Time	Origin	Policy	Related Incident
Ecomm-web-2	172.16.2.150	0	9 Failed sessions detected (Unanswered ARP): not redirected (Unanswered ARP request to a recently seen active IP address)	2020-05-04 17:50:39	DP-42046F0B-9E43-1E14-1108-D2F09341DFDC	RUL-DEFAULT	
Ecomm-web-2	172.16.2.150	0	14 Failed sessions detected (Unanswered ARP): not redirected (Unanswered ARP request to a recently seen active IP address)	2020-05-04 17:40:39	DP-42046F0B-9E43-1E14-1108-D2F09341DFDC	RUL-DEFAULT	
Ecomm-web-2	172.16.2.150	0	Failed session detected (Unanswered ARP): not redirected (Unanswered ARP request to a recently seen active IP address)	2020-05-04 17:31:14	DP-42046F0B-9E43-1E14-1108-D2F09341DFDC	RUL-DEFAULT	
Ecomm-web-2	172.16.2.150	22	Connection directed to service emulation	2020-05-04 17:31:10	gc-deception		INC-728EE533
Ecomm-web-2	172.16.2.150	22	Failed session detected (Unanswered SYNs): redirected to Deception Server.	2020-05-04 17:31:07	DP-42046F0B-9E43-1E14-1108-D2F09341DFDC	RUL-15C817B8	
Ecomm-web-2	172.16.2.150	22	Connection directed to service emulation	2020-05-04 09:31:08	gc-deception		INC-859B89E1
Ecomm-web-2	172.16.2.150	0	Failed session detected (Unanswered ARP): redirected to Deception Server.	2020-05-04 09:31:05	DP-42046F0B-9E43-1E14-1108-D2F09341DFDC	RUL-15C817B8	

Figure 21. Centra Redirections Log

To see dynamic deception incidents, we visited the *Incidents* menu and selected *Lateral Movements*. Here we can review all the types of deception interaction and behavior associated with the redirected traffic. We investigated several incidents, uncovering a wide range of useful data available for defenders and investigators, including summary information, a full session recording, screen shots, all of the files and processes affected, credentials used and even a full PCAP of the event, as shown in Figure 22 on the next page.

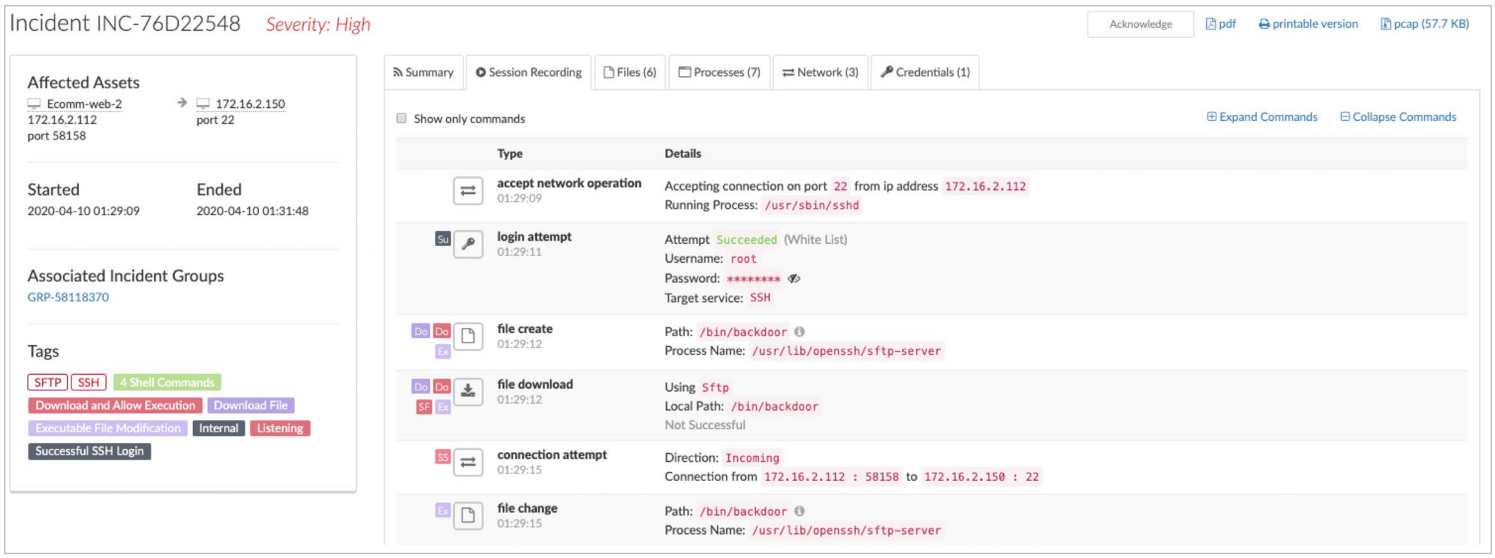


Figure 22. Dynamic Deception Incident Details

In the past, many organizations were leery of operating honeypots in their environments to avoid enticing hackers. Today, however, the technology has dramatically improved and matured, and attackers are already inside compromised networks looking for new systems and services to attack. Detection of these lateral movement scenarios is more critical than ever—and deception technology is rapidly gaining ground in identifying malicious activity sooner and with more detail to analyze attacks (or attempts) in progress.

Centra also includes file integrity monitoring capabilities. Security teams can specify files on critical systems for integrity monitoring, and Guardicore can track those files continuously for any violations. We looked at all detected integrity events in the *Integrity Log* pane within the *Activity* menu to see any violations. From there, just as in the previous examples, we followed up by visiting the *Incidents* menu and its *Integrity Violations* section. We delved into an example incident where two of the Ecomm systems have **`/etc/tomcat7/tomcat-users-xml`** files that were manipulated with bad file hashes (see Figure 23).

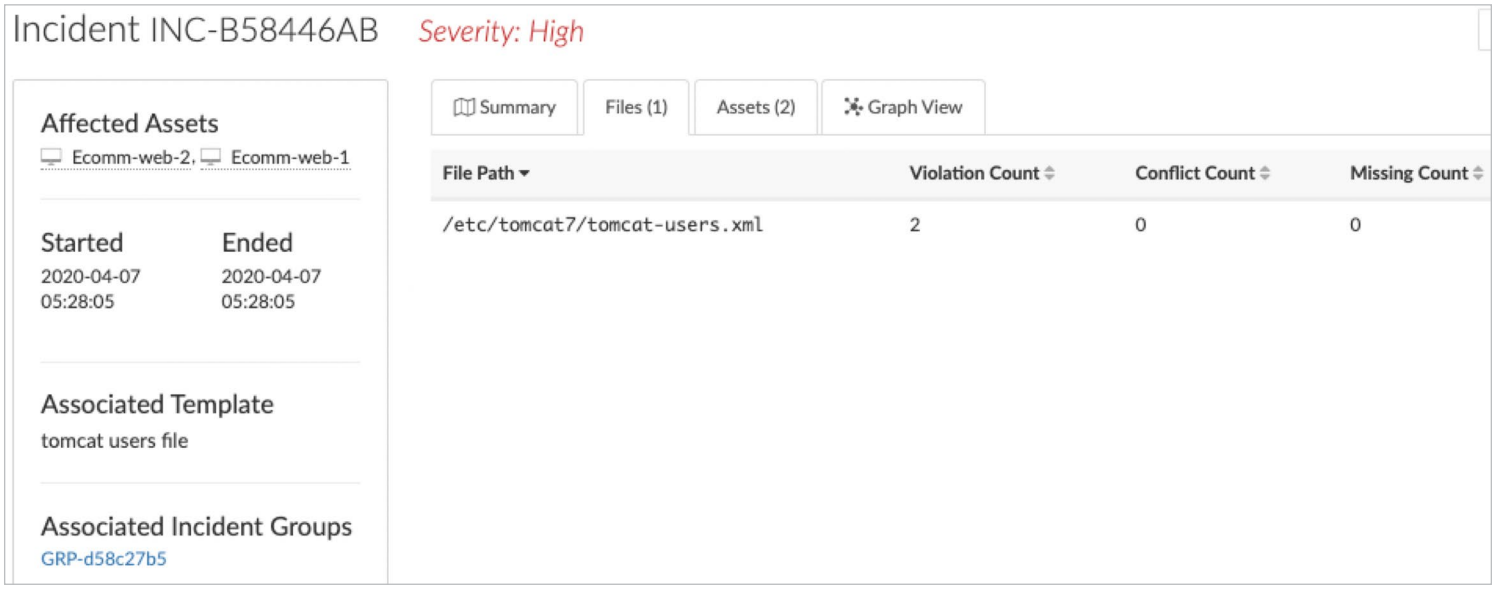


Figure 23. File Integrity Monitoring Incident Details

Finally, Guardicore Centra includes an extensive and well-documented RestAPI set. This can enable easy integration with many enterprise solutions including CMDB and deployment playbook platforms such as Chef, Puppet and Ansible, among others. Although we didn't integrate any other solutions as a part of this review, we did explore some of the API documentation, which seems current and highly detailed (see Figure 24).

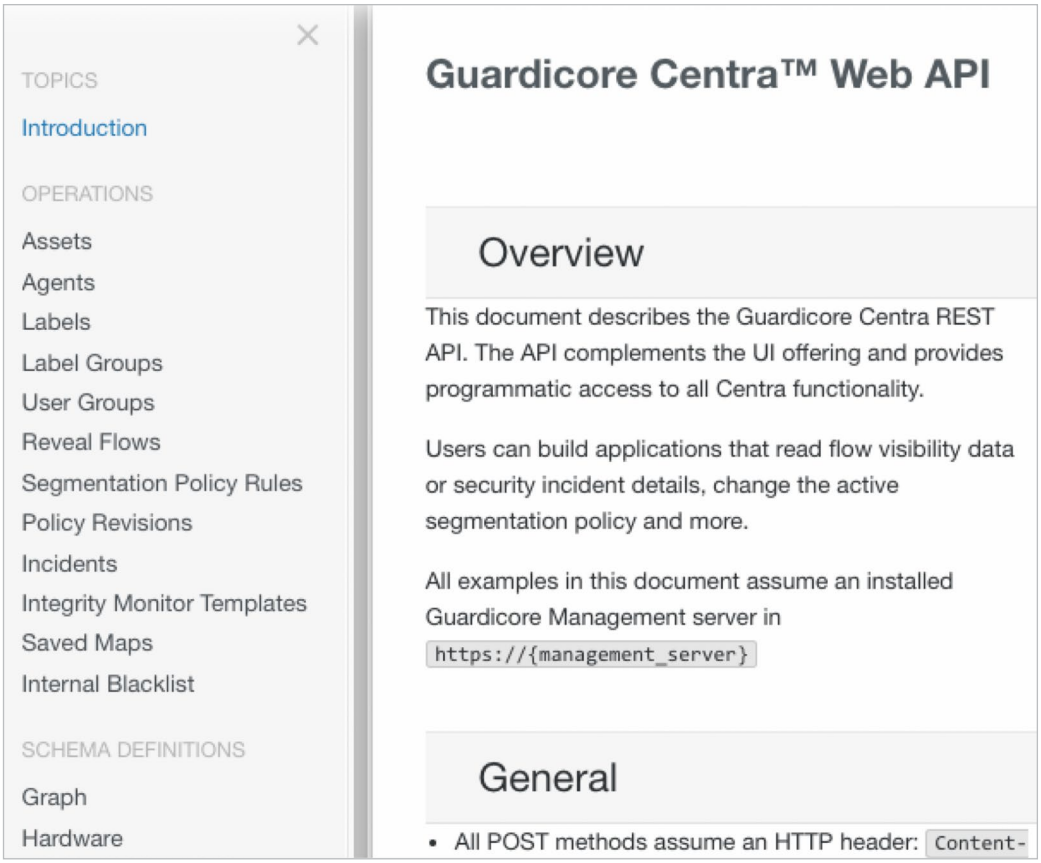


Figure 24. Centra API Documentation

Review Wrap-up and Conclusions

While the terms *micro-segmentation* and *zero trust* have been discussed often in IT operations of late, many organizations have struggled to find practical paths to implement technology that actually achieves the goals of micro-segmentation. Mapping assets, their actual behaviors and local components into logical policies has proven daunting, both conceptually and tactically. Based on our review of Guardicore Centra, we believe that these challenges are surmountable. The product was easy to use and offers a wide range of intuitive and powerful policies to implement ring fencing, internal micro-segmentation and much more.

Beyond just the micro-segmentation and access control outcomes, this solution provides a level of in-depth understanding and visibility into the environment that brings additional benefits in the form of detection and response capabilities. Reputation and monitoring services were useful, and Guardicore also supports additional threat intelligence data. Adding dynamic deception capabilities into the platform adds a whole new layer of depth and capability to this product. Many security operations teams should eagerly try Centra out for themselves. Deception technology can save security teams a lot of time and provide deep forensic data to boot.

As the rate of change in IT workload deployment and cloud service integration increases, the traditional models of access control increasingly fails us. Static, stationary firewalls and network segmentation tools and tactics just aren't keeping pace with the way organizations want to build and deploy infrastructure today. Micro-segmentation tools such as Guardicore Centra have a lot to offer and will likely help advance protective, detective and response activities as organizations wrestle with hybrid environments that include legacy platforms, convoluted internal networks and cloud service infrastructures.

About the Author

Dave Shackleford, a SANS analyst, senior instructor, course author, GIAC technical director and member of the board of directors for the SANS Technology Institute, is the founder and principal consultant with Voodoo Security. He has consulted with hundreds of organizations in the areas of security, regulatory compliance, and network architecture and engineering. A VMware vExpert, Dave has extensive experience designing and configuring secure virtualized infrastructures. He previously worked as chief security officer for Configuresoft and CTO for the Center for Internet Security. Dave currently helps lead the Atlanta chapter of the Cloud Security Alliance.

Sponsor

SANS would like to thank this paper's sponsor:

