

SECURITY IN THE NEXT- GENERATION DATA CENTER

Key Strategies for Long-Term Success

Table of Contents

Executive Summary	3
Introduction	3
The Strategic Security Imperatives of the Next-Generation Data Center	3
Scale	4
More Intra-Server Traffic	4
Support for Increased Bandwidth Usage	4
More Deployments	4
Increased Processing and Intelligence	4
Visibility and Context	4
Application-Level Visibility	5
Real-Time Context	5
Business Visibility	5
Visibility and Contextual Capabilities—Near Term Requirements	6
Intelligent Enforcement	6
Policy Enforcement in Virtualized Environments	6
Granular Policy Enforcement	6
Deployment and Enforcement Flexibility	7
Juniper Networks Security Solutions for the Next-Generation Data Center	7
Conclusion	7
About Juniper Networks	8

Executive Summary

In the data center today, several major trends are happening in parallel, with each representing a fundamental change in terms of how IT is managed. For the security teams responsible for safeguarding corporate IT assets, these trends present a host of challenges, necessitating several new capabilities and approaches to ensure ongoing, effective security. This white paper examines these trends, and it reveals the key capabilities that today's security teams require to effectively ensure that vital corporate assets remain secure, while at the same time optimizing access, cost, and administrative efficiency.

Introduction

Today's data center architectures are in the crosshairs of several significant technology trends, with each presenting fundamental security implications:

- **Large-scale consolidation.** To maximize economies of scale and cost efficiencies, enterprises continue to consolidate data centers. Consequently, extremely large data centers are increasingly the norm. This concentration of computing, storage, and networking is creating unprecedented scale requirements for network security.
- **Virtualization.** The nature of computing inside the data center has fundamentally changed, with workloads increasingly moving from dedicated physical servers to multiple virtual machines. As a result, a typical application workload is now completely mobile—it can be instantiated anywhere in the data center, and it can even be moved from one physical server to another while running. Furthermore, the increasing trend of desktop virtualization means that any number of clients can access a virtual desktop hosted on a server located in the data center. And, most importantly, virtual machines running on a single server communicate via an internal virtual switch (vSwitch). This has fundamental implications for traditional network security architectures, which were not designed with a focus on intra-server traffic.
- **Service-oriented architectures and application mashups.** Application architectures are evolving from being relatively monolithic to being highly componentized. The componentized application architectures emerging today allow for more reuse and, given that each component can be scaled separately, provide better scalability. One result of this emerging trend is that there is starting to be more “east-west” traffic (between components in the data center) than “north-south” traffic (between servers inside the data center and points outside the data center). This application evolution effectively acts as a traffic multiplier and can expose additional areas of vulnerability—further pushing the scale requirements of security mechanisms in the data center.
- **Fabric architectures.** Both in reaction to the above trends, and in an ongoing effort to realize improvements in administrative efficiency, network scalability and availability, and infrastructure agility, enterprises are increasingly looking to adopt fabric architectures, which enable many physical networking devices to be interconnected so that they are managed and behave as one logical device. Network security infrastructures will correspondingly need to adapt to the management and integration implications of these architectures.

These trends represent the characteristics of the next-generation data center, which will require a fundamental re-imagining of how security gets implemented. Consequently, traditional security approaches—which were characterized by a focus on relatively static patterns of communication, the network perimeter, and the north-south axis—will no longer suffice.

The Strategic Security Imperatives of the Next-Generation Data Center

Given the technology trends at play in the next-generation data center, effective network security approaches will need to address the following requirements:

- **Scale.** Network security will need to scale to accommodate increasing traffic, more processing-intensive intelligence to combat increasingly sophisticated threats, and more deployment options and scenarios.
- **Visibility.** To be effective, network security solutions will need to have more contextual visibility into relevant traffic.
- **Intelligent enforcement.** Security teams will need capabilities for efficiently enforcing policies on both physical and virtualized workloads.

These core requirements are detailed in the sections below.

Scale

The next-generation data center presents fundamental implications for the scalability of network security. Following is an overview of the areas in which this need for scale will be most evident.

More Intra-Server Traffic

In the wake of increased virtualization, industry experts estimate that network traffic will increasingly be comprised of traffic between two servers, as opposed to traffic between clients and servers. The decomposition of data center applications into a mashup of reusable components will also hasten this increase in server-to-server communications inside the next-generation data center. In fact, analysts estimate between 2010 and 2013, server-to-server traffic will grow from 5% to 75% of network traffic. Another implication of these trends is that enterprise security teams can expect that every gigabit of capacity entering the next-generation data center via the north-south axis will typically require 10 gigabits of network capacity on the east-west axis, and could scale up significantly from there. The increasing volumes and increasing dynamism of this server-to-server traffic will create a corresponding increase in the criticality of network security mechanisms, and their need to scale.

Support for Increased Bandwidth Usage

The traditional demands of perimeter protection will not go away in the next-generation data center. Plus, given the concentration of services provided by a large next-generation data center, the aggregate bandwidth connecting the data center to the Internet will commonly be measured in gigabits per second. These trends will only be exacerbated by the increased prevalence of rich media applications and their associated network bandwidth requirements.

More Deployments

Emerging trends will increase the usage of network security devices. For example, given extensive server virtualization within the next-generation data center, physical isolation can no longer be relied upon to ensure the separation of groups of applications or users, which is a requirement of many compliance mandates and security policies. Consequently, more network security mechanisms will be required to supply this isolation.

Given the addition of virtual desktops to the next-generation data center, campus and branch perimeter security mechanisms such as Web security now need to be delivered inside the data center—and with high levels of scalability.

Increased Processing and Intelligence

The increased sophistication of attacks means that the computing power and memory needed to secure each session entering the next-generation data center will expand substantially. Further, these sophisticated threats will also make it too difficult for perimeter security alone to determine all of the potential downstream effects of every transaction encountered at the perimeter. As a result, security teams will need to deploy additional, specialized network security services—for example, security specifically for Web services, XML, and SQL—closer to the enterprise's most valuable assets or "crown jewels."

Visibility and Context

Broadly speaking, there are three kinds of network security products deployed in the next-generation data center:

- Proactive—including security risk management, vulnerability scanning, compliance checking, and more
- Real time—featuring firewalls, intrusion prevention systems (IPS), Web application firewalls (WAFs), anti-malware, and so on
- Reactive—including logging, forensics, and security information and event management (SIEM)

Across the board, rather than operating solely based on IP addresses and ports, these solutions need to be informed by a richer set of contextual elements to enhance security visibility.

Application-Level Visibility

Visibility into applications is particularly vital—and challenging. Without visibility into what applications are actually present on a network, it is difficult to envision an effective security policy, let alone implement one. While in theory, data centers may be viewed as highly controlled environments to which no application can be added without explicit approval of the security team, operational realities often mean that the security team has limited visibility into all of the applications and protocols present on the data center network at any given time. Following are a few common scenarios that organizations contend with on a day-to-day basis:

- Application teams build virtual machine images with a particular application in mind, but the virtual machine templates they work with include extraneous daemons that send and receive packets on the network.
- The process for green lighting new applications is insufficiently policed and enforced, so the security team often finds out about applications well after they have been deployed.
- In virtual desktop environments, employees use their virtual desktops to access applications outside the data center, with the applications in use evolving on a continuous basis.

Given these realities, the next-generation data center security architecture will need to possess the capability to see and factor in the actual application being used (in lieu of TCP port 80).

Real-Time Context

Acquiring context in real time presents significant challenges. For example, for a firewall to be effective in the next-generation data center, it needs to determine the following contextual aspects:

- The identity of the user whose machine initiated the connection
- Whether the user is connecting with a smartphone, laptop, tablet PC, or other device
- The software—including OS, patch level, and the presence or absence of security software—on the user's device
- Whether the user is connecting from a wireless or wired network, from within corporate facilities, or from a coffee shop, airport, or some other public location
- The geographic location from which the user is connecting
- The application with which the user is trying to connect
- The transaction the user is requesting from the application
- The target virtual machine image to which the user's request is going
- The software—including the OS, patch level, and more—installed on the target virtual machine

Some of this context may be derived purely from the processing of packets that make up a session. For example, a WAF may be able to map the URL being requested to a specific application, or a hypervisor may be able to provide specific context about communications between virtual machines. However, other contextual information, such as information about the type of OS on the source and destination hosts, will need to be acquired out of band.

Business Visibility

Many aspects of business context might also affect security policy decisions. For example, policies may be contingent upon whether a service request is being made in relation to an end of quarter sale or a disaster recovery response. Clearly, stitching together the sequence of events required to identify this business context may prove very complex, but certain shortcuts may be possible to reduce some attack surfaces. For example, IT teams could use a global indicator in the data center that signals when disaster recovery is in progress and only permit certain actions—such as wholesale dumping or restoring of database tables—during those times.

Visibility and Contextual Capabilities—Near Term Requirements

The contextual visibility outlined above provides security teams with an overall view of what's going on in their network and allows them to set policies that mitigate risk and align the data center's risk profile with business requirements.

While acquiring some of these forms of context may not be possible immediately, there are some near term, must-have requirements. For example, to simply return to traditional levels of control, security administrators must be able to map virtual machine instances to IP addresses in virtualized environments. Any network security solution that cannot bridge this gap risks irrelevance in the next-generation data center.

Intelligent Enforcement

Enforcement requires that security teams can apply specific sets of security services to communications between groups or pools of resources, often virtual machines, in the next-generation data center. This requires mechanisms for ensuring that effective, timely control can be exerted. For example, this could mean that the traffic in question would have to traverse a security enforcement point that will have the necessary context and that will have been configured with the required policies.

For north-south traffic, it is fairly straightforward for perimeter firewalls to apply the necessary enforcement capabilities. For east-west traffic, enforcement isn't a trivial challenge. First, if enforcement policies neglect the virtual layer, threats can go undetected. If virtual-level security is applied in a way that is independent of physical-layer enforcement, businesses can still be vulnerable to gaps. Further, implementing physical-layer enforcement through a solution that is independent of security mechanisms operating at the virtual layer creates considerable complexity for administrators—complexity that can result in inefficiency and errors, and can also open the door to threats.

Policy Enforcement in Virtualized Environments

Traditionally, policy selection has been associated with the source and destination of traffic. For example, in the case of non-virtualized workloads, the switch port can authoritatively identify the source of a given traffic flow.

For virtualized workloads, however, the identification of the workload source presents a more dynamic challenge. Communication within a group of virtual machines on the same physical host can occur freely, though having visibility into this traffic may be required according to some security requirements. Within these virtualized environments, identifying the source and destination of traffic, mapping that traffic to specific policies, and ensuring that enforcement points execute the policies required can pose significant challenges.

To be effective, network security mechanisms need to be able to associate policies with groups of virtual machines, and consistently and accurately execute on those policies. To do so, security teams will need capabilities for supporting the virtualization technologies employed within the next-generation data center. In VMware environments, this requires integration with vCenter, which is used to create and manage groups of virtual machines. This integration is essential to enabling security teams to manage and monitor policies through a central console. Further, given the scalability demands of the next-generation data center, this central management infrastructure needs to have capabilities for aggregating information from multiple vCenter instances and from the physical security infrastructure, in order to maximize administrative efficiency.

Granular Policy Enforcement

To be effective, a security enforcement point needs to have the required visibility into the traffic to which policies need to be applied. Techniques such as VLAN partitioning can be used to ensure that a physical security appliance inspects all traffic crossing a security trust boundary, even in the case of virtual machine to virtual machine traffic that is occurring on the same physical host. However, this approach is suboptimal for two reasons:

1. In order to institute the requisite policy enforcement points, IT teams need to change the networking architecture, which requires tight collaboration between networking and security teams.
2. It represents a coarse-grained approach in which all traffic between VLANs has to be routed to a separate physical appliance before being routed to the destination VLAN. This approach doesn't enable finer grained filtering, so that only a subset of traffic gets routed to the physical security appliance. Further, even if such a capability were available, there would be no way for the physical security appliance to avoid having to process all of the packets in a given flow when it wants to implement a simple "permit" firewall rule.

For virtualized workloads, the use of protocols like Virtual Ethernet Port Aggregator (VEPA)—in conjunction with the switching infrastructure's ability to forward traffic from one virtual port to another virtual port through an intermediate physical appliance—can overcome the first hurdle, but this still remains a coarse-grained mechanism.

On the other hand, the performance and administrative overhead of deploying firewalls on each virtual machine can also prove problematic. To optimize the control of traffic within virtualized environments, security teams need a flow table and basic flow integrity checks that reside closer to the ingress of packets. In these virtualized environments, the closest place to the ingress virtual machine is the hypervisor which hosts the vSwitch.

Deployment and Enforcement Flexibility

Given the breadth of architectural choices in designing the next-generation data center, security teams need to have maximum flexibility in terms of where they deploy and apply security services.

For example, given that many physical network security appliances would not have the visibility required to inspect certain traffic between virtual machines, security teams need solutions that include the option of enforcing security in a hypervisor. However, in some cases, instead of or in addition to hypervisor-based security enforcement, security teams may also want to deploy security mechanisms on virtual or physical appliances which are logically inline or to which flows are specifically forwarded by the data center switching infrastructure.

It is also a critical requirement for organizations to be able to dynamically shift workload across hardware appliances and virtual machines. To realize this capability, security teams need a common language for defining security policy and selecting one of a collection of policies to apply to specific network traffic—regardless of virtual or physical security enforcement point. This is vital for reducing implementation complexity and minimizing total cost of ownership.

Juniper Networks Security Solutions for the Next-Generation Data Center

Juniper Networks offers comprehensive, flexible, and robust security solutions that are ideally tailored for the requirements of the next-generation data center. Juniper solutions offer an unparalleled combination of features, and they address these fundamental requirements:

- **Scale.** Juniper solutions offer the processing scalability required to combat today's sophisticated threats. With the Juniper Networks® SRX Series Services Gateways, Juniper offers firewalls that scale, both in terms of throughput and in the number of concurrent sessions, so they can accommodate the intensive security and bandwidth requirements of the next-generation data center.
- **Fabric architecture support.** Juniper provides switches for fabric architectures, and is leading the way in delivering cohesive security integration in these environments. The SRX Series product line is designed to work with the scalable Juniper fabric, ensuring optimal flexibility, scalability, and security in these emerging environments.
- **Visibility.** Juniper security solutions can incorporate the required context, in real time, to ensure effective, continuous policy enforcement. With AppSecure, Juniper provides a deep understanding and visibility into application behaviors and weaknesses to prevent application-borne threats that are otherwise difficult to detect, let alone stop. In addition, Juniper Networks Unified Access Control allows organizations to integrate user identities into their security policies. Juniper also offers hypervisor integration that enables security teams to incorporate visibility into virtualized components in their security mechanisms.
- **Intelligent enforcement.** Finally, Juniper offers solutions that can be deployed both on physical devices and in the hypervisor, enabling policy management and enforcement across the entire next-generation data center. As result, organizations realize stronger overall security, and security teams can implement faster configuration changes and reduce the errors that can create security exposures.

Conclusion

To achieve and sustain the required levels of security within the next-generation data center, security teams will need to address a number of critically important mandates that include the need for scale, context, and intelligent enforcement. Juniper Networks provides organizations with a host of integrated solutions that deliver on these key requirements, ensuring that enterprises can effectively, efficiently, and continuously address their security objectives.

About Juniper Networks

Juniper Networks is in the business of network innovation. From devices to data centers, from consumers to cloud providers, Juniper Networks delivers the software, silicon and systems that transform the experience and economics of networking. The company serves customers and partners worldwide. Additional information can be found at www.juniper.net.

Corporate and Sales Headquarters

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or 408.745.2000
Fax: 408.745.2100
www.juniper.net

APAC Headquarters

Juniper Networks (Hong Kong)
26/F, Cityplaza One
1111 King's Road
Taikoo Shing, Hong Kong
Phone: 852.2332.3636
Fax: 852.2574.7803

EMEA Headquarters

Juniper Networks Ireland
Airside Business Park
Swords, County Dublin, Ireland
Phone: 35.31.8903.600
EMEA Sales: 00800.4586.4737
Fax: 35.31.8903.601

To purchase Juniper Networks solutions, please contact your Juniper Networks representative at 1-866-298-6428 or authorized reseller.

Copyright 2011 Juniper Networks, Inc. All rights reserved. Juniper Networks, the Juniper Networks logo, Junos, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.