JUNIPER NETWORKS

LEARN ABOUT SOFTWARE-DEFINED SECURE NETWORKS (SDSN)

Learn about the newest security measures from Juniper Networks that provide end-to-end visibility into your entire network, both virtual and physical. It's called SDSN.

By Madhavi Katti, Keerthi Latha M R, Susan McCoy, & Sushma Sethuram

LEARN ABOUT:
SOFTWARE-DEFINED SECURE NETWORKS (SDSN)

As the scale and sophistication of cybercrime continues to increase, businesses need to rethink their defense strategies. To be truly effective, next-generation security must be built around automated and actionable intelligence that can be shared quickly to recognize and reduce threats, protecting the network and its users.

Juniper'sSDSN gives you end-to-end network visibility that secures the entire network, physical and virtual. It leverages cloud economics to find and stop threats faster than perimeter firewalls.

The unified SDSN platform combines policy, detection, and enforcement with a comprehensive product portfolio that centralizes and automates security:

- Policy: Simplified, centrally-managed policies are intelligible for all devices on a heterogeneous network.
- Detection: Threat intelligence is aggregated into a common, cloud-based service with policies that adapt to changing threat conditions, providing fast and effective protection.
- Enforcement: Updated policy is distributed across the network, dynamically and in real time, stopping rogue traffic and quarantining compromised endpoints.

SDSN transforms your network into a single, holistic defense domain where every element becomes an enforcement point. This is the future of the secure network.

Juniper Networks Books are singularly focused on network productivity and efficiency. Peruse the complete library at www.juniper.net/books.

Published by Juniper Networks Books

ISBN 978-1936779437

52000

9 781936 779437

JUNIPER NETWORKS

Learn About **SDSN**

JUNIPER NETWORKS

Published by Juniper Networks Books
Authors: Madhavi Katti, Keerthi Latha M R, Susan McCoy, Sushma Sethuram
Technical Reviewers: Thomas DiMicelli, Scott Emo, Rajan V K, John Lelane, Indira Upadhayaya
Editor in Chief: Patrick Ames
Copyeditors: Atanu Raychaudhuri, Nancy Koerbel
Illustrations: Karen Joice

ISBN: 978-1-941441-43-5 (paperback)
Printed in the USA by Vervante Corporation.

ISBN: 978-1-941441-44-2 (ebook)

Version History: v1, November 2016
2 3 4 5 6 7 8 9 10

http://www.juniper.net/dayone

This book is a compilation of eight separate *Learn About* documents that address various aspects of securing enterprise networks, both physical and virtual. They first appeared in Juniper's *Network Design and Architecture Center* on SDSN in 2016.

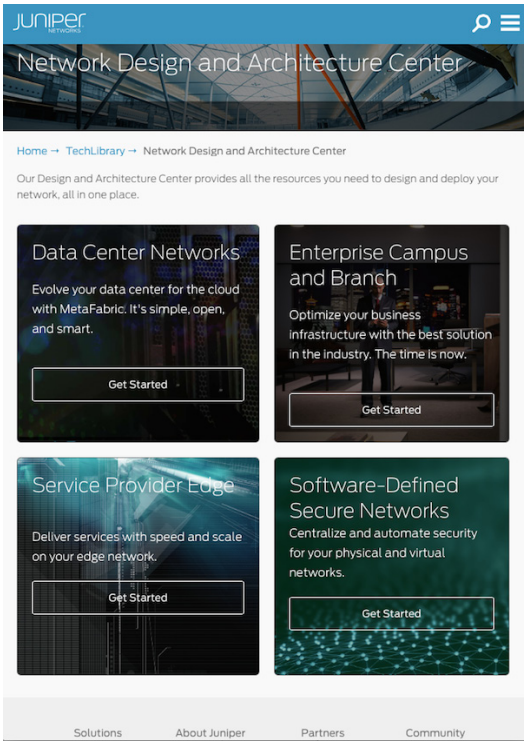

Chapter 1

SDSN and the Evolving Threat Landscape

A recent, in-depth report by economic and cybersecurity experts at RAND (http://www.rand.org/pubs/research_reports/RR1024.html) found that chief information security officers are faced with a chaotic and confusing landscape when deciding how to manage the risks (and costs) associated with providing security to their businesses.

What's more troubling is that RAND'S research shows even though companies are increasing spending on cybersecurity tools, they are uncertain as to whether or not such spending is making their infrastructure more secure. Many businesses do not know when, or whether, they have invested enough in their security strategy. Of even greater concern is the common belief that attackers are always gaining quickly on new protections.

Managing risk is a misunderstood concept in cybersecurity, usually focused on risks posed by threats and vulnerabilities instead of risks specific to business outcomes and operations. When managing risk, much of the emphasis – even the metrics used to demonstrate the value of security programs – is often based on the ability of a particular tool or program to stop a certain number of attacks, instead of on the metrics that matter most to the business.

Instead of measuring the quantity of blocked attacks, the goal of a comprehensive security program must be to understand the financial return between managing where to invest and the risk of not adopting various security measures. Therefore, although stopping attacks is an imperative, it must also be balanced with reducing the risk to business if an attacker were to get through.

What made these seasoned security experts nervous? Unfortunately, it is the ability of advanced threats to bypass traditional perimeter security defenses, enter a trusted network, and move about undetected. Table 1.1 lists a ZDNet Special Feature about the cybersecurity industry's 2015 predictions, in which the top prediction was about "new attack vectors and platforms and the evolution of existing cybersecurity solutions." (Source: ZDNet Special Feature: Security and Privacy: New Challenges)

Table 1.1 Cybersecurity Predictions for Year 2015

Number of Predictions	2015 Predictions
12	New attack vectors and platforms
11	Evolution of existing cybersecurity solutions
10	IoT and critical infrastructure
9	Mobile technology
9	Encryption and privacy
9	High-profile data breaches
8	Regulation, compliance, and cyberinsurance
8	Security strategy evolution
7	People and social networks
7	Advanced threat intelligence and prevention
6	State-sponsored and politically motivated attacks
6	Ransomware
5	Cloud services
4	Big data and analytics
4	POS and payment systems
3	Biometrics and multifactor authentication
3	Cybercrime

The traditional network security posture of building a strong perimeter defense, where devices at the edge are the primary means of defense for all types of threats, is becoming less and less effective. Why? Because the perimeter defense that checks everything coming inside is based on the trust/no trust model – *trust what's inside the network, don't trust what's outside coming in*. This model is no longer pertinent nor sufficient. Advanced threats can bypass traditional perimeter security defenses, and they can enter the trusted areas and stay there undetected. That greatly increases the *surface area* of the attack because the perimeter has now been lost.

How and *where* one needs to deploy security has changed.

The Perimeter Collapses

How has this happened? Here are some recent developments that have weakened the perimeter approach to network security:

Failing to deploy at all endpoints – Any smartphone or tablet that connects to a corporate network is an endpoint, and must be secured. If employees in an organization use VPN to connect from remote locations, their devices become endpoints, too. And with the proliferation of BYOD and the Internet of Things (IoT), trying to secure endpoints is nearly impossible.

Increasing threat sophistication – Threats have changed from phishing, malware, and morphing executables to security hackers who infiltrate enterprises to retrieve data for financial gain. The attacks are targeted and focused, and use advanced persistent threat (APT) processes. The attackers have the advantage of time on their side, while their victims have the disadvantage of organizational and network complexity working against them. It is crucial to control botnets' attack vector before it can be used as an APT and migrated into mobile devices.

Insider threats – You now need to secure every point of access in your network, because the threat is not just the intruder trying to break in. The threat can be your employee who has walked through your front door with malware on their device, or an employee who was developing software in a container and accidentally copied malware into the code. The security posture is quickly evolving into *zero trust* of anything entering or leaving the network.

Virtualization and cloud infrastructure – The burgeoning growth of cloud services adds new layers of infrastructure complexity. Security issues can include data breaches, handling security incidents, sensitive data access, exploited system vulnerabilities, malicious insiders, management console security, account control, and multitenancy issues. Failure to ensure appropriate security protection when using the cloud can result in higher costs and loss of business.

Managing security across multiple environments operating in an incoherent and unorchestrated manner – The use of different public/private cloud offerings by organizations can make network security operations extremely complex. The daily proliferation of security devices and policy enforcement points further adds to the complexity of new application deployments. If your management solutions are slow, unintuitive, or restricted in their level of granularity and control, your network security management becomes overly time-consuming and prone to error. The result? Your unorchestrated security management can actually make your network vulnerable to cyberthreats.

Coping with This Threat Landscape

To cope with today's broad threat landscape, you need threat intelligence and immediate threat migration. And that means a comprehensive security platform that can tie together and coordinate various threat analytics platforms. You also need a method of providing a simpler policy mechanism. Above all, you must be able to leverage the entire network, not just the perimeter, as a threat detection and threat migration solution.

The paradigm is changing to security solutions that can deliver comprehensive yet coordinated protection by:

- Integrating and deploying advanced security features to protect systems and data from spyware, viruses, malicious code, denial-of-service attacks, and more.
- Building a network flexible enough to deliver new services without causing a security gap and with improved performance and resilience (high availability).
- Using policy automation to adapt and enforce policy in real time and improve both compliance and business agility.
- Providing an automated, end-to-end network with endpoint security across the Web, e-mail, files, and applications.
- Creating comprehensive visibility into the entire network so that you can identify and address threats wherever they are detected, inside or outside the network perimeter.

The paradigm is also changing to solutions that can combat the new breed of hackers, by treating the internal network just like the Internet or an untrusted network, and by:

- Treating every port in the network as untrusted, including ports outside your network, inside your network, and between endpoints and cloud applications.
- Encrypting *all* communications with advanced encryption techniques to secure network communications.
- Understanding the new threat landscape and quarantining infected endpoints in real time, and making the network adaptable, to automatically detect aberrant behavior and immediately respond.

And finally, the paradigm is morphing into using the network itself as a detection and enforcement ecosystem by:

- Enabling every part of the network to be a detection as well as an enforcement point to respond to suspicious activity anywhere in the network, which is the most effective way to deal with threats and intruders.
- Centralizing the security policy engine so that it can determine trust levels between network segments by collecting real-time threat information.
- Closing the gap between threat intelligence and policy enforcement because threat intelligence loses most of its value if it is distributed too slowly, or if it does not reach all of the enterprise's enforcement points.
- Creating a unified security policy, with distributed new policies implemented in real time from a central location.

SDSN – The Path Forward

SDSN is Juniper's vision of applied threat intelligence and immediate security enforcement. It is the future of network security because by leveraging the cloud, it can more effectively and dynamically solve the litany of current network security issues cited earlier in this chapter.

SDSN works on the following principles:

- Leverage the entire network as points of threat detection and policy enforcement.
- Leverage the economy of the cloud to share threat intelligence at scale and to accelerate threat detection and make it adaptable in real time.
- Implement a centralized controller/policy engine that dynamically adapts policy and so stays ahead of constantly evolving threats and attacks.

Figure 1.1 illustrates the building blocks of a Software-Defined Secure Network that includes advanced firewalls for the branch and the data center, threat intelligence, orchestration, and cloud-based protection. You can see in Figure 1.1 that SDSN is the industry's only inside-out security model because it leverages the network as a sensor for delivery of context-aware threat alerts and then dynamically enforces security policy with software-defined segmentation designed to reduce the overall attack surface. SDSN promotes a *zero trust* model for information security that is fundamentally more secure, because even if one application on the network is compromised, you can isolate that infection or threat. You can protect the more critical assets inside your network.

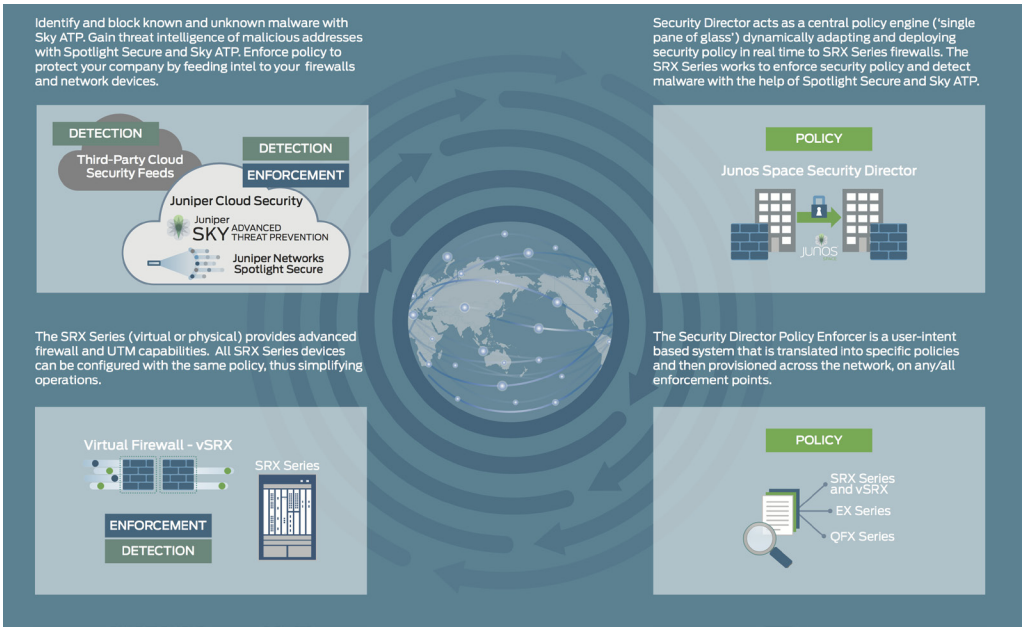

Figure 1.1 SDSN Building Blocks

Juniper Networks SDSN has the following salient features:

- *Visibility into network.* With SDSN, you get full visibility of all the traffic, whether it is North–South or East–West. The entire network infrastructure is operationalized and managed as a single enforcement domain, providing enforcement points across the network where policy can be deployed dynamically, and in unison, to block threats anywhere.
- *Comprehensive security.* With SDSN the firewalls, virtual or physical, are right-sized for their application in the network – for example, to provide advanced firewall and unified threat management (UTM) capabilities whenever needed. These capabilities are consistent across both physical and virtual platforms; therefore, with SDSN the device not only meets the security requirement but is also configured with the same policy as other devices, thereby simplifying the operation of the whole network.
- *Policy orchestration.* In an SDSN network, cloud-based security services provide the foundation for an open policy engine especially when you have security controllers that can push those policies into the network. By providing real-time feedback

between firewalls, the controller plus the cloud can deploy policy across network devices the instant it is needed.

- *Third-party integration*. SDSN is grounded in integrating third-party capabilities, in an effort to unite the *good guys* against the *bad guys*. The Open Convergence Framework (OCF) provides integration with technical alliance partners, enabling SDSN networks to choose their preferred threat intelligence information sources.

Threats can be detected faster with an SDSN approach because SDSN detects threats as they evolve, by leveraging threat intelligence from multiple sources (including third-party feeds) and tapping into the power of the cloud. When security policies are enforced consistently, even in global networks, network security can adapt dynamically to respond to real-time threat information.

Resources and References

Resources for learning more about SDSN, including product information, solution briefs, and press releases as well as Juniper's Network Design and Architecture Center, can be found here:

- http://www.juniper.net/documentation/en_US/design-and-architecture/software-defined-secure-networks/index.html
- http://www.juniper.net/us/en/solutions/software-defined-secure-networks/
- https://www.juniper.net/assets/us/en/local/pdf/solutionbriefs/3510569-en.pdf
- http://investor.juniper.net/investor-relations/press-releases/press-release-details/2015/The-New-Economics-of-Defense-First-of-Its-Kind-Heuristic-Model-Empowers-Companies-to-Make-Smart-Security-Investments/default.aspx

YouTube SDSN Video with an overview of the solution, features, and benefits:
https://youtu.be/dTMGw5Byi8E

A recent, in-depth report by economic and cybersecurity experts at RAND:
http://www.rand.org/pubs/research_reports/RR1024.html

ZDNet Special Feature about the cybersecurity industry's 2015 predictions:
http://www.zdnet.com/article/cybersecurity-in-2015-what-to-expect/

Chapter 2

Advanced Threat Prevention (ATP)

Malware is malicious software that disrupts network operations and gathers sensitive information on behalf of an unauthorized third party. While the majority of malware attacks are unfocused, new attacks are now using disguised threats that go after specific targets.

Targeted malware employs sophisticated methodology to evade traditional security defenses *in order to embed itself in the target's infrastructure*. Once attached, the malware can carry out a wide range of undetected malicious activities over months, or even years, including data theft, espionage, and disruption or destruction of network infrastructure and processes.

Unfortunately, there have been more than enough recent examples of malware attacks on major hotel chains, city infrastructures, and financial institutions. Examining some of these attacks in detail will illuminate the kind of advanced threats your network may be vulnerable to, and why technology and why technology like Juniper Networks Sky Advanced Threat Protection (Sky ATP) is so essential.

Point of Sale (POS) Malware

In December of 2015 Hyatt Hotels disclosed that they had experienced unauthorized access to their credit card payment system at over 250 of their hotel sites between August and December of that year. Hyatt didn't disclose the type of malware that was used, but the breach spawned discussions about POS attacks and a long-standing weakness in the Payment Card Industry's Data Security Standard (PCI DSS).

This weakness has since been addressed by updated PCI DSS standards, but the standards are voluntary and not mandated by law. As of October 2015, all businesses are liable for credit card fraud that results from a action at any of their locations, unless an EMV chip-enabled reader is present. But the EMV chip is only used in brick-and-mortar stores and does not apply to online transactions. As it traverses various systems, POS malware searches for any weakness across the lifetime of a transaction. In the UK, where EMV chips have been used since 2003, in-person fraud has decreased but online fraud and other types of fraud have increased.

Eric Merritt, a researcher at Trustwave, discovered evidence of widespread malware targeting all sorts of POS retailers that may have existed undetected since 2011. It's called *Cherry Picker*. "Cherry Picker knows what it wants – and if it can't find it on the system, it simply exits," Merritt wrote of the technique in a blog entry referenced by Chris Brook in a *Threat Post* article posted on November 13, 2015: "This implies that the malware author *already scouted the system* and knows exactly what process they are targeting." See https://threatpost.com/researchers-discover-two-new-strains-of-pos-malware/115350/.

According to an article written by Eduard Kovacs, "'Cherry Picker' – PoS Malware Cleans Up After Itself" in *SecurityWeek*, dated November 12, 2015 (see http://www.securityweek.com/cherry-picker-pos-malware-cleans-after-itself), Cherry Picker not only knows what it wants, it knows how to cover its tracks. "Cherry Picker relies on a new memory scraping algorithm," the article states. "It uses a file infector for persistence, and it comes with a cleaner component that removes all traces of the infection from the system."

Malware Targeting City Infrastructure

In December of 2015, just before the holidays, 80,000 customers in Ukraine experienced a 6-hour power outage. Experts widely describe the incident as the first known power outage caused by a cyberattack. "It's the major scenario we've all been concerned about for so long," said John Hultquist, head of iSIGHT's Cyber Espionage Intelligence Practice, in an article written by Dan Goodin in *Ars Technica* on January 6, 2016. See http://arstechnica.com/security/2016/01/first-known-hacker-caused-power-outage-signals-troubling-escalation/.

The outage is still being investigated, but a variant of a well-known Trojan horse called *Black Energy* was found on computers at the power plant. "The Black Energy trojan, together with an SSH backdoor and the destructive KillDisk component, which were all detected in several electricity distribution companies in Ukraine, are a dangerous set of malicious tools theoretically capable of giving attackers

remote access to a company's network, shutting down critical systems and, by wiping their data, making it harder to get them up and running again," said Robert Lipovsky, a senior malware researcher at ESET, in an article posted on January 11, 2016 on the *welivesecurity* website. See http://www.welivesecu-rity.com/2016/01/11/blackenergy-and-the-ukrainian-power-outage-what-we-really-know/.

As security experts continue to investigate, most agree that malware may not have directly caused the outage, but it was involved. "A new study of a cyberattack last month against Ukrainian power companies suggests malware didn't directly cause the outages that affected at least 80,000 customers. Instead, the malware provided a foothold for key access to networks that allowed the hackers to then open circuit breakers that cut power," wrote Jeremy Kirk from IDG News Service, in *PCWorld* on January 10, 2016. According to the article, malware alone did not cause the Ukraine power station outage: "They also conducted denial-of-service attacks on the utilities' phone systems to block complaints from affected customers." See http://www.pcworld.com/article/3020631/malware-alone-didnt-cause-ukraine-power-station-outage.html.

And in January 2016 a new wave of malware, similar to BlackEnergy, continued to target Ukraine's power grid, reinforcing the fact that as governmental infrastructures become more connected to the Internet, they also become more vulnerable.

Malware Targeting the Banking Sector

Malware has sometimes been stopped, only to evolve and reappear again. For example, a malware program called *Dridex* targeted the banking industry throughout 2015. Dridex would arrive on a system as an email with a Word attachment. When the user opened the attachment, a macro embedded in the document executed and triggered the download of Dridex onto the system.

Once the malware had infected an unsuspecting host, it would direct the victim's HTTP requests to a fake bank URL watching for authentication information that could be used on the real bank website.

"The technique, known as DNS cache poisoning, involves changing DNS settings to direct someone asking for a legitimate banking website to a fake site. DNS cache poisoning is a powerful attack. Even if a person types in the correct domain name for a bank, the fake website is still shown in the browser," wrote Jeremy Kirk in *PCWorld* on January 19, 2016 in the article "Dridex Banking Malware Adds a New Trick." See http://www.pcworld.com/article/3024247/dridex-banking-malware-adds-a-new-trick.html.

Dridex spread quickly and managed to achieve a high infection rate. But even after its discovery, Dridex still continues to evolve using its botnets to spread even further into networks of compromised hosts.

Ransomware – Malware Used in Extortion Schemes

In a paper entitled "2016 Threats Predictions," (http://www.mcafee.com/us/resources/reports/rp-threats-predictions-2016.pdf) McAfee Labs stated the following: "Ransomware will remain a major and rapidly growing threat in 2016. With upcoming new variants and the success of the *ransomware-as-a-service* business model, we predict that the rise of ransomware that started in the third quarter of 2014 will continue in 2016."

"On Monday, Nov. 24 (2014), a crushing cyberattack was launched on Sony Pictures. Employees logging on to its network were met with the sound of gunfire, scrolling threats, and the menacing image of a fiery skeleton looming over the tiny zombified heads of the studio's top two executives. Before Sony's IT staff could pull the plug, the hackers' malware had leaped from machine to machine throughout the lot and across continents, wiping out half of Sony's global network," writes Peter Elkind in the July 1, 2015 issue of *Fortune* in the article "Inside the Hack of the Century." See http://fortune.com/sony-hack-part-1/ for details on the hackers suspected of launching the attack and what their demands included.

At the RSA Conference in 2015, Stuart McClure, CEO of the computer security firm Cylance, spoke about how the Sony hack took place, explaining it as a combination of phishing emails, weak passwords, and a lack of server hardening.

The initial email, received by several Sony executives, contained fake Apple ID verification requests with a link to a fake domain that prompted them to enter their Apple ID and password information. Once the information was entered, the attackers took these credentials and coded them into a strain of malware known as *Wiper* in hopes that the Sony executives were using the same verification information on their corporate accounts. Apparently, some of them were; the attackers eventually crippled Sony's networks. The breach is estimated to have cost Sony upward of $171 million.

And, in February of 2016, Hollywood Presbyterian Medical Center paid a $17,000 Bitcoin ransom to hackers who used similar malware to take control of the medical center's computer systems and would not relinquish control until they were paid.

"The malware locks systems by encrypting files and demanding ransom to obtain the decryption key. The quickest and most efficient

way to restore our systems and administrative functions was to pay the ransom and obtain the decryption key," said hospital Chief Executive Allen Stefanek, in a *Los Angeles Times* article published on February 18, 2016. See http://www.latimes.com/business/technology/la-me-ln-hollywood-hospital-bitcoin-20160217-story.html.

The FBI is investigating the attack.

The Juniper Networks Solution to Advanced Threat Prevention

As malware evolves and its attacks become more specialized and highly targeted, a new category of advanced security has also emerged that can detect, analyze, and prevent these advanced threats, which are able to bypass traditional security methods.

Juniper Networks' solution for preventing advanced and emerging threats is Sky Advanced Threat Prevention (ATP), a cloud-based anti-malware solution coupled with the SRX Series firewall. Let's drill down and see how this unique combination makes Sky ATP so effective.

Sky ATP

Sky ATP is an add-on for the SRX Series. It provides anti-malware prevention for existing and new SRX Series customers as shown in Figure 2.1. In addition, Sky ATP includes malware detection and analysis, host analyzer, and command and control feeds. Each component in the solution has a role in detecting, analyzing, and blocking malware, but only the actual SRX Series device has a footprint in your network. All other components act as cloud-based services.

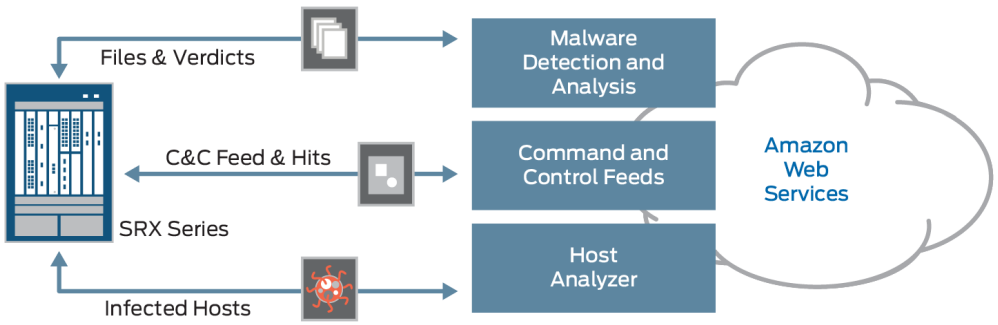

Figure 2.1 Sky ATP Solution Overview

Sky ATP's cloud-based design delivers protection against Day Zero threats as it analyzes ingress and egress traffic for malware and indicators of compromise. It can instantly provide deep inspection, actionable reporting, and inline malware blocking. Sky ATP's solution components are listed in Table 2.1.

Table 2.1 Sky ATP Solution Components

Malware Detection and Analysis	Deployed in the cloud.
	Serves as the inspection pipeline, performs malware detonation, and provides the logging infrastructure.
	Returns verdicts and provides analytics.
Command and Control Feeds	Provides cloud-delivered security intelligence, specifically Juniper's command and control (C&C) feed.
	Accepts C&C detections.
Host Analyzer	Correlates C&C detections with malware detections to identify compromised hosts.
	Provides a feed of compromised hosts to the SRX Series for quarantine.
SRX Series Firewall	Extracts suspicious content and sends samples to the Sky ATP service for analysis.
	Performs inline blocking based on verdicts from the Sky ATP service.
	Leverages C&C feeds and sends detections to the Sky ATP service.
	Sends collected data to the Sky ATP service for reporting and telemetry purposes.

Analyzing and Detecting Malware

Sky ATP detects malware by using an *analysis pipeline*, as shown in Figure 2.2, when files are sent to the Sky ATP service by the SRX Series device:

- Cache lookup: Determines if the file in question is a known bad file.
- Antivirus (AV) scanning: Runs the file through several well-known AV scanners.
- Static analysis: Checks the file for suspicious signs such as unusual instructions or structure.

- Dynamic analysis: Executes the file in a real environment to see what it does in a secure test bed. This is the most thorough method of analysis, and it is used when the other methods have flagged a file as suspicious.

The analysis pipeline assigns values to each step of the process: these values are combined to provide a progressively more accurate assessment.

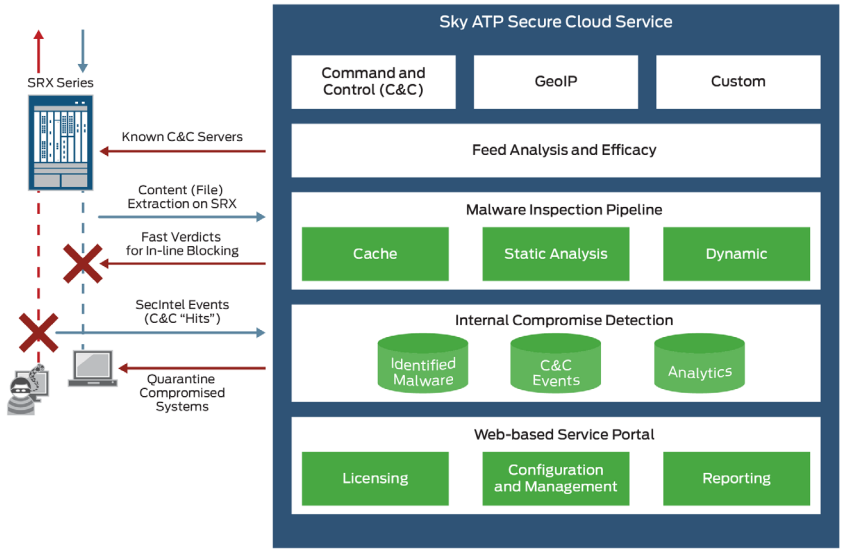

Figure 2.2 The Sky ATP Analysis Pipeline

Methods of Detection

Third-Party Antivirus Scanners

Antivirus (AV) scanners continue to play a role in detecting known viruses and malware. Although their reliance on signature matching provides a low Day Zero detection rate, by day one of a virus's appearance that detection rate increases to 55 percent. For viruses that have been in the wild for at least three months, the successful detection rate averages 75 percent. In order to improve the results of AV analysis, Sky ATP combines multiple AV scanners, and users can change which scanners they use at any time.

Machine Learning

Sky ATP uses a proprietary implementation of machine learning as another method of analysis that recognizes patterns and then correlates the information. The machine-learning algorithm is trained with features from thousands of malware samples as well as thousands of

samples. Sky ATP learns how malware acts and is regularly retrained to get smarter as threats evolve.

Static Analysis

Static analysis investigates suspect files for known information about their type or source. For example, the source URL from which the file originates will be investigated along with the file itself, which is broken down into specific features such as file structure, meta information, category of instruction used, and file entropy. Then each feature is fed into the Sky ATP's machine-learning algorithm and the technology improves itself.

Using Dynamic Analysis

Dynamic analysis sandboxes suspect files and executes them in a real environment where they can run uninterrupted for minutes. During that time, active deception encourages the malware to show itself, and a record of its activity is kept. The file is then fed into the Sky ATP machine-learning algorithms.

Dynamic Analysis Deception Techniques

Within the Sky ATP sandbox environment, various methods are used to draw out the malware and provoke it into action. So, for example, the sandbox must emulate a user environment with realistic patterns of user interaction, and high-value targets must appear with vulnerabilities so that the malware is sufficiently provoked. These targets could be stored credentials, vulnerable software, or tempting user files.

Because some malware is not so easily fooled, and will wait for specific signs that a real user is sitting at a computer before it shows itself, it's especially important that effective and realistic user actions, such as the following, are simulated within the sandbox:

- Faking a webcam feed
- Faking a microphone feed
- Moving the mouse
- Simulating key strokes
- Operating dialog boxes
- Installing and launching software
- Adding and removing USB sources

Once the malware exercises its payload, all actions taken by it are detected, analyzed, and documented.

The Sky ATP Dashboard and Web UI

The Web-based user interface for Sky ATP includes a dashboard that provides a visual summary of all information gathered on compromised content and hosts in real time (see Figure 2.3). The Web UI for Sky ATP cloud-based services (see Figure 2.4) allows you to customize information in various filtered windows, and track the devices in your network, as well as create specialized white lists and black lists, among many other features. For a complete list of features, visit the links at the end of this chapter.

As shown in Figures 2.3 and 2.4, the Sky ATP dashboard can display, among other features:

- The top infected hosts, including IP address, domain, and threat level.
- A file scanning summary with percentages of blocked malware, Day Zero malware, unknown files, and clean files for a chosen amount of time (for example, one week).
- A summary of top users downloading malware, which provides a list of those users, including the number of downloads they make and their email addresses.
- A list of infected file types (for example, PDF, XLS, DOC, CSV, RTF, and others) with a graphical comparison of the number of downloaded files for each type that was blocked.
- A map of the world with shaded areas displaying countries with the highest number of C&C servers and malware sources.

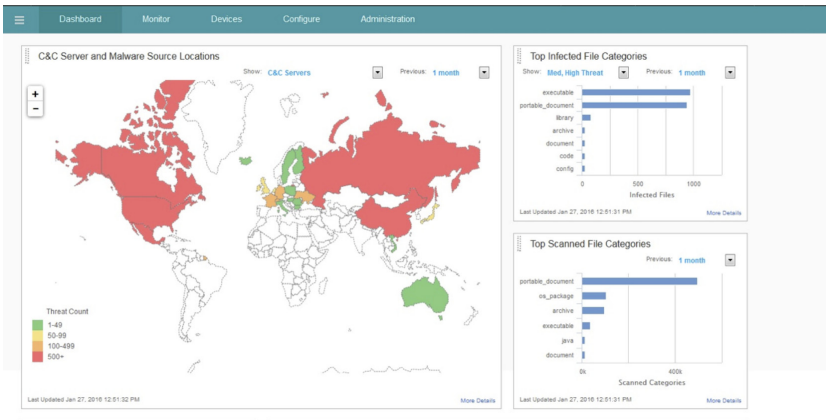

Figure 2.3 Sky ATP Dashboard

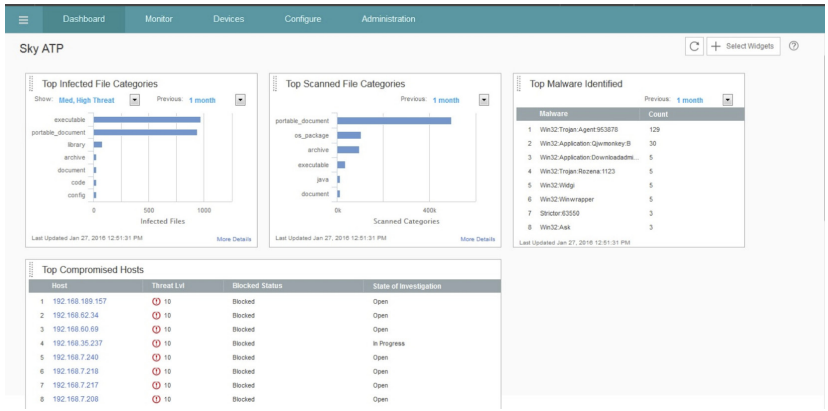

Figure 2.4 Sky ATP Web UI

Resources and References

Start here at the Sky ATP product page on the Juniper Networks website:
http://www.juniper.net/us/en/products-services/security/sky-advanced-threat-prevention/

The datasheet for Sky ATP provides a great product overview:
http://www.juniper.net/assets/us/en/local/pdf/datasheets/1000549-en.pdf

Sky ATP is an add-on service to the SRX Series:
http://www.juniper.net/us/en/products-services/security/srx-series/

Learn more about security intelligence feeds used by Sky ATP:
http://www.juniper.net/us/en/products-services/security/spotlight/

Read about the Hyatt breach in detail:
http://www.esecurityplanet.com/network-security/hyatt-breach-affected-250-hotels-worldwide.html

Read why consumers should be angry at negligent retailers:
https://www.securestate.com/blog/2014/01/14/why-chip-and-pin-isnt-the-answer-to-retailers-problems

Read about Cherry Picker malware in detail:
http://www.securityweek.com/cherry-picker-pos-malware-cleans-after-itself

https://threatpost.com/researchers-discover-two-new-strains-of-pos-malware/115350/

Chapter 3

Cloud-Enabled Enterprises

The cloud is central to transforming enterprise networks because it offers a springboard to flexibility and agility hitherto unavailable to administrators managing smaller networks. When smaller networks adopt cloud technologies as a primary operating model, administrators need a secure on-ramp to cloud-based applications deployed in their private clouds or their on-premise data centers, or hosted in their remote locations. Up until now, this situation has created uncertainty surrounding cloud security and inhibited businesses ready to adopt cloud-enabled networking. That is why securing the cloud-enabled enterprise has been one of the top priorities at Juniper Networks. This chapter explains Juniper's solutions to help you secure cloud-based enterprise networking.

What Is a Cloud-Enabled Enterprise?

A cloud-enabled enterprise network has a common, unified infrastructure that supports a diverse set of devices, applications, and people. It provides reliable, scalable, secure, and readily available access to resources. So it does not matter where these data points reside – in the cloud, the data center, or the WAN. They are all unified in one infrastructure.

Back in 2012, cloud technology was identified as one of the fastest moving technologies by Gartner Inc.'s Hype Cycle for Emerging Technologies (http://www.gartner.com/newsroom/id/2124315). Today, Verizon notes that nearly 90% of enterprise networks already use the cloud for mission-critical workloads, and that half of that number will use cloud models for at least three-quarters of their workloads by 2018 (Source: Verizon Enterprise Solutions' State of the Market: Enterprise Cloud 2016 report – http://www.verizonenterprise.com/resources/reports/rp_state-of-the-market-enterprise-cloud-2016_en_xg.pdf).

Cloud architectures help enterprises increase productivity and competitiveness through virtualization and automation technologies that simplify network life-cycle management while delivering operational efficiencies. They provide a chance for an enterprise network to:

- Reduce CapEx and OpEx
- Simplify centralized management
- Scale at will
- Create business agility
- Support new business models
- Collaborate more effectively

The good news is that administrators are excited by opportunities to reduce capital costs and the chance to divest themselves of infrastructure management. The bad news is that numerous challenges must be addressed before an enterprise network can become cloud-based. Some of the many challenges businesses face when incorporating cloud architecture are:

- Static and inflexible multilayer architectures
- Operational complexity
- Lack of regulatory compliance and standards
- Lack of control over data and performance
- Security, which is the most immediate challenge

Let's examine the security challenges in a cloud-enabled enterprise.

Security Challenges in Cloud-Enabled Enterprises

Enterprises are migrating toward cloud-enabled networks much faster than expected and this trend has created an immediate need to extend the level of security found in their *traditional* networks to the new cloud landscape.

One of the biggest security concerns administrators face when moving information over to the cloud is the perception that they will lose control of their data. Other concerns include identity and access management, Web application security, browser security, securing data in transit (to the cloud and back), data leaks, compliance issues, remote access for users, securing endpoints, malware, mobile BYOD, network viruses, and more.

The risks of cloud networking are real, and they include:

- Data breaches that can result from a flaw in a cloud application design or other application vulnerability.
- Data loss as a result of a malicious attack, an accidental deletion, or a physical problem in the data center.
- Account hijacking, including the use of phishing, fraud, or social engineering to obtain a user's private login information.
- Insecure interfaces, such as the application programming interfaces (APIs) used by cloud services to provide authentication, access control, encryption, and other key functions. Vulnerabilities in these interfaces can increase the risk of a security breach.
- Denial-of-service (DoS) attacks, in which a malicious actor prevents users from accessing a targeted application or database.
- Malicious insiders, such as employees or contractors, who use their position to gain access to private information stored in the cloud.
- Abuse of services, which involves hackers who use the limitless resources of the cloud to crack an encryption key, stage a distributed denial-of-service (DDoS) attack, or perform other activities that would not be possible with limited hardware.
- Shared vulnerabilities, including platforms or applications accessed by different users in a multitenant environment.

Thus, while cloud-enabled networking offers many benefits to a business enterprise, failing to properly anticipate the security risks of working in the cloud, or even rushing the migration process, can expose organizations to considerable risk.

Remember that much of enterprise networking has traditionally relied on the perimeter safety net, and now that the perimeter has shifted to the cloud, your network security posture must change. For example, in a cloud-enabled enterprise, you need to adopt security models suitable to cloud environments and consider end-to-end embedded security models, as opposed to traditional perimeter-based models. Here are some of the critical security challenges enterprises need to address while shifting to a cloud-enabled environment:

- *Protecting against threats* – You need to monitor and inspect packets for the presence of malware or malicious traffic and send appropriate alerts to cloud service managers, as well as to

subscribing customers. You need to use intrusion detection and prevention (IDP) systems. Also, you must have virtualization-specific antivirus protections in place that can deliver on-demand and on-access scanning of virtual machine (VM) disks and files (with the capability to quarantine infected entities that have been identified).

- *Protecting data in storage and data in transit* – To protect confidential business, government, or regulatory data, you must ensure that the data remains secure both at rest – data residing on storage media – as well as when in transit. In addition, you need to address confidentiality, integrity, and authentication concerns to ensure that data is secure, accessible by the correct recipients, and not modified in transit.
- End-to-end visibility – To secure your network, you need to have end-to-end visibility so that you can first understand it, and then be able to program it, in order to defend against threats. And as your network becomes more virtualized, having visibility and control across both your physical *and* virtual infrastructures becomes critical.
- Elasticity – You need to be able to scale your resources up and down in real time, readily, and cost-effectively. The challenge here is to provide detailed access and predefined security controls without compromising on your network's performance, security, or stability.
- *Authentication of users, applications, and processes* – You need hardened access control procedures to keep unauthorized access at bay while allowing authorized users in with a minimum of fuss. In any cloud environment, authentication and access control are critical because the cloud, with all of its data, is potentially accessible to anyone on the Internet.
- *Incident response* – You need a dynamic anti-malware solution that can perform incident response activities such as detecting and blocking malware by monitoring ingress and egress network traffic, looking for malware and other indicators of compromise, reporting incidents, performing dynamic attack analysis, collecting and preserving data, and ensuring remediation of problems and restoration of service.
- *APIs and orchestration* – Security tasks in the cloud typically require a number of APIs to accomplish tasks. For example, when a new VM instance is created, a security policy is needed to allow traffic flow to that instance.

- *Security intelligence* – You need to increase your focus on detecting attacks and attackers early, instead of at the point of breach, and you need to share security intelligence across all your security architectures (in contrast to perimeter products that function without information sharing). Better security intelligence enables you to better leverage the live data feed that is, in turn, used for policy enforcement. Time is of the essence when protecting your critical data from being compromised.
- *Multitenancy management and isolation enforcement* – Your cloud provider (or your cloud team) must enforce the isolation of customer-specific traffic, data, and resources because one tenant cannot impose risks on others in terms of data loss, misuse, or privacy violations. Also, by isolating areas of the virtualized network into security zones, you limit the scope of exploits and facilitate consistent policy enforcement throughout the cloud. Ask for VMs and hypervisors in the cloud as they can provide stronger network separation and security.
- *Cloud legal and regulatory issues* – You must have strong policies and practices as part of your business objectives to address any legal and regulatory issues that might arise. Your cloud or network administrators must be able to see their aggregate compliance posture at a glance. And if a noncompliance alert is triggered, they must be able to drill down within servers or VMs to identify the exact condition that caused the alert.

Cloud-enabled enterprise networks bring a whole new level of virtualization, agility, management, and security issues; but if implemented correctly, they also provide a new simplicity to enterprise campus and branch networking.

Juniper Networks Unite

Juniper Networks has introduced an innovative reference networking architecture called *Juniper Unite*, which comprises enterprise switching, security, routing, and software technologies. Juniper Unite also contains third-party WLAN, unified communications and collaboration (UCC), and network access solutions supported through the Open Convergence Framework (OCF) as shown in Figure 3.1.

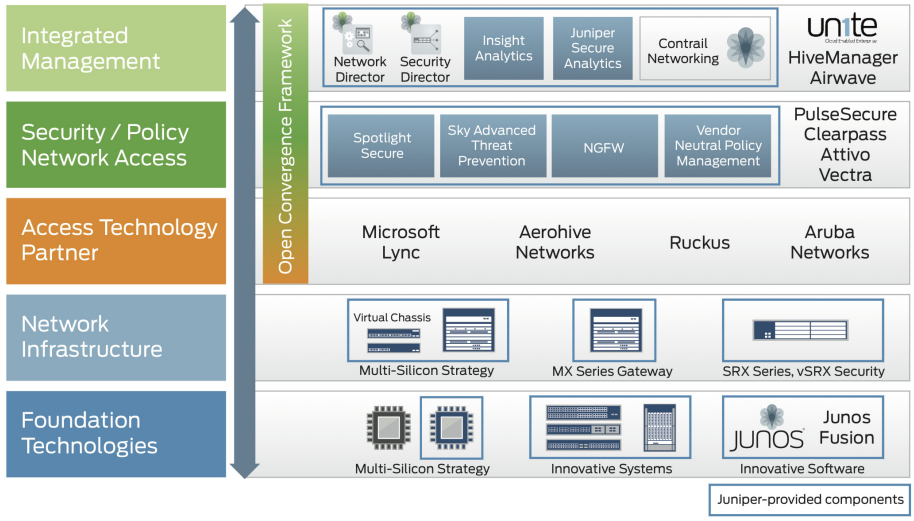

Figure 3.1 Juniper Unite

The Juniper Networks Unite enterprise architecture provides a simplified, automated, and secure design to build and manage cloud-enabled enterprises. Based on Juniper's switching and security solutions, it provides UTM, next-generation firewalls, and malware detection and eradication tools. Let's focus on Juniper's security solutions, as shown in Table 3.1, for the remainder of this chapter.

Table 3.1 Juniper's Cloud-Enabled Enterprise Security Solution

Cloud-Enabled Requirement	Juniper Solution	Description
Firewall	SRX Series	• Protects against application-borne security threats and manages bandwidth usage. • Offers a range of user role-based firewall control solutions that support dynamic security policies. • Secures against network-based exploit attacks aimed at application vulnerabilities. • Protects against malware, viruses, phishing attacks, intrusions, spam, and other threats through antivirus and antispam, as well as by means of Web filtering and content filtering. • Offers multiple deployment options – appliance, chassis-based, and virtual – all offering the same full-feature capability. • Provides a foundational platform for open policy enforcement.

Virtual Firewall	vSRX	• Provides full visibility and granular access control over all traffic flowing through VMs. • Acts as a barrier to secure perimeter access to a network. • Provides dedicated security services and assured traffic isolation within the cloud, offering customizable firewall controls as an additional managed service. • Enforces a granular virtualized security policy consistent with and integrated into physical server security, including Juniper Networks SRX Series Services Gateways and Juniper Networks Secure Analytics.
Threat Prevention	Sky ATP	• Delivers protection against sophisticated *zero-day* threats. • Inspects ingress/egress traffic for malware and indicators of compromise. • Delivers deep inspection, actionable reporting, and inline malware blocking. • Integrates with SRX Series and enables detection and prevention of threats. • Defends against sophisticated malware leveraging deception techniques. • Outputs actionable (compromised host quarantine) rich reporting. • Augments sandboxing (with additional analysis to detect evasive malware).
Threat Intelligence	Spotlight Secure	• Uses threat detection and advanced malware prevention in an open-platform environment. • Connects multiple intelligence sources to the enforcement points. • Provides a single point of administrative control. • Supports policy engine controls for prioritization and categorization of threats.
Management	Junos Space Security Director	• Delivers scalable security management. • Provides security policy management through an intuitive, centralized, Web-based interface that offers enforcement across emerging and traditional risk vectors. • Enables easy management of all phases of the security policy life cycle through a single Web interface. • Offers extensive security scale, granular policy control, and policy breadth across the network.
Analytics	Secure Analytics	• Combines network performance and security management into a single enterprise and service provider solution that integrates log, threat, and compliance management for both Juniper and non-Juniper products. • Delivers a complete picture of network visibility and access, making troubleshooting and cloud security optimization easier and faster.

To better understand how Juniper's cloud-enabled enterprise security is relevant, let's walk through a use case for enterprise expansion and workload distribution that shows how Juniper's Unite security solution can address both of these common concerns.

Use Case: Securing and Simplifying an Enterprise Network

This use case is an example of a hybrid cloud enterprise network that combines aspects of both private (internal) and public (external) cloud computing environments. In this use case, administrators typically outsource noncritical information and processing to the public cloud, while keeping business-critical services and data within their control.

The issue with a hybrid cloud is that some processes remain on premise and some move to the cloud. It is an architectural choice involving new and extended security processes that typically require multiple hybrid-cloud administrators to collaborate in order to achieve comprehensive data protection.

However, with the Juniper Unite architecture, the solution is based on high-performance threat detection and prevention technologies as shown in Figure 3.2. It is a design for the most effective security detection and remediation for existing and future hybrid cloud applications.

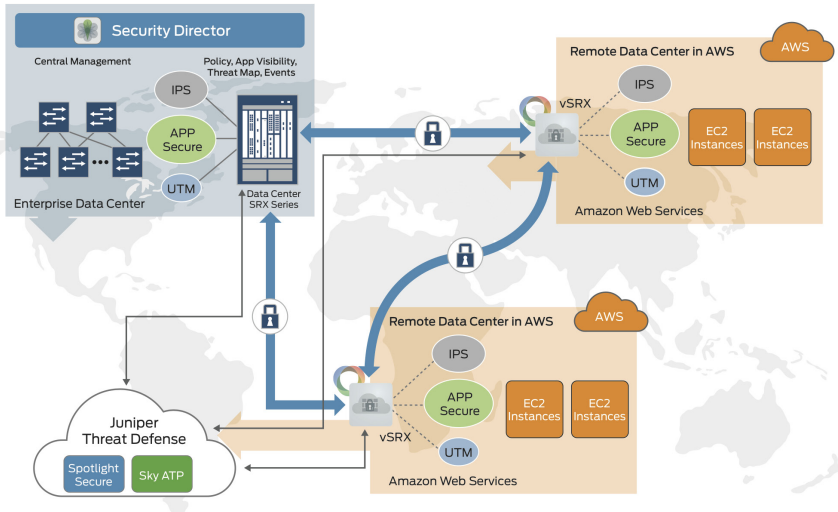

Figure 3.2 Juniper Hybrid Cloud

In Figure 3.2, the SRX Series and vSRX firewalls secure the network with a suite of cloud-based threat defense services and feeds that centrally defend all locations. By leveraging security feeds from cloud-based advance threat intelligence platforms, such as Sky ATP, the SRX Series can detect known and zero-day threats while enforcing application security, intrusion prevention, and UTM.

For large-scale deployments, the Juniper Unite design enables APIs for cloud management automation, which in turn enables rapid provisioning and elastic scalability. With a wide range of programmatic APIs supported on Junos OS, in-house or for-hire DevOps resources can easily automate deployment and management activities through simple scripts, streamlining the entire workflow.

Figure 3.2 shows that Junos Space Security Director centrally manages all security policies across the infrastructure. The vSRX virtual firewalls deployed in the remote data centers register with Security Director, whether installed at headquarters or in the cloud. New security policies are then centrally added or updated from Security Director and deployed across all data centers as the vSRX provides data isolation of cloud tenant resources by shrink-wrapping each customer VM, or group of VMs, in distinct security policies. (A mobile Security Director app is accessible to security administrators or CIOs who want to monitor security updates in their network remotely.)

The vSRX virtual firewalls in the remote data center branches connect to the SRX Series firewalls in the head office through IPsec VPN for secure data transportation. Additionally, the introduction of Juniper Secure Analytics into the environment provides a comprehensive suite of tools for monitoring all security, networking, and server infrastructure.

Finally, the ability to extend familiar and well-known security policies once used in the physical data center to both private and public clouds is a critical benefit that enables IT managers to leverage existing enterprise personnel to manage their cloud and hybrid-cloud infrastructures.

Resources and References

Links to relevant resources for the cloud-enabled enterprise:

- http://www.juniper.net/assets/us/en/local/pdf/solutionbriefs/3510567-en.pdf
- http://www.juniper.net/assets/us/en/local/pdf/solutionbriefs/3510558-en.pdf
- https://www.juniper.net/us/en/local/pdf/whitepapers/2000465-en.pdf
- http://newsroom.juniper.net/press-release/juniper-networks-unveils-latest-security-innovations
- http://newsroom.juniper.net/press-release/juniper-networks-introduces-juniper-unite-for-the-cloud-enabled-enterprise

Trusted Computing Group's White Paper, 2010: http://www.trustedcomputinggroup.org/cloud-computing-security-natural-match/.

Security and Privacy Issues in Cloud Computing by Jaydip Sen (Innovation Labs, Tata Consultancy Services Ltd., Kolkata, INDIA): https://www.safaribooksonline.com/library/view/architectures-and-protocols/9781466645141/978-1-4666-4514-1.ch001.xhtml.

Chapter 4

Application Visibility and Control

This chapter explains how application visibility and control functionality plays an important role in protecting critical high-speed networks against the penetration and proliferation of existing and emerging application-level threats by identifying the applications traversing the network using application identification.

It introduces you to the basics of Juniper Networks AppSecure security suite, reviews each of the core AppSecure services, and explains why those services are becoming increasingly necessary for successful businesses.

The Application Landscape

The rapidly evolving world of contemporary business has benefited from many technologies that were unheard of just 10 years ago. But every new technology brings new security challenges, and with the huge numbers of users, devices, and data being deployed to take advantage of the latest technologies, enterprises are becoming increasingly vulnerable to data loss, malicious attacks, and network instability. Let's begin by discussing some security challenges related to the evolution of applications.

Web-based applications have changed the dynamics of security. In the past, specific applications were associated with specific protocols and ports, and setting and enforcing policies at the host level was relatively straightforward. Now, given the reliance on Web applications, virtually all traffic is HTTP-based (ports 80/443) as shown in Figure 4.1.

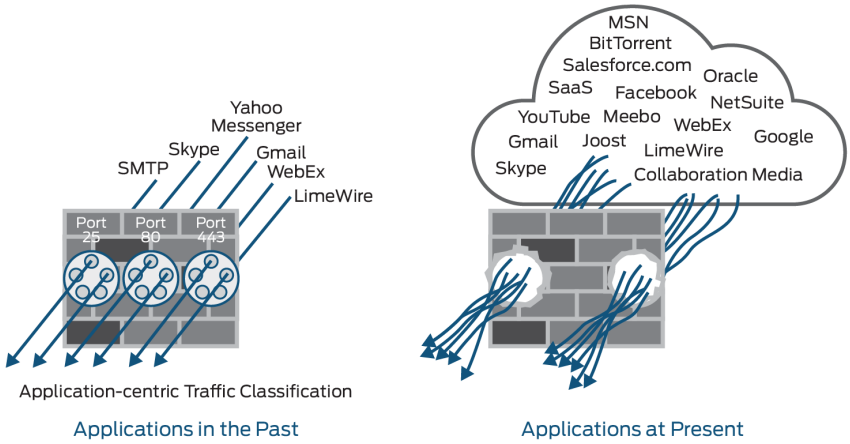

Figure 4.1 Applications Landscape – Past and Present

Use of nonstandard ports and encryption are two of the means by which applications have become more accessible, but cyberattackers implement the same technology to create cyberthreats or hide those threats within the application traffic itself.

Consequently, network security solutions operating solely on basic Internet Protocol (IP) layer information are unable to distinguish between permitted and malicious activity. Further, the very advantage of Web applications—the fact that they can be accessed from anywhere by employees, contractors, partners, and service providers through the firewall—creates its own set of access control challenges.

User-centric applications designed primarily for personal communications, such as instant messaging, peer-to-peer file sharing, Webmail, social networking, and IP voice/video collaboration, present a specific set of challenges as many of these applications evade traditional security mechanisms by dynamically changing their communications ports and protocols, or by tunneling within other commonly used services (for example, HTTP or HTTPS).

The concept that any application can be used on any port is one of the fundamental changes in the application landscape, driving the migration from port-based firewalls to next-generation firewalls.

Yet another challenge is maintaining core business applications, as they are heavily targeted by cyberattackers using multifaceted attacks. Organizations need more control over the applications and traffic on their networks to simultaneously protect their assets against attacks and manage bandwidth use. An effective security solution needs to deliver the right security services in order to provide administrators with visibility and control over the applications traversing their networks.

Requirement for Application Visibility and Control

In response to major industry trends such as mobility and virtualization, applications are increasingly appearing in a dynamic environment that includes mobile devices, mobile apps, hosted virtual desktops, and hybrid clouds. Since users coming in from a variety of media must be accounted for, achieving effective application visibility becomes a challenge.

Businesses are often left vulnerable to threats, or rendered unable to respond to threats, by the complexity caused by voice, data, video, and applications running on the same network. So it is essential that the network be aware of each application traffic type and provide the appropriate priority, routing, and bandwidth required to ensure the maximum user quality of experience. Factors adding to these complexities include the following:

- Applications are often highly extensible, and often include features that may introduce unwarranted risk. Such applications represent both business and security risks and your challenge will be to determine how to strike an appropriate balance between blocking some and securely enabling others.
- Converged solutions, such as peer-to-peer applications, are also driving new traffic patterns that are foreign to the way today's networks were provisioned. Ensuring that networks are application-aware enables them to flexibly adapt to new applications and traffic patterns as they emerge. In addition to the highly publicized legal concerns surrounding peer-to-peer file sharing applications, these applications can rob network bandwidth and leave the majority of users with a poor network and application experience.
- Use of nonstandard ports, for example when a Web server is running on a port other than those commonly associated with HTTP (that is, 80 and 443), applications can set up sessions at nearly 5000 other ports. Many applications that use SSL never use port 443, nor do they use SSL-defined ports.

Application visibility and control is necessary in order to:

- Identify applications, and allow, block, or limit applications – regardless of the port, protocol, decryption, or any other evasive tactic.
- Identify users, regardless of device or IP address, by using granular control of applications by specific users, groups of users, and machines that the users are operating. This helps organizations control not only the types of traffic allowed to enter and exit the network, but also what a specific user is permitted to send and receive.

- Support all inbound and outbound SSL decryption capabilities. This includes recognition and decryption of SSL on any port, inbound and outbound, policy control over decryption, and the necessary features to perform SSL decryption across tens of thousands of simultaneous SSL connections. This helps an organization identify and prevent threats and malware in encrypted network streams.
- Integrate with IPS so that it applies appropriate attack objects to applications running on nonstandard ports. Application identification improves intrusion detection and prevention (IDP) performance by narrowing the scope of attack signatures for applications without decoders. It can be activated and deactivated as required.
- Support the exact same firewall functions in both a hardware and virtualized form factor.
- Scan all applications on all ports in real time for viruses and malware to protect against known and unknown threats embedded across applications.

How Does Application Visibility and Control Work?

A growing number of applications running over HTTP, including video and even hosted applications, are causing strain on the network, leading to higher infrastructure costs and making the network more difficult to manage. Also, the increase in complexity of applications is making it more difficult for network administrators to optimize the performance of these applications running on thier network.

It is no longer effective to block or allow TCP and/or UDP ports, since most applications do not map to individual ports. For example, controlling traffic on an HTTP or HTTPS port is ineffective against complex social networking sites and cloud applications.

To overcome this problem, you need to use application identification (App ID) as the primary classification engine and then add an application signature pattern-matching engine that operates at Layer 7 and inspects the actual content of the payload for identifying applications.

App ID performs a deep packet inspection (DPI) of traffic on the network and on every packet in the flow that passes through the application identification engine until the application is identified. Application findings such as IP addresses, hostnames, and port ranges are saved in the application system cache (ASC) to expedite future identification.

Let's review the different mechanisms used by App ID to identify traffic.

Application Signatures

Context-based signatures are used to first look for unique application properties and related transaction characteristics to correctly identify the application regardless of the protocol and port being used. Application signature mapping, or signature mapping, is a precise method of identifying the application that generated traffic on the network. Signature mapping operates at Layer 7 and inspects the actual content of the payload. The payload of the first few packets is compared to the content of the database. If the payload contains the same patterns as an entry in the database, the application related to the traffic is identified as the application mapped to that pattern in the database entry.

Heuristic Detection

Evasive applications such as peer-to-peer applications do not provide any obvious patterns for matching. Heuristics detection looks at the behavior of the traffic in an analytical fashion to detect what application is running; it can examine the byte stream to determine if it is encrypted by measuring the randomness of the payload bytes. Any application stream that is encrypted (or compressed) will exhibit a highly randomized byte stream.

Alternate Mapping Techniques

In some cases, an alternative method of identifying an application might be required. For example, if traffic on a network is generated by an internal proprietary application, a predefined application signature will not exist. Application identification will identify the application of the traffic as unknown. To keep this traffic from being handled as unknown, Layer 3 or Layer 4 information specific to this application can be mapped to the application name, overriding the application identification process.

Custom Signatures

User-defined custom application signatures can also be used. Custom application signatures are unique to your environment and are not part of the predefined application package. You can create custom signatures using hostnames, IP address ranges, and ports, which allows you to track traffic to specific destinations.

SSL Inspection

Application identification detects encrypted applications based on the patterns in HTTP contexts. However, some HTTP sessions are encrypted in Transport Layer Security (TLS) and Secure Sockets Layer

(SSL). The SSL inspection feature is used for identifying applications that use HTTP over SSL/TLS or HTTP. If App ID determines that SSL encryption is in use, the traffic is decrypted and then passed to other identification mechanisms as needed. An SSL proxy must be enabled for application identification of HTTPS traffic to take place.

Application Protocol Decoding

Once the application is identified, it is further decoded at protocol level, if necessary. Protocol decoders are used to check protocol integrity and protocol contextual information by looking for anomalies and ensuring that RFC standards are met. An anomaly can be any part of a protocol, such as the header, message body, or other individual fields that deviate from RFC standards for that protocol.

When protocol contextual information is available, protocol decoders check for attacks within those contexts. By using protocol contextual data, rather than the entire packet, for attack detection, protocol decoders improve overall performance and accuracy.

About Juniper Networks AppSecure

The Juniper Networks AppSecure (AppSecure) suite of application-aware security services for the SRX Series, which was born from App ID technology, classifies traffic flows, while bringing greater visibility, enforcement, control, and protection to your network security.

App ID enables you to see the applications on your network and learn how they work, their behavioral characteristics, and their relative risk. Using several different identification mechanisms, App ID detects the applications on your network regardless of the port, protocol, and encryption (TLS/SSL or SSH) or other evasive tactics used. The number and order of identification mechanisms used to identify the application will vary, depending on the application.

AppSecure uses a sophisticated classification engine to accurately identify applications regardless of port or protocol, including applications known for using evasive techniques to avoid identification. It gives you the context to regain control of your network traffic, set and enforce policies based on accurate information, and deliver the performance and scale required to address your business needs.

The Onbox Application Identification Statistics feature adds application-level statistics to the AppSecure suite. Application statistics allow an administrator to access cumulative statistics as well as statistics accumulated over user-defined intervals.

Figure 4.2 shows the different AppSecure service modules.

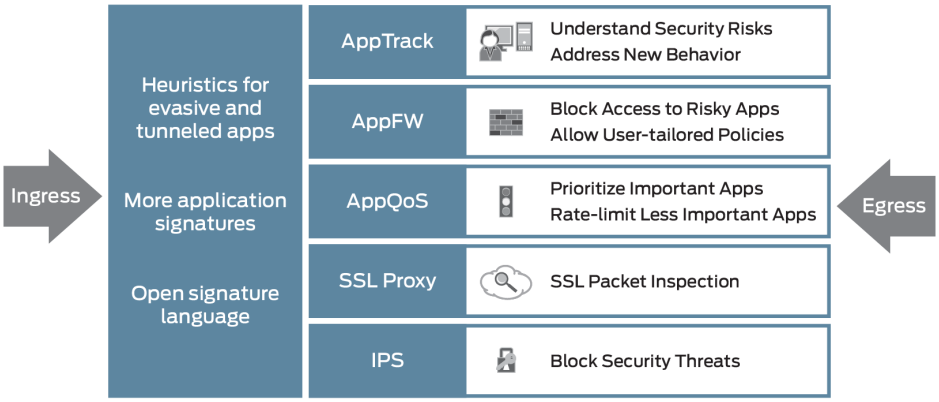

Figure 4.2 AppSecure Service Modules

The services enabled by AppSecure include AppTrack for detailed visibility of application traffic, AppFW for granular policy enforcement of application traffic, and AppQoS for prioritization and metering of application traffic. SSL proxy provides visibility of encrypted traffic to allow deep packet inspection (DPI). These modules perform various tasks on the traffic based on the result of the App ID.

AppSecure works with additional content security through integrated UTM, IPS, and Sky ATP on the SRX Series for deeper protection against malware, spam, phishing, and application exploits.

Essential Capabilities for a Complete Solution

AppSecure service modules are capable of addressing the application visibility challenges faced by many organizations by monitoring and controlling traffic for tracking, prioritization, access control, detection, and prevention, based on the application ID of the traffic.

Here are some examples of situations in which App ID is used along with AppSecure in order to solve common problems with application visibility and control.

Application Awareness and Control

Requirement (1): Increased use of cloud-based services, mobile devices, and media-rich applications puts more strain on a network, leads to higher infrastructure costs, and makes a network more difficult, and critical, to manage. The performance of your applications and business services depends on the performance of your network, so enhanced visibility into the nature and behavior of applications running on your

network is a must. Administrators need to ensure their network delivers optimal performance for the applications that matter most to their business.

Requirement (2): Internet and social media applications are a common source of vulnerabilities and attacks. Organizations have the challenge of managing or controlling a vast array of Web applications without hindering productivity. More and more, applications such as instant messaging applications, peer-to-peer file sharing, or Voice over Internet Protocol (VoIP) are capable of operating on nonstandard ports or can hop ports.

In order to enforce application-specific firewall policies, organizations need to detect all applications, regardless of the port on which they occur, by inspecting traffic to establish the true identity of applications.

Solution: App ID provides granular control over applications, including video streaming, peer-to-peer communication, social networking, and messaging, as well as identifying services, port usage, underlying technology, and behavioral characteristics within applications. The App ID module matches applications for both client-to-server and server-to-client sessions.

Once the application is determined, AppSecure service modules can be configured to monitor and control traffic for tracking, prioritization, access control, detection, and prevention based on the application ID of the traffic.

Devices enabled with AppSecure provide your network team with enhanced visibility into application behavior. As a result, administrators can gain improved control over the prioritization, routing, load balancing, and optimization of application traffic in order to improve user experience and reduce costs.

Application Control with AppQoS

Requirement: Organizations need to control access to applications not only because of the vast number of vulnerabilities they introduce to a network, but also because they consume an excessive amount of bandwidth, resulting in a considerable loss of productivity.

Sometimes simply permitting or denying traffic is not a granular enough response. For instance, you might have traffic that you do not want to explicitly block, but at the same time, you do not want to give it free rein on your network. Examples are applications that impact productivity, such as online games, or consume large amounts of bandwidth, such as peer-to-peer apps or streaming video. You might also have business-critical applications that need to be prioritized over other applications so they are not forced to deal with insufficient bandwidth, resulting in poor application control and frustrated users.

Solution: Identify and control access to specific applications so as to achieve optimal bandwidth utilization for business-critical applications.

AppQoS allows you to do this by providing the ability to invoke it on top of your firewall rule base. AppQoS provides the ability to prioritize, rate-limit, perform DSCP rewrite on, set loss priority for, and queue traffic. It provides the granularity of the stateful firewall rule base (including User Role Firewall and Dynamic Application identified by App ID) to match and enforce quality of service (QoS) at the application layer. This results in prioritizing business-critical applications, queuing up noncritical applications, and selectively allowing business-critical applications while blocking undesirable or malware-infected applications based on network policy, user, and time.

Application Enforcement with AppFW

Requirement: Traditionally, applications like HTTP, SMTP, and DNS use well-known standard ports and are easily controlled by a stateful firewall. However, it is possible to run these applications on any port, as long as the client and server are using the same protocol as the well-known ports.

Additionally, with the growing popularity of Web applications and the shift from traditional, full client-based applications to the Web, more and more traffic is being transmitted over HTTP. Network administrators must be able to detect evasive applications and enforce protocol and policy control at Layer 7. An application firewall must be able to identify not only HTTP, but also any application running on top of it, allowing you to properly enforce your organization's policies. For example, an application firewall rule could block HTTP traffic from Facebook, but allow Web access to HTTP traffic from Microsoft Outlook.

Solution: Application Firewall (AppFW) refers to the ability to take the results from the App ID engine and leverage them to make an informed decision to permit, deny/reject, or redirect the traffic. AppFW sits on top of the existing stateful firewall engine that makes decisions based on the standard seven-tuple (from-/to-zone, source/destination IP address, source/destination port, and protocol). This still allows you to enforce traditional firewall controls on the traffic while layering AppFW to ensure that the application conforms not only to the well-known port information, but to what is actually being transmitted between the client and the server. AppFW provides an auxiliary rule base that is tied to each firewall rule for maximum granularity with the ability to leverage the standard match criteria of the firewall rule, plus the application identity. You can permit, deny, or reject

applications, and also use a special redirect feature for HTTP and HTTPS. The redirect action provides a better user experience; instead of explicitly blocking the application, the user can be redirected to a custom Webpage or an externally hosted URL.

Application Visibility with AppTrac

Requirement: Administrators need visibility and control of applications and Websites (including related sub-Websites) resident in all parts of their networks, from the wired or wireless edge all the way through the core and the data center, as well as of application traffic from the enterprise to the private cloud, public cloud, or any service on the Internet. Administrators need to optimize the network for each and every application, enhance security for those applications, and provide data for business analytics. Administrators require that they can log/report, as well as enforce actions on sessions based on the result of App ID. Administrators must be able to send the results of App ID via syslog so that these results can be leveraged both on box and on an external device such as Juniper Secure Analytics (JSA Series) appliance which can provide a rich logging and reporting experience based on the data in these logs.

Solution: AppTrack is essentially a logging and reporting tool that can be used to share information for application visibility. After App ID identifies an application, AppTrack not only keeps statistics on the box for application usage, but also sends log messages via syslog providing application activity update messages. Because these log messages are sent by syslog, they can be consumed by Juniper products like the JSA Series appliances as well as third-party devices.

Resources and References

- The Juniper TechLibrary includes everything you need to understand and configure all aspects of AppSecure. See http://www.juniper.net/techpubs/en_US/junos15.1x49-d40/information-products/pathway-pages/security/security-application-identification.html.
- Resources for the AppSecure including datasheets, white papers, and solution briefs. See http://www.juniper.net/us/en/products-services/security/appsecure/, and http://www.juniper.net/assets/us/en/local/pdf/solutionbriefs/3510505-en.pdf.
- *The Juniper SRX Series*, published in 2013 by O'Reilly Media. See http://www.juniper.net/us/en/training/jnbooks/oreilly-juniper-library/srx-series/.

Chapter 5

Secure Analytics

Networks are growing larger and more complex than ever before. At the same time, multiple threats to the security of those networks are emerging and spreading rapidly. As shown in Figure 5.1, there are also more possible points of entry into any given network because of the increase in user mobility, the number of remote locations that might exist, and the sheer number of devices accessing the network.

The digital market economy, with its continual barrage of new applications and technologies, also creates additional risks and invites a slew of new attacks on networks. In some organizations *security breaches* can go completely undetected for months, while others have IT departments with staff dedicated to protecting a network against malicious activity. They must analyze data from a multitude of sources in order to understand what threats are facing a network, then they must determine what actions to take to address those threats.

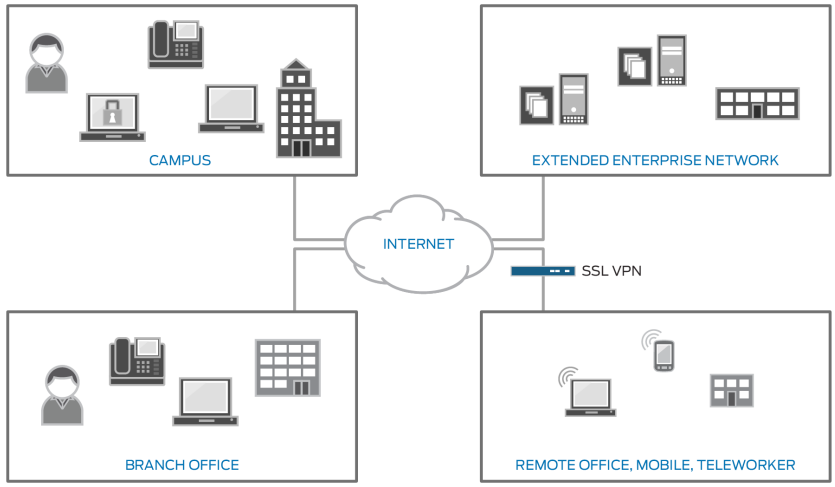

Figure 5.1 Enterprise Network

What IT staffs need is a complete, holistic solution that provides layered security to protect from threats that occur at all layers and at every location of a network, including branch offices, campuses, and extended enterprises. Without such a solution, IT professionals cannot fully manage all the threats a network can incur. They need:

- Comprehensive visibility that can analyze everything happening in the network.
- Analytics that will analyze and investigate potential threats in near real time.
- Actionable intelligence that will identify targets, threats, and incidents.

IT departments also need to keep abreast of compliance requirements, providing:

- Accountability that can survey the reports on who did what and when.
- Transparency that can provide visibility into the security controls, business applications, and assets that are being protected.
- Measurability that can provide metrics and reporting around IT risks within a company.

Introduction to SIEM

Security information and event management (SIEM) software provides a powerful way for organizations to detect the latest security threats to their networks before they can cause damage. SIEM provides a holistic view of an organization's IT security by providing real-time reporting coupled with long-term analysis of security events.

SIEM software logs event records from sources throughout a network. Those logs provide important forensic tools to an IT staff, which the software then helps to analyze. Complete log collection also helps address many compliance reporting requirements.

Parsing and normalization maps log messages from different systems into a common data model, enabling IT professionals to better connect and analyze related events, even if those events are initially logged in different source formats. Additionally, correlation links log events from disparate systems or applications, which greatly speeds up not only the detection of, but the reaction to, security threats.

SIEM aggregation can also reduce the volume of event data by consolidating duplicate event records and then reporting on the correlated, aggregated event data in real time, comparing it to long-term summaries.

How SIEM Works in an Attack

Let's begin with a look at a basic network attack as shown in Figure 5.2.

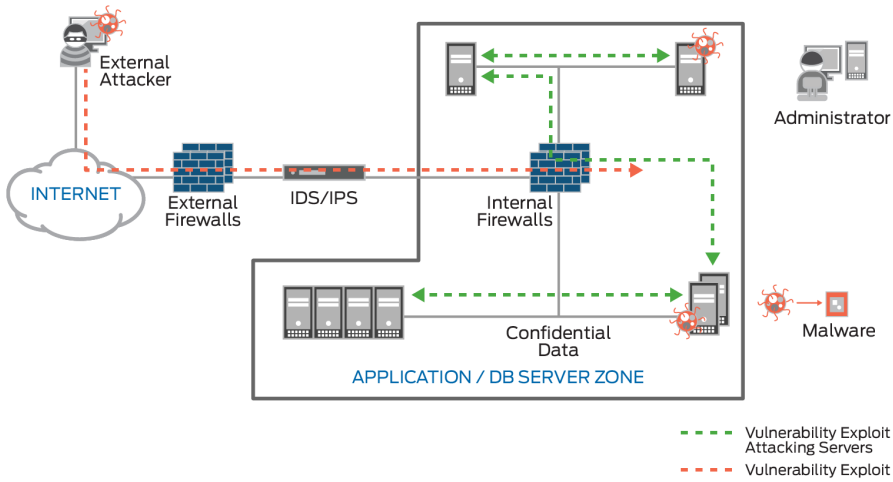

Figure 5.2 Example of a Basic Attack to a Network

In Figure 5.2, the attacker on the left scans the perimeter defenses to find a hole in the network. The attack bypasses network defenses and compromises Web servers using a vulnerability exploit. From the Web server, the attack pivots to the database server, which holds confidential data and installs malicious software that opens a backdoor for the attacker to steal data.

How would one detect such an attack without using SIEM? Figure 5.3 shows the steps in a traditional network defense.

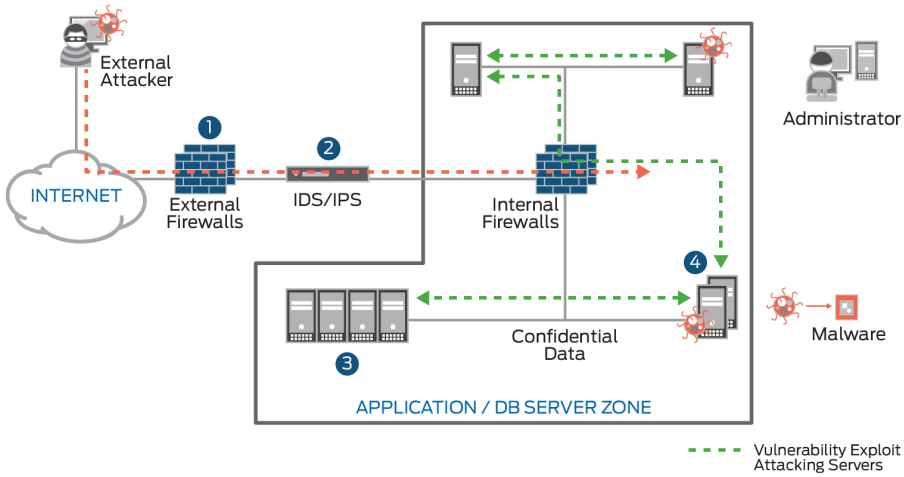

Figure 5.3 Analyzing the Basic Attack Without Using SIEM

You can see, in Figure 5.3, that the network uses:

- Firewall logs with events for reconnaissance, scanning, and so on.
- Intrusion detection service (IDS) or intrusion prevention system (IPS) logs that have exploit signatures triggering (both behavior and anomaly).
- Web or application server logs (to access inbound or outbound traffic).
- And of course, database logs.

In Figure 5.4, when the same attack occurs in a network using SIEM, the software provides insight into all the IT components (gateways, servers, firewalls, and so on).

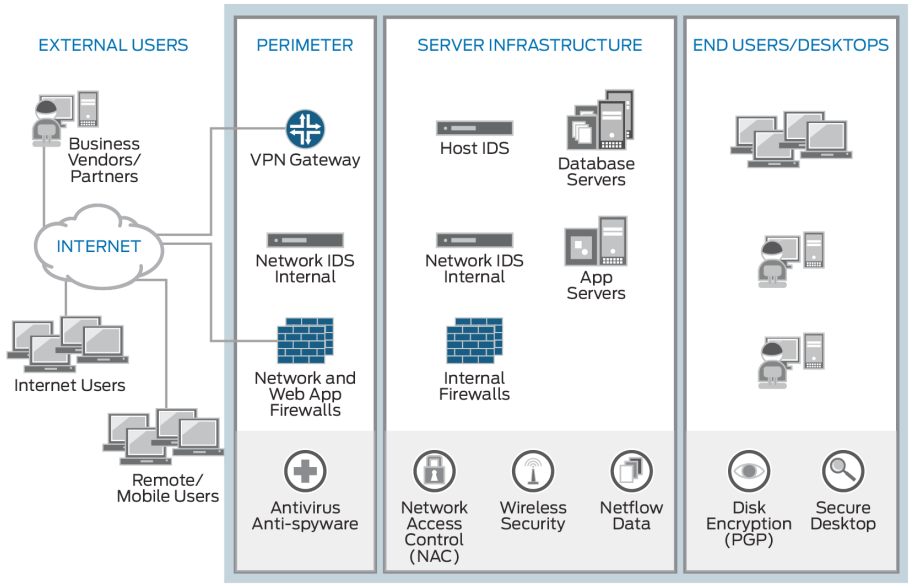

Figure 5.4 SIEM Holistic View

SIEM software centrally collects, stores, and analyzes logs from *perimeter* to end user. It monitors for security threats in real time for quick attack detection, containment, and response with holistic security reporting and compliance management.

It's time for SIEM software in any network that is open to attacks.

Juniper Networks Secure Analytics

Once you realize the value of a SIEM and its functionality, you need to understand how JSA can support SIEM security and compliance requirements.

A JSA Series appliance is a SIEM appliance that solves many requirements of IT staffs around the world. To better understand how JSA works, let's briefly review its key components and how they operate as a SIEM solution.

Event Collection and Processing

JSA combines many key SIEM features (see Table 5.1) but the core components of the JSA Series are an event processor, a flow processor, an event collector, and a magistrate (console).

An event is a record from a *log source*, such as a firewall, a router, a server, an IDS, or an IPS, that describes an action on a network or a host.

As shown in Figure 5.5, JSA event processing involves the following steps:

1. Log sources typically send syslog messages (but they can use other protocols, too).

2. The event collector receives the raw events as log messages from a wide variety of external log sources.

3. Device Support Modules (DSMs) in the event collectors parse and normalize raw events as the raw log messages remain intact.

4. The classification engine and the *rules* are responsible for processing events received by JSA and comparing them against defined rules, keeping track of systems involved in incidents over time, generating notifications to users, and generating offenses.

5. Event processors receive the normalized events and raw events to analyze and store them.

6. The magistrate correlates data from event processors and creates offenses.

7. Event storage (Ariel) is a time series database for events and flows where data is stored on a minute-by-minute basis. Data is stored where the event is processed.

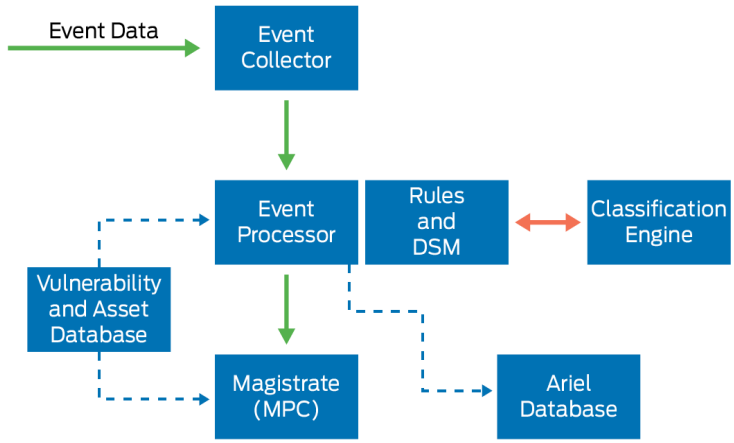

Figure 5.5 Event Collection and Processing Flow Diagram

Flow Collection and Processing

A flow is a communication session between two hosts that provides information about network traffic and can be sent to JSA in various formats, including network taps, span or mirror ports, flowlog files, NetFlow, J-Flow, sFlow, and Packeteer.

The flow processing (see Figure 5.6) involves the following steps:

1. The flow collector reads different types of flow data and creates flow records to be processed.

2. The event collector completes a number of flow processing functions, such as:

- Removing duplicate flows when multiple flow collectors are providing data to flow processor appliances.
- Recognizing flows from each side and combining them into one record. When data is not received from both sides, the event collector then analyzes and combines the external flow sources, such as NetFlow, that might only report ingress or egress traffic, as well as instances where span traffic enters a network from a single point, and exists through another, creating asymmetric reporting of data to flow collectors.
- Monitoring the number of incoming events and flows to the system to manage input queues and licensing.
- Applying routing rules for the system, such as sending data to offsite targets, external syslog systems, JSON systems, other SIEM products, and so on.

3. Classification engine and the rules are responsible for processing events received by JSA and comparing them against defined rules, keeping track of systems involved in incidents over time, generating notifications to users, and generating offenses.

4. Event processors parse the message's fields (IP address, ports, and so on) and store data in the Ariel database.

As you can see, JSA goes beyond traditional SIEM products and network behavior analysis (NBA) products to create a command-and-control center that delivers threat analytics, log analytics, and complete compliance measurability.

When it comes to secure analytics, JSA Series appliances can protect your network. Let's look very briefly at all the features and benefits of the JSA Series.

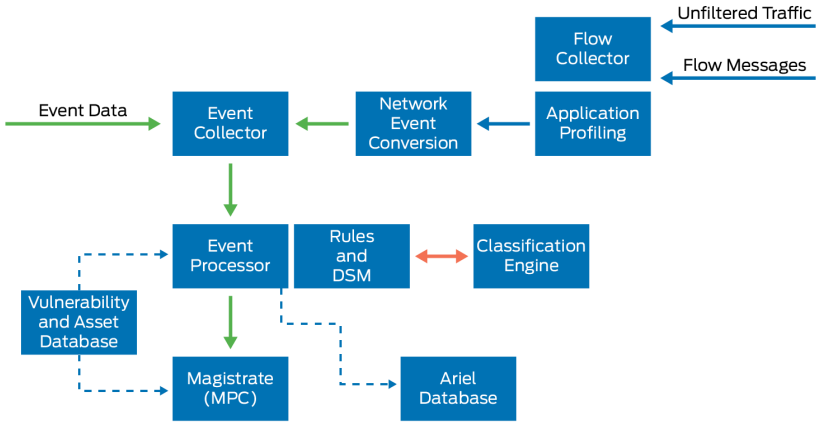

Figure 5.6 Flow Collection and Processing Flow Diagram

JSA Appliance Features and Benefits

JSA Series appliances come in several form factors to enable you to scale their features and benefits:

- JSA Virtual Appliance – A virtualized platform that can be deployed as an all-in-one appliance or in a distributed setup as a console, or as an event or a flow processor. A JSA virtual appliance can also be deployed as a store and forward event collector.
- JSA3800 – An enterprise-class appliance that provides a scalable network security management solution for medium-to-large size companies, including globally deployed organizations. It is also the base platform for an enterprise-class scalable secure analytics solution. JSA3800 can be deployed as an all-in-one appliance or in a distributed setup as a dedicated event, flow, or combination processor. It can also be deployed as a store and forward event collector.
- JSA5800 – An enterprise and carrier-class appliance that provides a scalable network security management solution for medium-size companies and scales to support large, globally deployed organizations. JSA5800 can be deployed as an all-in-one appliance or in a distributed setup as a console or dedicated event or flow processor. It can also be deployed as a store and forward event collector.
- JSA7500 – An enterprise and carrier-class appliance that provides a scalable network security management solution for large, globally deployed organizations. JSA7500 can be deployed as a console or distributed event or flow processor. It can also be deployed as a store and forward event collector.

Table 5.1 details some of the major features and benefits of owning and using JSA appliances, many of which go beyond the SIEM discussions in this chapter.

Table 5.1 JSA Features and Benefits

Features	Description	Benefits
All-in-one appliance	Event collection, flow collection, event processing, flow processing, correlation, analysis, and reporting are all embedded within JSA.	All core functions are available within the system and it is easy for users to deploy and manage in minutes. The architecture provides a streamlined solution for secure and efficient log analytics.
Distributed support	Ability to scale to large distributed deployments that can support up to 5 million events per second.	Gives users flexibility to scale to large deployments as their business grows and can be easily deployed in large distributed environments.
HDD implementation	Utilizes SAS HDD in RAID 1 and RAID 10 setups.	SAS HDD is designed for 24x7 operations. RAID 1/10 implementation provides best possible performance and redundancy.
Quick install	Comes with an easy, out-of-the-box setup wizard.	Users can install and manage JSA Series appliances in a couple of steps.
Automatic updates	Automatically downloads and deploys reputation feeds, parser updates, and patches.	Users do not need to worry about maintaining appliance and OS updates and patches.
High availability	Users can deploy all JSA Series appliances in HA mode.	Users can deploy JSA with full active or passive redundancy. This supports both deployment scenarios: all-in-one and distributed.
Built-in compliance reports	Out-of-the-box compliance reports are included with the JSA.	Provides more than 500 out-of-the-box compliance reports.
Reporting and alerting capabilities for control framework	Control Objectives for Information and related Technology (CobiT) International Organization for Standardization (ISO) ISO/IEC 27002 (17799), Common Criteria (CC) (ISO/IEC 15408) NIST special publication 800-53 revision 1 and Federal Information Processing Standard (FIPS) 200.	Enables repeatable compliance monitoring, reporting, and auditing processes.

Features	Description	Benefits
Compliance-focused regulation workflow	Payment Card Industry Data Security Standard (PCI DSS) Health Insurance Portability and Accountability Act (HIPAA) Sarbanes-Oxley Act (SOX) Graham-Leach-Bliley Act (GLBA) Federal Information Security Management Act (FISMA)	Supports multiple regulations and security best practices. Includes compliance-driven report templates to meet specific regulatory reporting and auditing requirements.
Management-level reports on overall security state	The JSA reports interface allows you to create, distribute, and manage reports that are generated in PDF, HTML, RTF, XML, or XLS formats.	Users can use the report wizard to create executive and operational level reports that combine any network traffic and security event data in a single report.
One-stop support	Juniper Networks Technical Assistance Center (JTAC) supports all aspects of JSA.	Users do not need to go to several places to get support, even for multivendor issues.

JSA Use Case

As a final step, let's review a use case for JSA, and follow the requirements and the solution. This use case concerns the PCI DSS that was created by major credit card companies to ensure privacy and security of credit card holders. All organizations that deal with credit card processing and transactions need to comply with these standards to avoid fees and penalties, and this use case will show you how JSA addresses the six main PCI DSS objectives.

PCI DSS Requirements

The PCI DSS standard outlines six relatively broad control objectives for network security:

- Build and maintain a secure network.
- Protect cardholder data.
- Maintain a vulnerability assessment (VA) program.
- Implement strong access control measures.
- Regularly monitor and test networks.
- Maintain an information security policy.

It is not an easy task for IT administrators to implement these standards across their network as there is no single product that complies

with all six standards. Many SIEM and log management products claim to answer all these concerns, but the PCI DSS standard calls for more than the collection and correlation of logs. Insight into the network from the passive monitoring of network communications must be put in place in conjunction with aggregation and correlation of logs from the security and network infrastructure.

The Solution

JSA is a network security management platform that facilitates the comparison of data from the broadest set of devices and network traffic. It combines log management, SIEM, and network behavior anomaly detection (NBAD), into a single integrated end-to-end network security management solution. This allows administrators to get a complete picture of their network security posture. This surveillance capability brings together all pertinent PCI DSS data for the purpose of executing and maintaining an organization's PCI DSS program. Table 5.2 details the JSA approach to meeting PCI requirements. Whether it's for the PCI industry, the Federal Information Security Management Act (FISMA), or any other compliance-driven organization, JSA has a complete solution.

Table 5.2 JSA Approach to Meeting PCI Requirements

PCI Requirement	JSA Approach
Build and maintain a secure network	▪ Detection and classification of protocols and applications within the network. ▪ Automatic policy creation through learning normal traffic behavior and acceptable protocols, alerting when traffic deviates from normal patterns, and alerting when new servers, databases, protocols, or applications are discovered in the DMZ. ▪ Layer 7 visibility detects and alerts risky or secure protocols running over nonstandard ports, which indicates suspicious behavior. ▪ Real-time intuitive views of network traffic by protocol or application allow for in-depth analysis and troubleshooting. ▪ Stores flows like NetFlow, sFlow, and jFlow and allows for detailed forensic searching of network communications associated with risky or mistrusted protocols. ▪ Default PCI report templates and a flexible reporting wizard provide in-depth reports on PCI-related networks and services.

PCI Requirement	JSA Approach
Protect card holder data	▪ Send alert and notification of any suspicious attempts to access sensitive data. ▪ Detect unencrypted data even in the absence of intrusion detection systems. ▪ Store the content from flows, which allows detection of unencrypted user name and passwords, or information on potential data theft. ▪ Logging from encryption technologies such as SNMPv3 devices.
Maintain VA program	▪ Automatic correlation of antivirus data with other logs and network information for accurate detection and prioritization of threats. ▪ Reporting and real-time viewing of antivirus logs. ▪ Integration with vulnerability management and assessment tools used for creation of asset/host profiles. ▪ Asset profiles are centrally stored within the JSA and used for detection of new hosts on the network, new services running on a host or network, and accurate prioritization of threats based on vulnerability information. ▪ Use real-time passive profiling to augment vulnerability data, which is typically not kept up-to-date, by using network communications to profile which services are running on hosts and keep asset profiles current.
Implement strong access control measures	▪ Complete auditing and alerting for access, configuration changes, and data changes to systems and databases with cardholder data. ▪ Detection of multiple logins that are followed by a failed login from suspicious or unknown hosts. ▪ Default, out-of-the-box authentication log correlation rules allow for easy identification of regulatory compliance servers and quick configuration of internal policies.
Regularly monitor and test networks	▪ Out-of-the-box customizable access and authentication rules allow for easy detection of threatening or invalid access attempts. ▪ Deep inspection analyzes all log data and network communications to monitor and audit all activity around an access offense. ▪ File integrity monitoring and notification through log analysis. ▪ Backup and archive of access audit trails. ▪ Provides continuous monitoring of security, systems, and processes. ▪ Real-time alerting and notification of changes to the network, threats or violations that impact meeting compliance, and views and historical reports of all collected network and log data. ▪ Up-to-date vulnerability information through the use of passive profiling of network communications.

PCI Requirement	JSA Approach
Maintain an information security policy	▪ Continuously analyzes all network and security data for identification of threats and vulnerabilities. ▪ Automatically learns all assets and hosts on the network and provides user identity profiles and running services profiles based on passive vulnerability assessment and active vulnerability assessment. ▪ Default built-in policy rules map directly to PCI requirements. ▪ Easy-to-use customizable rules engine that allows organizations to build their own compliance intelligence for monitoring and notification of specific violations. ▪ Offenses provide documented and historical perspective of all analysis and data associated with a PCI-related incident.

Resources and References

- This technical documentation includes everything you need to understand and configure all aspects of JSA: http://www.juniper.net/techpubs/en_US/release-independent/jsa/information-products/pathway-pages/jsa-series/product/
- JSA7500 introduction video. In this video, learn about JSA7500 and its components: https://www.youtube.com/watch?v=mcVUm2MsN2g
- This white paper contains PCI DSS objectives for network security and the solution using JSA: http://www.juniper.net/us/en/local/pdf/whitepapers/2000260-en.pdf
- This technical documentation includes all the information you need to understand and configure all aspects of QRadar software: http://www-01.ibm.com/support/docview.wss?uid=swg21614644
- This IBM page can help you to learn more about QRadar software and its features: http://www-03.ibm.com/software/products/en/qradar-siem
- This is the official IBM Security Support channel. It provides presentations and videos—such as IBM Security QRadar Open Mic webcasts—created by the IBM Support team: https://www.youtube.com/playlist?list=PLFip581NcL2XlvaEyrZMm3Nf1-Mc5-wRk
- This site provides WordPress posts on SIEM from the professionals: http://siemthoughtsonsecurity.wordpress.com/tag/what-is-a-siem/
- This site provides articles on SIEM from the professionals: http://www.networkworld.com/article/2180119/tech-primers/5-reasons-why-siem-is-more-important-than-ever.html

Chapter 6

Intrusion Detection and Prevention

Security has long been important to network technology. But there is an increasing focus on it today because most business networks are designed to provide access to the Internet and other public networks in order to perform their core operational functions. As shown in Figure 6.1, a typical business network has several access points to other networks, both public and private. Thus, securing organizational networks as well as their multiple access points is now fundamentally important, if not critical, for businesses to survive.

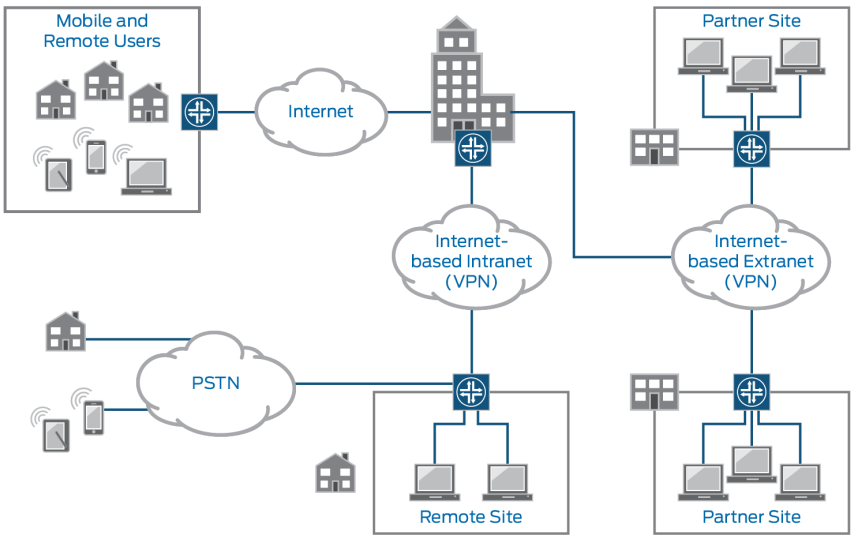

Figure 6.1 The Network Today

The challenge is maintaining the security of these networks while keeping them open to their customers. And of course, the threats are constantly changing; the network-based attacks that caught your attention only yesterday can and will continue to evolve. Currently, attacks are so sophisticated that they can thwart the best security systems, especially those that still operate under the assumption that networks can be secured by encryption or firewalls. Unfortunately, those technologies alone are not sufficient to counter today's attacks.

For example, Figure 6.2 illustrates the frightening frequency and sophistication of cyberattacks. Today's attackers have increased knowledge and understanding of the technology, infrastructure, and systems of their victims. In addition, the amount of knowledge an attacker needs to have about your network in order to launch a sophisticated attack is decreasing. This means sophisticated attacks are growing more severe each day. (*Source: Citi Online Academy, Digital Security – Cyber Security and Fraud Prevention*)

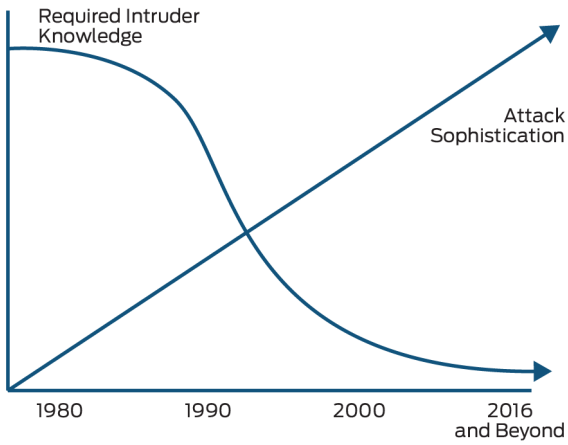

Figure 6.2 Attack Sophistication vs. Intruder Technical Knowledge

While your network still needs firewalls and encryption to improve security, you also need security systems that will watch your network and detect suspicious activity (such as attackers gathering intelligence about your network) 24 hours a day. When these tools observe any suspicious activity or event, they produce alerts for the network administrators. Often, they can detect attacker activity even before the attack begins.

Intrusion detection is the process of monitoring the events occurring in your network and analyzing them for signs of possible incidents, violations, or imminent threats to your security policies. And *intrusion*

prevention is the process of performing intrusion detection and then stopping the detected incidents. These security measures are available as IDS and IPS, which become part of your network to detect and stop potential incidents.

How Does IDP Work?

IDP constantly watches your network, identifying possible incidents and logging information about them, stopping the incidents, and reporting them to security administrators. In addition, some networks use IDP systems for other purposes, such as identifying problems with security policies, documenting existing threats, and deterring individuals from violating security policies. IDP systems have become a necessary addition to the security infrastructure of most organizations, precisely because they can stop attackers while they are gathering information about your network.

Most IDPs typically record information and produce reports. But many IDPs can also respond to a detected threat by attempting to prevent it from succeeding. This process can use several different response techniques such as involving the IDP in stopping the attack itself, changing the security environment (for example, reconfiguring a firewall), or even changing the content of the attack.

Figure 6.3 illustrates the following general components of an IDP solution (note that specific network architecture will differ depending on the exact type of IDP):

- *Sensors* or *agents* monitor and analyze activity on the host, node, or network (the term *agent* is typically used for host-based IDP technologies).
- *Management servers* are available as either software or as an appliance that is the core of the IDP solution. A management server:
  - Manages the agents and sensors, collects data from them, analyzes the data received, and identifies intrusion attempts.
  - Compares events from multiple management servers to see if there are correlations between triggered events (multiple servers are common in larger networks, but not required in smaller deployments).
- *Database servers* act as a centralized repository for event information recorded by sensors, agents, and/or management servers. Many IDPs provide support for database servers.

- The *console* is the administrator program interface to the IDP system and is used to configure agents or sensors, run updates, and monitor and analyze events.

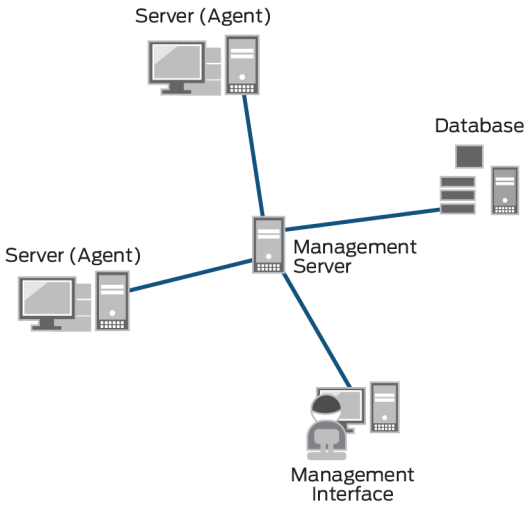

Figure 6.3 IDP Components

All of this observation generates lots of data. So in addition to monitoring and analyzing events in order to identify undesirable activity, most IDP technologies archive the recorded data locally, although it might also be sent to separate systems, such as centralized logging servers, security information and event management (SIEM) solutions, or enterprise management systems. Reports summarizing monitored events and details can be provided on any event of interest. And once data is flagged as suspicious, the IDP system notifies the security administrators through e-mail, text, messages on the IDP user interface, Simple Network Management Protocol (SNMP) traps, system log messages, and user-defined programs and scripts.

Network-based IDP and *host-based* IDP are two different types of IDP technology characterized by the types of events they monitor and the ways in which they are deployed, as depicted in Figure 6.4:

- *Network-based* IDP monitors network traffic for particular network segments or devices, and analyzes the network and application protocol activity to identify suspicious activity.
  - Wireless components monitor wireless network traffic and analyze it to identify suspicious activity involving the wireless networking protocols.

- Network behavior analysis (NBA) examines network traffic to identify threats that generate unusual traffic flows, such as distributed denial-of-service attacks, certain forms of malware, and policy violations (for example, a client system providing network services to other systems).

- *Host-based* IDP monitors the characteristics of a single host, and the events occurring within that host, for suspicious activity.

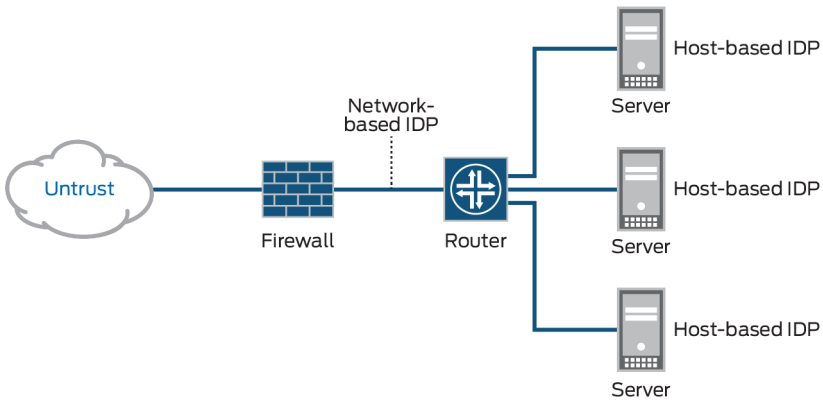

Figure 6.4 Deploying Network-Based and Host-Based IDPs in a Network

Most IDP technologies use multiple detection methodologies, either separately or integrated with, to provide wider and more accurate detection. Table 6.1 lists three IDP detection methodologies (*signature-based*, *anomaly-based*, and *stateful protocol analysis*) that are typically used to detect incidents.

Table 6.1 Common IDP Detection Methodologies

Detection Method	Description	Example	Benefits
Signature-Based Detection	Signature-based detection compares signatures against observed events to identify possible incidents.	A telnet attempt with a root username, which is a violation of an organization's security policy or an e-mail, with *Free Pictures* in the subject line, and an attachment with the filename freepics.exe, all of which are characteristics of a known form of malware.	This is the simplest detection method because it compares only the current unit of activity (such as a packet or a log entry, to a list of signatures) using string comparison operations.

Detection Method	Description	Example	Benefits
Anomaly-Based Detection	Anomaly-based detection compares definitions of what is considered normal activity with observed events in order to identify significant deviations. An IDP using anomaly-based detection has profiles that represent the normal behavior of such things as users, hosts, network connections, or applications. The profiles are developed by monitoring the characteristics of typical activity over a period of time.	A profile for a network might show that Web activity comprises an average of 13% of network bandwidth at the Internet border during typical workday hours. The IDP then uses statistical methods to compare the characteristics of current activity to thresholds related to the profile, such as detecting when Web activity consumes significantly more bandwidth than expected and alerting an administrator of the anomaly.	This detection method can be very effective at spotting previously unknown threats.
Stateful Protocol Analysis	Stateful protocol analysis compares predetermined profiles of generally accepted definitions for benign protocol activity for each protocol state against observed events in order to identify deviations. Unlike anomaly-based detection, which uses host or network-specific profiles, stateful protocol analysis relies on vendor-developed universal profiles that specify how particular protocols should and should not be used. The *stateful* in stateful protocol analysis means that the IDP is capable of understanding and tracking the state of network, transport, and application protocols that have a notion of state.	When a user starts a File Transfer Protocol (FTP) session, the session is initially in the unauthenticated state. Unauthenticated users should only perform a few commands in this state, such as viewing help information or providing usernames and passwords. An important part of understanding state is pairing requests with responses, so when an FTP authentication attempt occurs, the IDP can determine if it was successful by finding the status code in the corresponding response. Once the user has authenticated successfully, the session is in the authenticated state, and users are expected to perform any of several dozen commands. Performing most of these commands while in the unauthenticated state would be considered suspicious, but in the authenticated state performing most of them is considered benign.	This analysis identifies unexpected sequences of commands, adds stateful characteristics to regular protocol analysis, and adds reasonableness checks for individual commands (for example, min/max lengths).

Choosing IDP Systems

IDP technologies can provide a wide array of security capabilities for your network. Look for these common, but necessary, security capabilities:

- **Information gathering:** Systems that identify hosts and the operating systems and applications being used, as well as identifying general characteristics of the network.
- **Logging:** Your IDP should perform extensive logging of data related to detected events. This data can be used to confirm the validity of alerts, investigate incidents, and correlate events between the IDP and other logging sources. You should know that specific types of IDPs log additional data fields, such as network-based IDPs that perform packet captures, and host-based IDPs recording user IDs. Your IDP should permit administrators to store logs locally and send copies of logs to centralized logging servers (for example, syslog, security information and event management software). Also, your IDP should ideally synchronize its clocks using the Network Time Protocol (NTP), or through frequent manual adjustments, so that log entries have accurate timestamps.
- **Detection:** IDP technologies should typically offer extensive detection capabilities. The types of events detected and the accuracy of detection can vary greatly depending on the type of IDP technology being used. Most IDPs require at least some fine-tuning and customization, such as setting prevention actions to be enabled for particular alerts, to improve their detection and effectiveness.
- **Prevention:** Finally, most IDPs should offer multiple prevention capabilities. While the specific capabilities vary by IDP technology type, your IDP should allow administrators to specify the prevention capability configuration for each type of alert, including enabling or disabling prevention, as well as specifying which type of prevention capability should be used.

Juniper Networks IDP

Juniper Networks uses its SRX Series Services Gateways for intrusion detection and prevention services and its IDP policy configuration lets you selectively enforce various attack detection and prevention techniques on the network traffic passing through your chosen SRX Series device (the IDP-enabled device). You can define policy rules to match a section of traffic based on a zone, a network, or an application, and then take active or passive preventative actions on that traffic. The SRX Series device contains a full set of IDP signatures to secure networks against attacks.

Juniper Networks regularly updates the predefined attack database and makes it available on the Juniper Networks website. This database includes attack object groups that you can use in IDP policies to match traffic against known attacks. Although you cannot create, edit, or delete predefined attack objects, you can use the CLI to update the list of attack objects that you can use in IDP policies.

The SRX Series device can forward packet capture (PCAP) data from its traffic to a Juniper Secure Analytics (JSA) appliance using the PCAP Syslog Combination Protocol. With the PCAP Syslog Combination Protocol, the JSA appliance is capable of receiving both syslog and the additional PCAP data once configured with the SRX Series.

This Juniper Networks IDP system is shown in Figure 6.5, in a very small site deployment that larger networks can scale. The SRX Series device displays the visibility of incoming or outgoing traffic and the JSA appliance collects events, allowing real-time streaming of events and monitoring of events through a common dashboard.

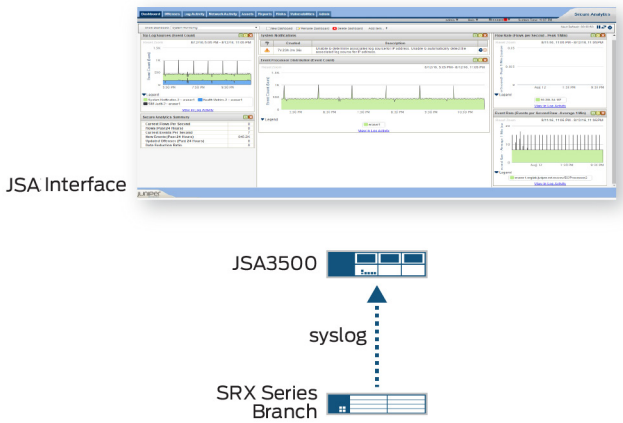

Figure 6.5 Small Site Deployment – JSA Appliance

Packet capture data is forwarded to the JSA appliance on a specified port, which is separate from the port that receives forwarded syslog data. The data contained in the packet capture and the outgoing port from the SRX Series is all configured using the SRX Series user interface.

Use Case: Protect Server and Application Vulnerabilities

Let's employ a simple use case to examine how the Juniper IDP system works. Assume that *Company X* is hosting its own commercial website as shown in Figure 6.6. Traffic is sent to the SRX Series services gateway for monitoring.

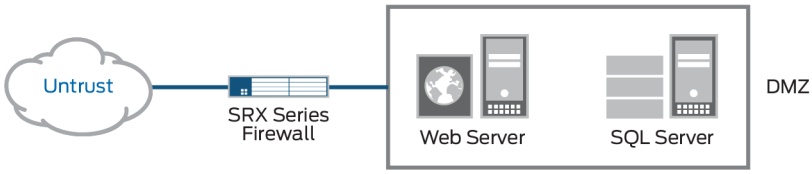

Figure 6.6 Company X Network Overview

When the traffic is sent to the SRX Series, it is discovered that the company's website is vulnerable to a specific SQL injection attack as shown in Figure 6.7. Packet capture provides the following details of the attack:

- The external connections are coming from the UNTRUSTED zone.
- The webserver (10.10.10.80) is located in our DMZ zone.
- The attack happens over TCP/80 or HTTP.
- The attack uses the GET command.
- The attack uses the following pattern:

```
form.php?q=1+UNION+SELECT+VERSION
```

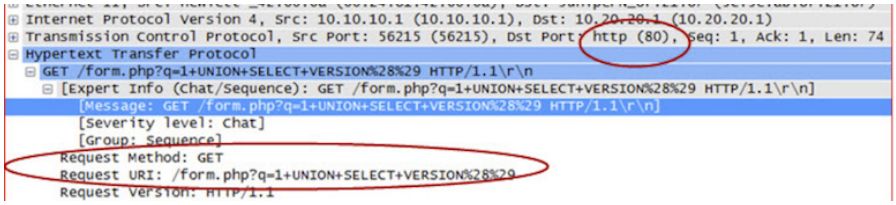

Figure 6.7 Packet Capture

Once the SRX Series services gateway has been spotted, and alerts are sent, the administrator can create custom attack objects to detect SQL injection as shown in Figure 6.8.

```
root@SRX# set security idp custom-attack HTTP:CUST_SQL_INJECT

root@SRX# edit security idp custom-attack HTTP:CUST_SQL_INJECT

[edit security idp custom-attack HTTP:CUST_SQL_INJECT]

root@SRX# set severity major

root@SRX# set attack-type signature protocol tcp destination-port match equal value 80

root@SRX# set attack-type signature direction client-to-server

root@SRX# set attack-type signature context http-get-url

root@SRX# set attack-type signature pattern "/form\.php\?q=1/+UNION/+SELECT/+VERSION"
```

Figure 6.8 Creating a Custom Attack Object

Each rule is composed of match conditions, objects, actions, and notifications. When you define an IDP rule, you must specify the type of network traffic you want IDP to monitor for attacks by using the following characteristics – source zone, destination zone, source IP address, destination IP address, and the application layer protocol supported by the destination IP address. The rules are defined in rulebases, and rulebases are associated with policies.

Figure 6.9 shows the result – the network drops the attack when the SQL injection attack is attempted.

```
name="SERVICE_IDP" application-name="NONE" rule-name="1" rulebase-
name="IPS" policy-name="COMPANY_X" repeat-count="0" action="DROP"
threat-severity="HIGH" attack-name="HTTP:CUST_SQL_INJECT" nat-source-
address="0.0.0.0" nat-source-port="0" nat-destination-address="0.0.0.0" nat-
destination-port="0" elapsed-time="0" inbound-bytes="0" outbound-bytes="0"
inbound-packets="0" outbound-packets="0" source-zone-
name="UNTRUSTED" source-interface-name="fe-0/0/7.0" destination-zone-
name="DMZ" destination-interface-name="fe-0/0/6.0" packet-log-id="0"
message="-"]
```

Figure 6.9 The Attack Is Now Known to the Network

Resources and References

- The Juniper TechLibrary documentation includes everything you need to understand Juniper's IDP system. See http://www.juniper.net/techpubs/en_US/junos15.1x49-d40/information-products/pathway-pages/security/security-idp-index.html.
- A tech note on Juniper SRX Series device forwarding of packet capture (PCAP) and syslog data to the JSA appliances. See http://www.juniper.net/techpubs/en_US/jsa2014.7/information-products/topic-collections/jsa-managing-juniper-pcap-data.pdf.
- The SANS Reading Room maintains, and makes available at no cost, a wide collection of research documents about various aspects of information security. It features over 2,460 original computer security white papers in 96 different categories. See http://www.sans.org/reading-room; http://www.sans.org/reading-room/whitepapers/detection/network-ids-ips-deployment-strategies-2143.
- The TechTarget network of technology-specific websites gives you access to industry experts, independent content, and analysis. See:
  - http://searchsecurity.techtarget.com/Do-you-need-an-IDS-or-IPS-or-both
  - http://searchsecurity.techtarget.com/feature/Enterprise-benefits-of-network-intrusion-prevention-systems
  - http://searchsecurity.techtarget.com/feature/The-basics-of-network-intrusion-prevention-systems
- Webopedia is an online tech dictionary for IT professionals, educators, and students. It also provides in-depth articles, study guides, and links to sources of further information on the topic, where applicable. See http://www.webopedia.com/DidYouKnow/Computer_Science/intrusion_detection_prevention.asp.
- The Computer Security Resource Center (CSRC) is the primary gateway for gaining access to NIST computer security publications, standards, and guidelines plus other useful security-related information. See http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-94.pdf.

Chapter 7

Network Security Management

Security teams must support internal and external compliance mandates, enable new services, optimize performance, ensure availability, and support the ability to troubleshoot efficiently on demand – all with no room for error. That's a lot to balance when managing network security.

In addition, an ever expanding matrix of users, devices, locations, and applications makes it difficult for IT staff to ensure that access controls and other security mechanisms are consistently applied to the same user without fail. And mobile workers need *anytime*, *anywhere* access to a broad array of applications, further taxing the security infrastructure.

This exponential growth in network traffic, combined with changes in mobile user behavior and an onslaught of new cloud services and applications, means the avenues available to malicious network attackers are expanding.

Managing security policy in these complex environments can be prone to error and overly time-consuming, especially if the management solutions employed are slow, unintuitive, or restricted in their level of granularity and control. Poor policy management can also lead to security misconfiguration, making the network vulnerable to sophisticated threats and regulatory noncompliance.

In order to successfully confront these challenges, network administrators will need to learn how to consistently deploy thousands of security policies across their network. To assist them in this task, they will need a new set of tools and technologies that can provide:

- More visibility into network security policies.
- A platform that will enable them to troubleshoot policy issues quickly and efficiently.
- The capacity to do more, with less resources.
- The option to deploy thousands of devices and VPNs in less time.

Overcoming Network Security Management Challenges

Security practitioners can no longer resort to CLI or device-level tools to implement next-generation firewall/perimeter security or manage policies. They need consolidated, easy-to-use interfaces that can help them implement security across the entire network. Usability is extremely important; it has become one of the key criteria in deciding between network security equipment vendors.

Data centers are undergoing a dramatic transformation. As Nemertes Research has noted, "workloads in today's data center move dynamically and start and stop based on real-time performance needs." See https://www.nemertes.com/reports/securing-physical-virtual-cloud-continuum. Unfortunately, in many enterprises, the current security architecture looks much the same as it did 15 years ago.

To keep pace with changes, networks need new security and compliance controls that span physical and virtual environments and dynamically enforce policy regardless of application type or user location. In evaluating next-generation security solutions, administrators must look for the following capabilities, as shown in Figure 7.1. This applies to all data centers regardless of their size.

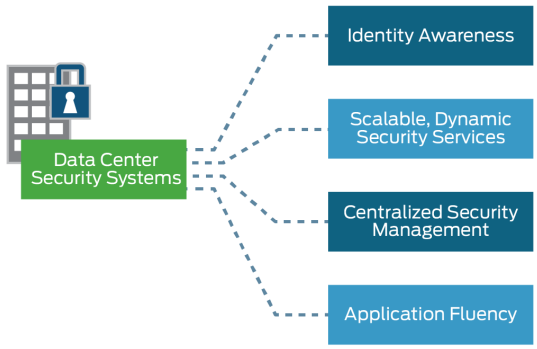

Figure 7.1 Security Systems for Data Center Networks

Network administrators must configure security policies on numerous platforms, both physical and virtual, sift through a flood of logged events to determine which ones require attention, and compile data to demonstrate compliance with government and industry-specific mandates. Today, network administrators must monitor multiple individual management consoles and create and maintain policy scripts for each different security platform – a manual and error-prone process that makes it difficult to apply policy consistently across the network. Similarly, each security system typically generates its own logs, creating silos of event data. Many solutions handle traffic flows and security events separately, which makes it virtually impossible to spot networkwide threats and anomalies.

Given the complexity of today's networks and the rapid evolution of threats from multiple sources, administrators need security management solutions that can:

- Automate security device and service provisioning.
- Abstract and centralize policy definition.
- Provide policy life cycle management.
- Deliver a unified solution for managing traffic flow and security events.
- Provide a single management interface for both physical and virtual systems.
- Correlate data from diverse sources on the network.

Evolving Security Management Solutions

With so many disparate vendor devices and hosts, security teams need a normalized, comprehensive view of their network, including routing rules, access rules, Network Address Translation (NAT), VPN, and more; hosts, including all products (and versions), services, vulnerabilities, and patches; and assets, including asset groupings and classifications. Although administrators need a holistic view of their network in order to see how all the pieces fit together, they must also be able to easily access the information on rules, access policies, and configuration compliance for a particular device.

This information must be provided in a simple, clear, and understandable format. The network components that impact the device come from various vendors, creating data of different vendor languages that must be deciphered, correlated, and optimized to allow administrators

to streamline rule sets. For example, network administrators need to be able to block or limit access by application and to view violations of these access policies. Beyond accessing a specific device, network administrators must also be able to drill down to each device level, accessing information on users, applications, vulnerabilities, and more. This provides administrators with a broader network view and enables them to focus on particular devices for management.

Junos Space Security Director Solutions

Juniper Networks Junos Space Security Director is an application that runs on the innovative, intuitive, and intelligent Junos Space Network Management Platform, providing detailed visibility into application and user performances, and reducing risk while enabling users to move quickly from knowing something is wrong, to doing something to fix the problem.

With Security Director, network administrators can provision identity-based and role-based policies across the entire network. It uses an innovative approach that abstracts the network security policy and then applies it to a group – effectively protecting an entire security domain. Security Director has an easy-to-use wizard-driven interface, granular configuration options, and predefined profiles for rapidly deploying devices and security services. Using Security Director, administrators can easily provision complex identity-based policies across their entire network, greatly improving its operational scale and efficiency, and enhancing overall policy consistency and security because of minimal operator error.

As shown in Figure 7.2, Security Director provides solutions to these problems by:

- Enabling efficient policy management for multiple security devices.
- Providing highly scalable device management to keep up with business growth.
- Allowing comprehensive security policy management for granular protection, including firewall, application security, VPN, and NAT, from a single location.

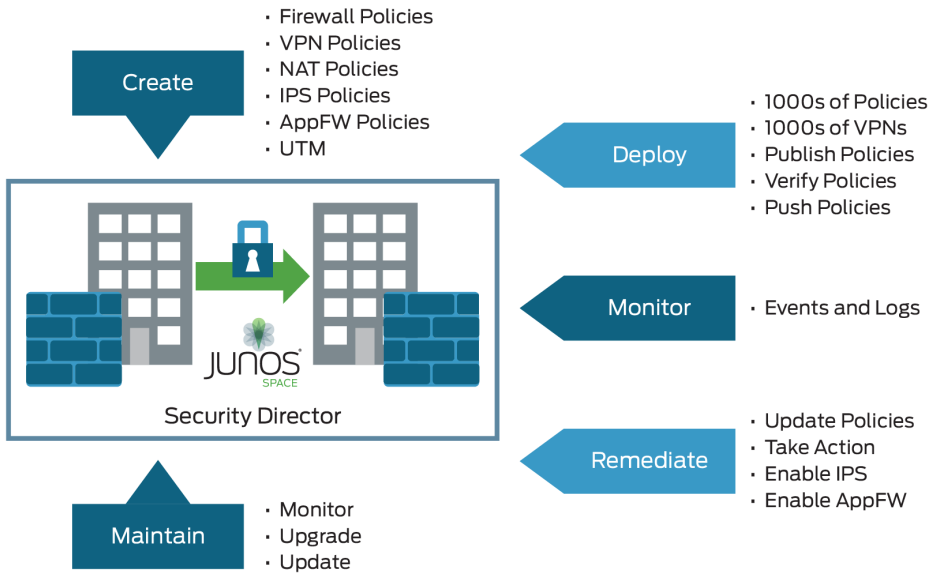

Figure 7.2 Junos Space Security Director Components

Efficient Security Policy Management

Security Director provides efficient security policy management supporting highly scalable policy implementation based on support for the Junos Space Network Management Platform. This efficient security policy management can be achieved by enabling and easing policy management (implementation and validation) across multiple devices, and reducing the chances of errors and misconfigurations in policy enforcement.

A single policy can be applied to multiple Juniper Networks SRX Series devices; you can apply a complex policy with thousands of rules to one or more SRX Series devices, and then scale the number of rules as required.

Highly Scalable Device Management

Security Director supports the management of thousands of devices running Junos OS. It can instantly scale to many devices by simply adding or deleting nodes on the fabric. In addition to managing many devices, Security Director can also increase the number of concurrent administrators.

Comprehensive Security Policy Management

In addition to enabling highly scalable device management, Security Director also enables companies to configure multiple security functions from a single location. Security Director can rapidly manage a comprehensive set of functions, including firewall, VPN, NAT, IPS, and application firewall, from a single management console. It eases policy configuration by providing read-write APIs for firewall policy, objects, and VPNs. A comprehensive and intuitive Google-like search mechanism is built into the main UI, which enables network administrators to quickly locate policy terms or issues, even in the rules, for faster maintenance.

Security Director Dashboard and Web GUI

The Web-based user interface for Security Director includes a dashboard that provides customizable, information-rich widgets offering visually intuitive displays that report security device status at a glance, as shown in Figure 7.3.

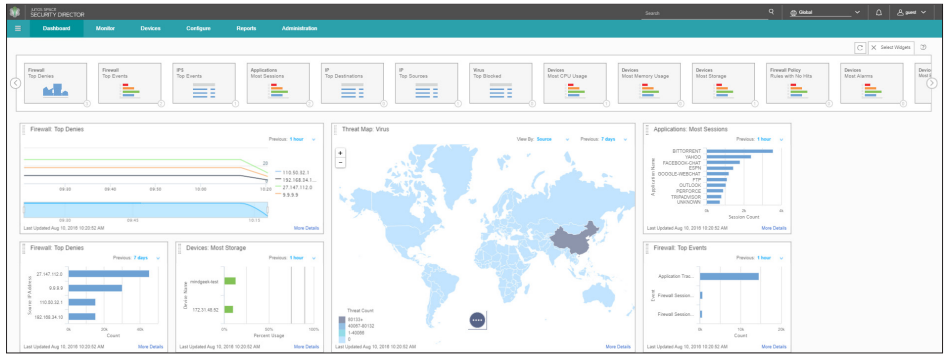

Figure 7.3 Security Director Dashboard

The Security Director Dashboard can display, among other features:

- A palate providing predefined widgets that display firewall, threat, IPS, application throughput, and device-related information.
- A quick view of important statistics for SRX Series devices, such as alarms, consumption for most CPU cycles, or RAM for a specific time period, and more.
- A threat map widget showing the number of IPS events detected per geographic location. The More Details option enables you to drill down to pre-filtered events.

As shown in Figure 7.4, the Web UI for Application Visibility displays an actionable intelligence feature that eliminates the need to manually create and manage the required firewall rules.

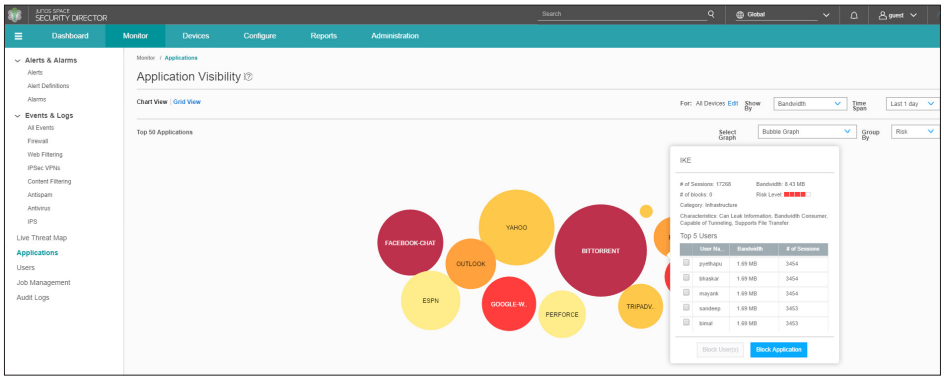

Figure 7.4 Security Director Application Visibility

The Security Director innovative application or user visibility charts display:

- Interactive or graphical summary of applications.
- Visual representation of the types and relative amounts of traffic passing through your network in a graph view. You can block users or applications by selecting the specific application from the graph.

As shown in Figure 7.5, the Web UI to detect threats shows the attack vectors for currently active IPS and virus attacks.

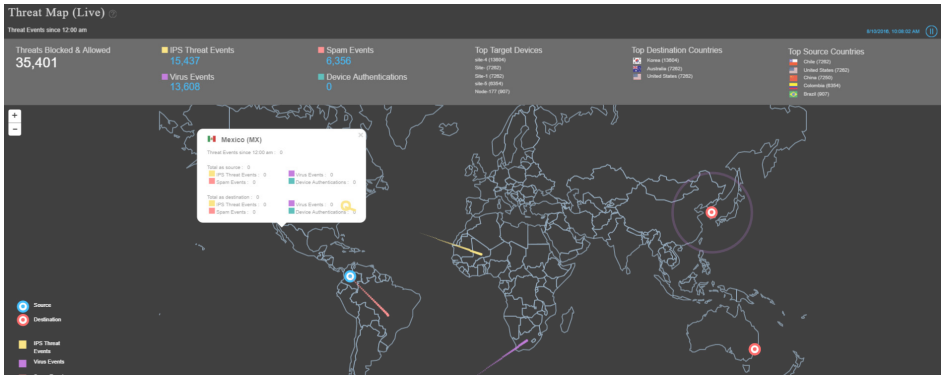

Figure 7.5 Security Director Live Threat Map

The Security Director live threat map provides:

- Live animation of the threat origin.
- Ability to zoom into a region for filtered threat view details.
- List of threat criteria.
- Insight bar to drill down for more detail.

For a complete list of dashboard features, visit the links at the end of this chapter.

Security Director Use Case

Let's review a use case for Security Director, and follow the requirements and the solution. This use case concerns a multinational financial services firm with corporate clients around the globe.

The firm wants to accelerate its business transformation and improve its capital position while reducing its Basel III risk and costs. At the same time, it wants to enhance its competitive positioning across its many businesses, so the financial services firm deploys a simple, open, and smart data center based on Juniper Networks MetaFabric architecture. The MetaFabric architecture is delivered through a combination of switching, routing, and security platforms while leveraging network orchestration, software-defined networking (SDN), and open APIs to simplify integration within the technology ecosystem.

The intelligence behind the robust security of the MetaFabric architecture is Security Director, which gives network administrators the power to centrally configure and manage application security, firewalls, IPS, VPNs, and security policies by using a single, intuitive interface.

MetaFabric architecture for this customer includes:

- QFabric System
- SRX Series Services Gateways
- MX Series 3D Universal Edge Routers
- Junos Space Network Management Platform, including Network Director and Security Director
- Junos Space Service Now and Service Insight

By deploying MetaFabric architecture, the financial services firm is able to:

- Improve application performance to better serve clients.
- Cut OpEx by simplifying the turn-up process for new data centers.
- Use automation to reduce IT workload.
- Reduce global data center footprint by more than forty percent.

With a highly intelligent infrastructure in place, this financial services firm has a rock-solid foundation for the future, with the ability to incorporate SDN capabilities when ready.

Resources and References

- Start here at the Security Director product page on the Juniper Networks website: http://www.juniper.net/us/en/products-services/security/security-director/
- The datasheet for Security Director provides a great product overview: http://www.juniper.net/assets/us/en/local/pdf/datasheets/1000332-en.pdf
- Security Director is an application that leverages the Junos Space Network Management Platform: http://www.juniper.net/us/en/products-services/network-management/junos-space-platform/
- Security Director Onboarding Guide: http://www.juniper.net/techpubs/en_US/junos-space15.2/topics/task/operational/junos-space-security-director-guide-onboarding.html
- Security Director technical documentation: http://www.juniper.net/techpubs/en_US/release-independent/junos-space-apps/junos-space-security-director.html

Chapter 8

Service Provider Edge Security

The digital universe containing banking data, corporate financials, health information, tax statements, and personal photos and videos is doubling approximately every two years, according to *USA Today*. Within that changing digital universe, Internet service providers (ISPs) and telecom cloud providers are transforming their service delivery models, customer experience, and operations environment. During this enormous transformation undertaking, it is vital that network security be given top priority.

Network Protection and Security

Protecting a network is not merely implementing firewalls anymore, and it is no longer safe just playing defense. Today's networks are subject to both active and passive attacks from various malicious sources that cause infection throughout the network. Active attacks, such as DoS attacks, IP address spoofing (or masquerade attacks), and malware created to target both physical and VMs, are the most complex security threats to manage because they target the control plane (the part of a network that carries signaling traffic and is responsible for routing) of network elements.

Security must be ingrained everywhere – in the protocols, the systems, the elements, the provisioning, and in the business surrounding the network. ISPs must implement security measures similar to those that financial institutions and governments use to protect their networks, devices, and data.

Network security is effective only if appropriate defensive mechanisms are established. Here are the most common categories of defensive mechanisms:

- *Corrective mechanisms* – Used to reduce the consequences of an incident by limiting the damage. They can be used during or after an incident. An example of a corrective mechanism is restoring system backups to rebuild a compromised system.
- *Detective mechanisms* – Used to detect and react appropriately to any incidents that occur. In the event of an attack, a detective mechanism signals the preventive or corrective mechanisms required to fix the issue. Typically, network security monitoring, including intrusion detection service (IDS), is used to detect attacks on a network and its supporting communications infrastructure.
- *Deterrent mechanisms* – Used to reduce attacks on a network. Similar to a *no trespassing* sign, deterrent mechanisms reduce the threat level by informing potential attackers that there will be adverse consequences for them if they persist in their actions.
- *Preventive mechanisms* – Used to strengthen a network against incidents usually by reducing and eliminating vulnerabilities. For example, strong authentication of users makes it less likely that unauthorized users can access the network, and more likely that users are positively identified.

Universal Edge Service Provider Network Security

An enhanced broadband network gateway (BNG) brings services and access methods together on a common platform, or a universal edge platform. The Juniper Networks MX Series 3D Universal Edge Router portfolio supports a universal set of services, enabling service providers to eliminate service silos, integrate community Wi-Fi access to provide an enriched residential subscriber experience, and consolidate their edge networks onto a single universal edge platform. Figure 8.1 shows a universal edge implementation example for a converged ISP network. Mobile IP, residential, and business services are consolidated onto a universal edge platform that provides universal access to any application from any device, anywhere.

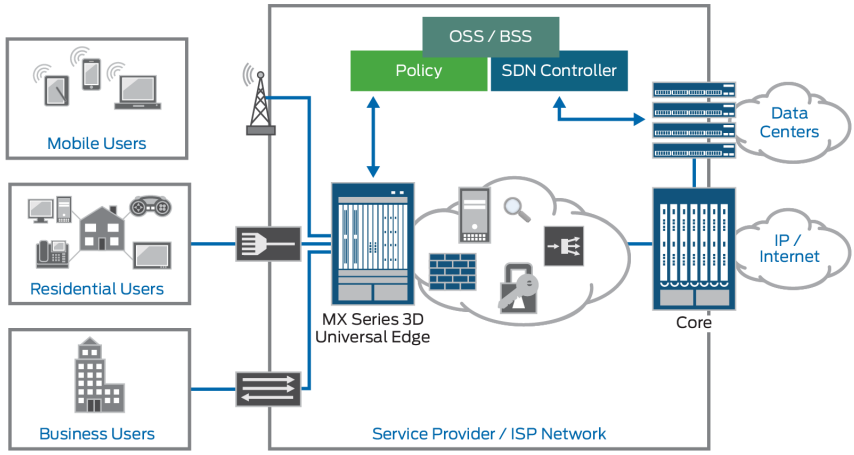

Figure 8.1 Converged Service Provider Network

The MX Series routing platforms concurrently support routing, switching, security, and Layer 4 through Layer 7 services under one Junos operating system (Junos OS), and within one element management system (EMS), which reduces network complexity and operational costs. Additionally, MX Series routers such as the MX2020 – the highest-capacity and highest-throughput edge router available – address the steadily growing interest among ISPs for collapsed edge and core networks, and can satisfy those requirements as well.

Router Security

Many functions are performed by MX Series Routing Engines, from processing routing protocol updates, to driving the command-line interface (CLI) and running Junos OS. Given that the Routing Engine is critical for the operation of the router and the network, it must be protected from unwanted activities such as malicious traffic seeking to gain unauthorized access, unintentional routing protocol updates from neighboring devices, or even legitimate traffic that exceeds a given bandwidth limit.

Junos OS provides configuration automation tools that enable dynamic prefix lists, and as your network scales and changes, the security framework automatically changes and adjusts without requiring additional input from your network administrator.

Mobile or Wireless Network Security

The old network security model of securing the perimeter of a mobile network no longer works. In reality, network complexities and communications undermine security between internal and external environments, and mobile networks are vulnerable to both inside-out attacks and inside-in attacks. Most threats and attacks originate from the Internet (making interfaces to the Internet the most critical to secure), while other common sources of threats are compromised subscriber devices and roaming peers. Today, those in charge of network security have to assume zero trust among network elements, and ISPs want to operate their network as a single-enforced domain where every element – not just those at the perimeter – becomes a policy enforcement point.

Therefore, in order to better combat and contain security threats in the mobile network, ISPs are moving toward a more distributed architecture, with detection and enforcement enabled everywhere. As the threat environment morphs and accelerates, you can have automated and centralized security polices with decentralized enforcement on switches and routers driven by dynamic and real-time security updates. Using software-defined networking (SDN) and Network Functions Virtualization (NFV), you can detect threats and enforce security policies with a high level of automated security, unified threat detection, and real-time protection.

Juniper Networks provides the building blocks for a new security model and a vision for SDSN as shown in Figure 8.2.

By using the Juniper Networks SRX Series Services Gateways (or the physical firewalls), and Juniper Networks vSRX virtual firewalls in combination with Spotlight Secure and Sky ATP, you can organize and coordinate threat intelligence, while employing simple common-sense polices. And by upgrading the Gi/SGi firewall, ISPs can modernize the perimeter to become adaptable, and use cloud economics for instant intelligence leading to more effective detection.

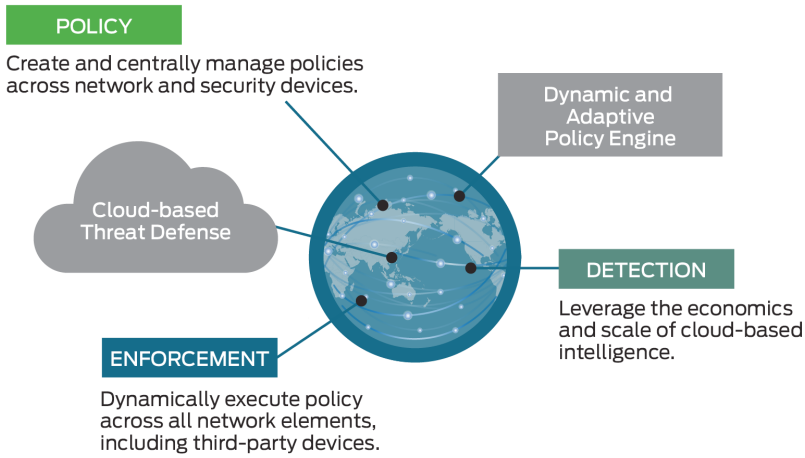

Figure 8.2 Juniper Networks Security Building Blocks

Table 8.1 lists Juniper's comprehensive suite of products and how they centralize and automate security regarding threat intelligence.

Table 8.1 Juniper Networks Building Blocks of Products

Secure Network Areas	Products	What Product Provides	
Dynamic Adaptive and Policy Engine	Security Policy Controller	Policy	Dynamically adapts policy and deploys in real time
Evolved Packet Core	Junos Space Security Director	Policy	
	SRX Series Physical Firewall	Detection, Enforcement	Consistent firewall capabilities, both physical and virtual
	vSRX Virtual Firewall	Detection, Enforcement	
	MX Series Routers with Junos Network Secure (supported with SDSN products)	Detection, Enforcement	
	QFX Series Switches (supported with SDSN products)	Detection, Enforcement	
Juniper Cloud Security	Spotlight Secure Threat Intelligence	Detection, Enforcement	Instant threat intelligence, detection, and defense; controls supply of feeds to inform firewalls
	Sky Advanced Threat Prevention	Detection, Enforcement	
Third-Party Cloud Security Feeds	Other vendors	Detection	

By combining SRX Series physical firewalls, vSRX Series virtual firewalls, Junos Space Security Director, Security Policy Controller, MX Series routers (with Junos Network Secure), and Juniper Networks QFX Series switches, ISPs can deploy a policy engine that communicates with their network, gathers analytics from network data, customizes the user interface to provide data correlation, and uses network edge as detection and enforcement points.

Residential Broadband Edge Security

Today's residential broadband subscribers want personalized online digital experiences, with the freedom to use any device (from phones to tablets to gaming consoles to big-screen TVs) to consume, control, and share content at any location, at any time, and with high quality and resolution levels. With service providers offering IP-based data, video, and voice content to their subscribers on a myriad of devices, and making content easier to access, network security is a very significant issue from the point of view of both a subscriber and a service provider.

For residential subscribers with *always-on* connections, hacker attacks are a significant issue. Hackers can attempt to breach security at odd hours when no one is likely to notice, and these connections often use static IP addresses, thereby making it easy for hackers to consistently return to the site to alter, refine, and exploit their attacks.

Therefore, residential broadband service providers need to meet two major security objectives:

- Defend their own network by shielding network resources and operations against malicious external parties, such as DDoS attacks launched against, and from within, their network (IP address spoofing).
- Protect their subscribers against threats from outside networks.

By mitigating these security threats, ISPs can deliver a reliable and highly available service, which is key to growing their business, retaining the loyalty of subscribers, and maintaining profitability. Figure 8.3 shows a typical residential broadband network topology for IP-based data, video, and voice services.

The MX Series 3D Universal Edge Routers with subscriber management software license provide an industry-leading broadband network gateway (BNG) that supports IPTV/video, broadband Internet, integrated voice, Layer 2/Layer 3 wholesale, and content services, among many other services. BNGs are located at IP edge points where they aggregate, authenticate, and terminate subscribers. Because the BNG acts as an initial service point, service providers primarily

implement Layer 3 security. For downstream devices, ISPs need to carefully consider and implement Layer 2 and transport-related security.

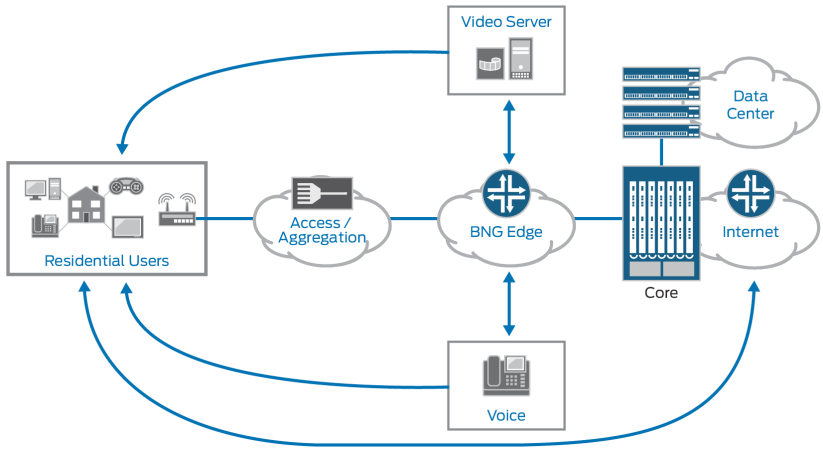

Figure 8.3 Residential Broadband Network Topology

In addition to physical MX Series routers, ISPs can also use Juniper Networks vMX (a virtual version of the MX Series 3D Universal Edge Router that runs as software on x86 servers) as a BNG. Services that were traditionally deployed on a physical MX Series router are easily provisioned and enforced in a vMX environment to deliver a consistent model of visibility and security, and can enable automated provisioning and context-sharing across virtual and physical security platforms.

Business Edge Security

The traditional method for service providers to provision services to business customers on the universal edge is by connecting through secure VPN tunnels or firewalls (refer back to Figure 8.1). ISPs can use Junos VPN Site Secure to configure, set up, and apply security to IPsec VPN tunnels, and the Junos Network Secure software to provide stateful firewall services that integrate with the MX Series router to protect edge services. This configuration eliminates external firewalls that consume router ports and management resources, and can be used as a first line of defense in layered security architectures. ISPs can authenticate users by using methods such as token authentication (RSA SecurID), smart card authentication, machine authentication, and credential provider authentication.

For ISPs, protecting their infrastructure elements and services from targeted attacks is the utmost concern for securing the business edge. Table 8.2 lists several of the most common attacks along with measures to prevent them.

Table 8.2 Common Attacks and Preventative Measures

Attack Type	Purpose	Preventative Measure
Unauthorized access	Attacker attempts to gain access to network infrastructure elements for reconfiguration, traffic redirection, or malicious traffic injection.	• Limit management access, including physical access and network access, to network elements such as loopback filters and allow/deny expressions for user groups. • Disable controls when not in use. • Implement secure communication for management access using SSH for system administrators, SCP (Secure Copy Protocol) to secure file transfers between local and remote host or between two remote hosts. • Hide infrastructure elements from end user.
Software and hardware security flaws within code	Attacker uses flaw to control, reconfigure, and manipulate code.	
Hijacking of management and network control communications	Infrastructure elements, used to run management routines and protocols, are accessed by attacker who then establishes new control and management sessions, or intercepts existing ones, to impact forwarding behavior.	Configure control protocols used in the network, such as IS-IS, OSPF, BGP, and so on, to mitigate hijacking threats.
DDoS	Attacker attempts to deny valid users access to network or server resources by using up all of the resources of the network element or server. DDoS attacks from multiple sources enable a much greater amount of traffic to attack the network.	For VPN and DIA (Direct Internet Access) services, the most vulnerable point of attack is located at the PE (provider edge) to CE (customer edge) interface connection. To mitigate, isolate failure domains by configuring PE to provider interfaces to dedicated line cards, then distribute users between several line cards, and police control plane traffic from each customer.
DIA and IP address spoofing	Attacker attempts to hide the identity of the sender or impersonates another computing system.	Configure unicast reverse path forwarding (RPF) check in strict mode at the PE to CE interface and over all feasible paths to effectively block any traffic from prefixes that are not reachable through the configured interfaces.

Today, many ISPs are deploying hosted private cloud services at the business edge, and transforming themselves into cloud-hosting service providers. According to an IDC study, the worldwide market for hosted private clouds is expected to reach over $40 billion by 2019 (https://www.idc.com/getdoc.jsp?containerId=259066). A hosted private cloud includes a dedicated hosted private cloud and an on-demand private cloud. Hosted private clouds offer a cost-effective and robust solution to quality and reliability issues associated with the public cloud. They provide enhanced security and quality-of-service (QoS) features by bringing the cloud platform behind the firewall. Figure 8.4 shows the Juniper Networks universal edge solution for hosted private cloud over VPN.

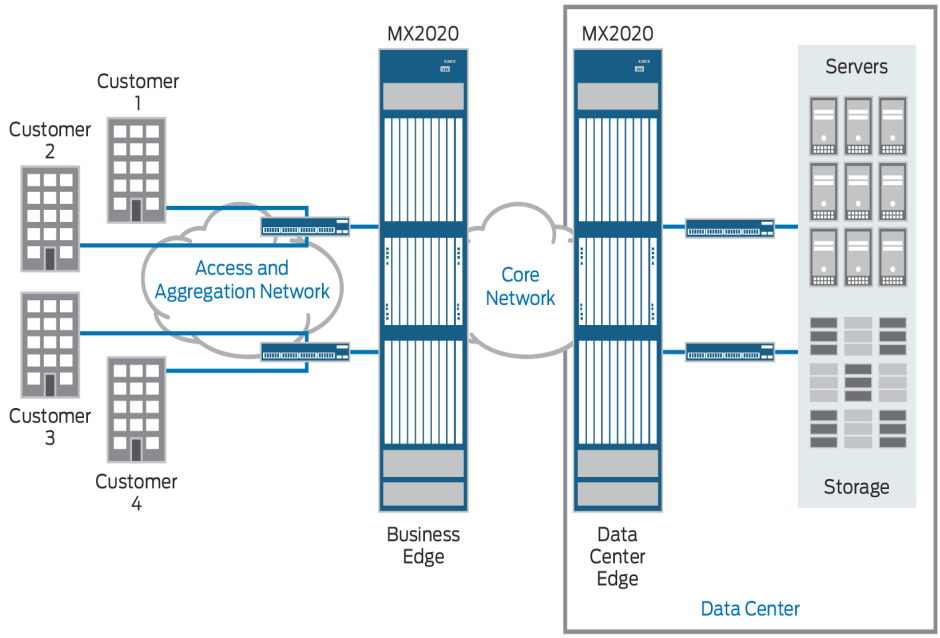

Figure 8.4 Universal Edge Solution for Hosted Private Cloud

The universal edge solution seamlessly extends customer VPNs to cloud data centers. Cloud customer security is extended to virtual machines. The MX Series routers in a Virtual Chassis configuration provide collapsed layers of data center routing, switching, and security devices. As more and more customers embrace the agility and flexibility of automated provisioning for virtual private resources, the on-demand private cloud is trending toward becoming the standard operating model for hosted private clouds.

Figure 8.5 shows an example of a service provider offering virtual security services (DDoS and URL blocking) on demand to their customers by using a virtual high-end provider edge (vHEPE) router to connect the ingress (incoming) PE device and the URL blocking device (scrubber).

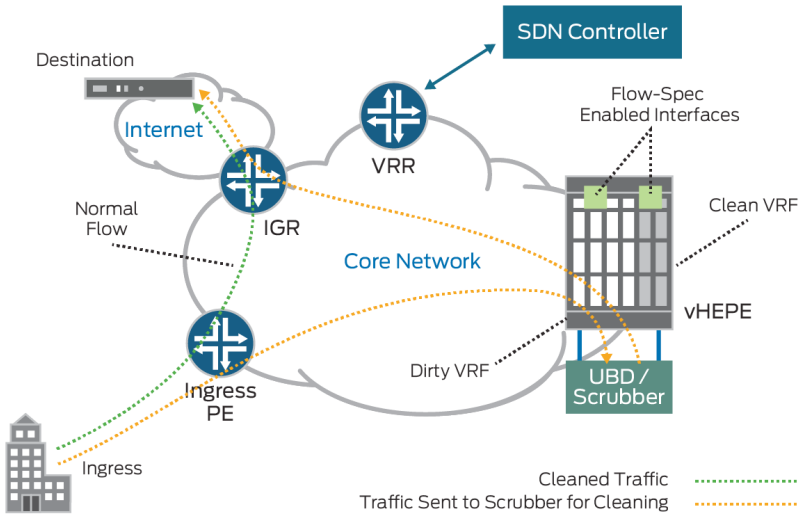

Figure 8.5 Virtual Security Services – DDoS and URL Blocking

The condition, or origin, of the URL request traffic coming from the PE device to the HEPE device, determines the traffic's destination. If the traffic:

- Is in a dirty virtual routing and forwarding (VRF) table, then the scrubber device begins filtering.
- Belongs to a blocking list, then some blocking messages are sent out.
- Belongs to an "allowed list" blocking list, then the traffic is first sent to the clean VRF table, and then forwarded to its original destination.

Similarly, for DDoS, if a server is under attack, dirty traffic with known signatures is dropped, while clean traffic is forwarded to its original destination.

In Figure 8.5, the vMX plays the role of a virtual HEPE, diverting traffic to the virtualized scrubber applications, which can be aligned with the virtual scrubber device to leverage the same cloud infrastructure and orchestration tool.

Telecom Cloud Providers and Security

By 2020, nearly 40% of the information in the digital universe will be touched by cloud computing providers (http://preservationmatters.blogspot.com/2014/05/the-digital-universe-in-2020-big-data.html). Telecom cloud providers provide cloud services (or Web services) over the Internet. The most common cloud service resources, collectively known by the acronym SPI, are:

- Software as a Service (SaaS) – A software distribution model designed for customers where applications are hosted by a service provider and made available to customers over the Internet. This represents the end result of cloud computing.
- Platform as a Service (PaaS) – Provides hardware and software tools, as well as operating services delivered over the Internet without downloads or installation, that can be used as a service by IT personnel to develop applications for end users.
- Infrastructure as a Service (IaaS) – Similar to PaaS providers; however, it hosts the infrastructure equipment used to support operations (such as storage, hardware, software) on behalf of its customers, providing IT personnel more control over the OS of the business.

Threats to cloud platforms are ongoing and affect everyone who uses cloud services for business or personal reasons. Multilayer, multidomain security is critical for a telecom cloud deployment. Table 8.3 lists some of the biggest security issues and vulnerabilities specific to cloud computing and telecom cloud providers today.

Table 8.3 Cloud Computing Security Issues

Security Issue	Description/Example	Preventative/Corrective Measure
Data breaches	An organization's sensitive internal data is accessed as a result of malicious and intrusive action. For example, a new data breach called side channel timing exposure allows a user on one VM to listen for activity that signals the arrival of an encryption key on another VM on the same host.	Implement secure hypervisor and VM operations.
Data loss or leakage	When a disk drive dies without the customer having a backup drive, and the customer has lost the key that unlocks their encrypted data.	Implement disk-level data protection using RAID, backup/replicated backup, data replication, and journaled/checkpoint-based replication.

Security Issue	Description/Example	Preventative/Corrective Measure
Account/service traffic hijacking	Phishing, loss of passwords, and compromised credentials can lead to loss of control of an account. If an account in the cloud is hijacked, an attacker can use it as a base and take advantage of a customer's reputation to enhance himself/herself at the customer's expense.	Practice defense in depth, prohibit sharing of account credentials, and implement strong two-factor authentication.
Insecure interfaces and APIs	As layers are added to APIs to reach value-added services, increasing complexity might create additional exposure allowing attackers to circumvent policy.	Implement OAuth, an open authorization service for Web services that controls third-party access and is an Internet Engineering Task Force (IETF) standard.
Malicious insiders	A system is still vulnerable to inside attackers if encryption keys are not stored with the customer and are available only at data usage time.	Keep encryption keys on your own premises, not in the cloud.
Cloud services abuse	Attackers take advantage of a collection of a large number of servers to manipulate existing cloud services to launch DDoS attacks or add malware.	Plan how to detect inappropriate use, clearly define what constitutes abuse, and determine how to prevent it in the future.
Shared technology	A misconfigured operating system or application within the cloud's shared infrastructure can lead attackers to compromise data beyond a customer's immediate surroundings.	Institute an in-depth defensive and monitoring strategy for errant/destructive behaviors relating to storage, network, application, and user access.

However, in spite of the security threats, cloud security is being promoted as a benefit instead of an unknown and risky hindrance for cloud adoption. Telecom cloud providers must adopt a defense-in-depth strategy to ensure, protect, and secure data by implementing layers of security technologies and business practices. Service providers are building security into the platform at the IaaS level, similar to the way that customers expect power to be built into the platform.

Juniper Networks high-end SRX Series firewalls provide front-end security to telecom cloud data centers around the world. Combined with Juniper's AppSecure and UTM capabilities, the SRX Series provides an effective barrier between the outside world and sensitive telecom cloud applications and services. Within the data center, Juniper provides multiple layers of security-focused virtualized network functions (VNFs), such as the vSRX.

For example, public cloud service providers (which can host large numbers of VMs for their customers) can deploy the vSRX to protect their customers by placing the virtual firewall in front of each customer's individual hosting environment, keeping the hosting environments separate from each other, as shown in Figure 8.6.

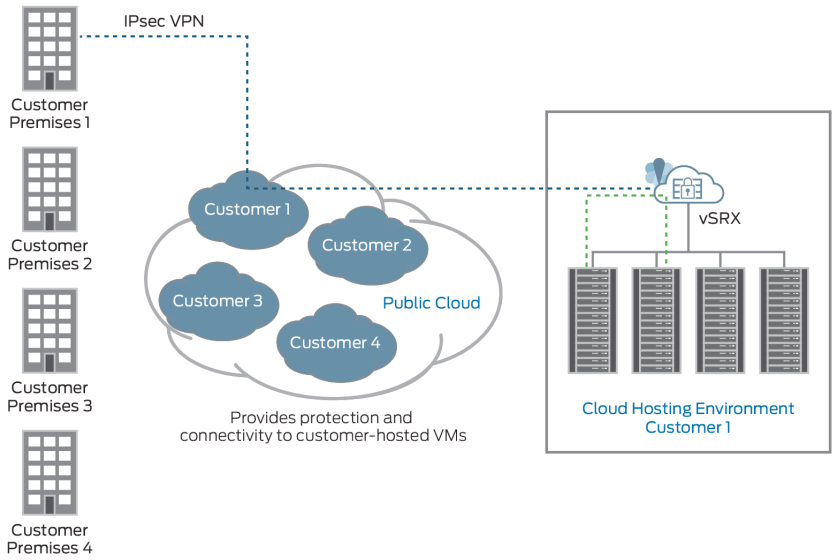

Figure 8.6 Public Cloud Service Provider

By combining Junos Space products (such as Security Director and Security Policy Controller), Juniper Networks Secure Analytics virtual appliances, and VMware products in the cloud infrastructure, along with virtualizing specific objects with vSRX, public cloud service providers can tailor security between network elements and offer security services to customers with multiple remote locations.

Traditional cloud architectures rely on relatively unsecure VLANs, where rogue or compromised network elements can impact other workloads. However, using the Contrail Cloud Platform to provide IP VPN (or IPsec VPN) connectivity between virtual objects, ensures that workloads and VMs can communicate only with predetermined network elements, resulting in a more highly secure virtualized network.

With the vSRX, public cloud service providers can provide their customers with the security required, both inside their virtualized data centers, and at the customer and business network edges.

Summary

ISPs face a wide range of security issues originating from a multitude of known and unknown sources. ISPs must ensure that their infrastructure is secure and that their customers' data and applications are protected. To protect the growing digital universe, ISPs must devise and follow a thorough, multilayered, and defense-in-depth approach to security by considering all information traversing the network and in the cloud, and not just threats solely identified at the perimeter and edge. A change in mindset is required – it's not *network security* anymore, it's *secure networks*!

Resources and References

- The following book introduces you to all the fundamentals of the Juniper Networks Dynamic Subscriber Management solution and shows you how to get it up and running in a day: http://www.juniper.net/us/en/training/jnbooks/day-one/networking-technologies-series/dynamic-subscriber-management/.
- This technical library includes everything you need to understand and configure for all aspects of Junos OS Broadband Subscriber Management and Services: http://www.juniper.net/techpubs/en_US/junos16.1/information-products/pathway-pages/subscriber-access/index.html.
- Eliminate threats at the network edge with industry-leading scalability and performance. Service providers can securely scale services while converging edge routing, switching, and security in a single MX Series 3D Universal Edge Router: http://www.juniper.net/us/en/products-services/network-edge-services/security/.
- The Juniper-sponsored J-Net Communities forum is dedicated to sharing information, best practices, and questions about Juniper products, technologies, and solutions. Register to participate at this free forum: http://forums.juniper.net/jnet.