SECURITY CONSIDERATIONS FOR CLOUD-READY DATA CENTERS

Table of Contents

Table of Figures

Executive Summary

Data centers have evolved significantly as organizations consolidate servers, applications, and other resources, and as they adopt new technologies as a means to reduce costs and increase efficiency. Technologies such as server virtualization, distributed application tools, and IP-based storage are helping organizations maximize their data center resources, while at the same time making it more difficult to protect these critical assets.

In addition to cyber theft and increasing levels of malware, organizations must guard against new vulnerabilities introduced by data center technologies themselves. To date, security in the data center has been applied primarily at the perimeter and server levels. However, this approach isn't comprehensive enough to protect information and resources in new system architectures.

To effectively manage the new risks, organizations should reevaluate their data center security practices and implement new network-centric capabilities to ensure the integrity of their services. Because the network touches every device in the data center, it is an ideal location for security. A network-centric approach to providing security in the data center delivers benefits such as scalability, unified security policy definition and enforcement, visibility into application traffic, and reduced operations overhead.

Introduction

Data centers are evolving quickly in response to operational pressures and technology innovations. To reduce costs and gain flexibility, organizations are consolidating data centers and adopting new technologies ranging from virtualization to new application architectures and cloud computing.

The current economic environment is accelerating these trends. With data center consolidation, organizations can achieve efficiencies by moving "high touch" systems from satellite offices to either central data centers or third-party cloud providers. Consolidating servers, applications, and other resources leads to higher utilization of these resources and eliminates the need for IT staff in many locations. Outsourcing some applications to the cloud can make it easier to support teleworkers and employees in remote offices who need secure access to centralized resources that are always available.

At the same time, new collaboration tools, including telephony presence, instant messaging, wikis, blogs, and social networking, are bringing employees closer despite physical distance. In addition, use of service-oriented architecture (SOA) and other distributed approaches to application development are resulting in highly distributed applications.

These trends are having a major impact on data center architectures and the problems that need to be addressed to provide adequate security for data and systems residing in them. In traditional data center models, applications, compute resources, and networks have been tightly coupled, with all communications gated by security devices at key choke points. However, technologies such as server virtualization and Web services eliminate this coupling and create a mesh of interactions between systems that create subtle and significant new security risks.

To secure modern data center applications, organizations need comprehensive, scalable and elastic security tools in the network that combine application fluency and identity-based controls with centralized policy and compliance management. Understanding these new security challenges is key to implementing appropriate security solutions for the virtualized, cloud-ready environment.

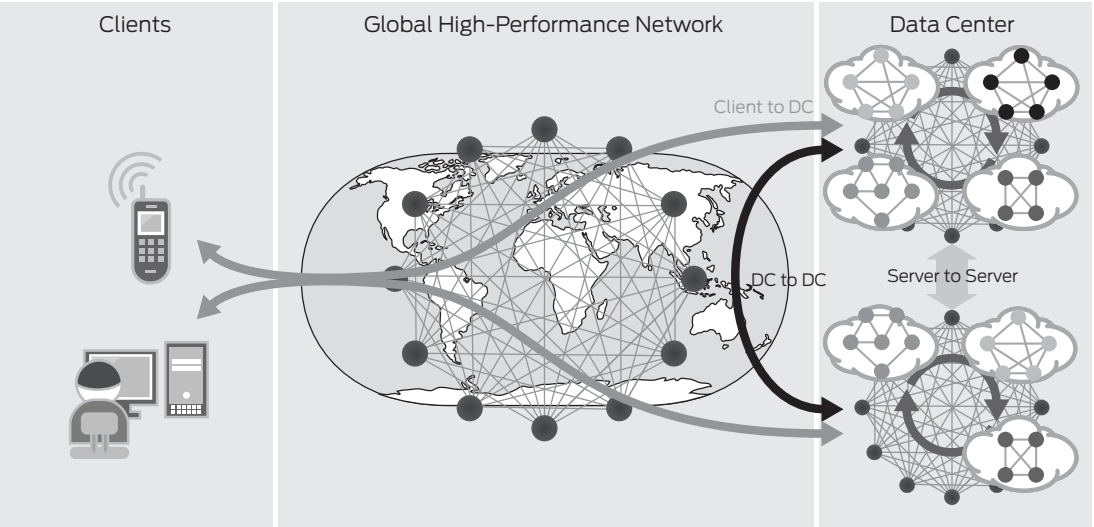

Figure 1: The increasing risk in the new data center infrastructure

New Security Challenges in the Data Center

In information networks and systems, security controls address the confidentiality, integrity, and availability of information. As organizations consolidate applications, data, and other resources within a few large data centers, they increase the risk that a single system breach represents. Whereas a single server has conventionally housed one application, virtualized servers today host multiple applications or components or both. The compromise of a single physical server can affect numerous applications and a large number of users.

Increased presence of four new technology elements in current data center networks pose significant new security challenges that must be addressed:

- Server virtualization
- Distributed application architectures
- IP-based storage networks
- Sophisticated denial of service attacks with botnets

Virtualization

Server virtualization has been key to data center consolidation, allowing enterprises to squeeze the most out of server resources, and reduce the need for floor space, electricity, and cooling. However, the difficulty of monitoring and controlling what goes on inside a virtual machine (VM) and between virtual machines presents a new and significant security risk.

Virtual server technologies enable multiple OS instances to run on a single host machine. Each of these operating systems has its own “virtual CPU,” memory, and I/O resources, creating a virtual machine. More recently, server virtualization has been extended to include the ability to run multiple OS instances over a cluster of hosts. As a result, a pool of resources beyond the finite capacity of a single server is available for applications. In addition, virtual machines can migrate from host to host and be replicated across hosts, allowing for dynamic resource allocation as demand rises and falls.

To facilitate communication between virtual machines on the same host, vendors of this technology have developed software that simulates an Ethernet switch. A virtual switch runs on each physical host and allows connectivity between guest VMs. Within this internal server network, some VMs may be connected to the same broadcast domain with other VMs connected to disparate broadcast domains.

Switches in VM hosts cannot be monitored by the same devices used to monitor the physical network … traffic passing back and forth between VMs on a single server does not travel out into the rest of the data center network and thus cannot be seen by regular network-based security platforms. This lack of visibility into and control over intra-server virtual machine communication creates significant security risk for the organization. The lack of enforcement points inside the virtualized domain allows worms and other malicious traffic to propagate unchecked between virtual machines and potentially onto the physical data center network. Such lack of visibility also presents compliance problems. Organizations need a new set of security tools and best practices to secure this new infrastructure tier.

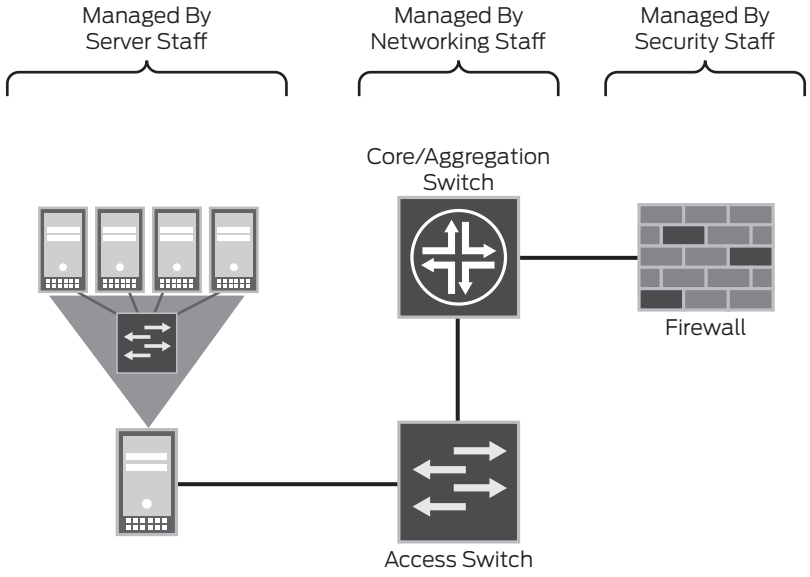

Figure 2: Virtual server networking

Distributed Application Architectures

There has been a major shift away from monolithic application development and toward distributed application architectures, which rely on common, reusable sub-components. While these architectures make application development faster and more efficient, they pose security risks by creating highly distributed communication patterns, with multiple flows per transaction.

For example, many Internet applications are mashups in which different elements are delivered by different servers. An "igoogle" session, for instance, requires accessing multiple specialized applications running on different servers, with the result aggregated and delivered to the user's browser. Many of these applications are built with a service-oriented architecture. SOA separates functions into distinct units, or services, which developers make accessible over the network so that they can be combined and reused in the production of applications. These services communicate with each other by passing data from one service to another, or by coordinating an activity between two or more services.

As a result, each mashup application front-end typically connects to a large set of back-end applications, creating a fan out hierarchy per user session with a high number of TCP connections per client interaction. Securing this type of highly distributed environment requires the ability to correlate all of these flows with one user performing one transaction. User identity and credentials must be sent and shared among the different mashup applications and back-end services, so that the user information is presented consistently and appropriate access is granted to each application element based on user credentials.

Despite their many advantages, mashups, SOA, and similar distributed application technologies make it difficult to enforce access entitlements. Data privacy is another security issue in highly distributed application environments. Since client communications are now targeted at a larger set of systems, the possibility of an eavesdropper intercepting a communication stream increases, making encryption a requirement for communication.

The use of XML within SOA-based applications presents another security challenge. XML, is characterized by large data sets and hence packets. Transporting these large data sets (and packets) between servers requires multiple, high bandwidth TCP sessions, each of which must be inspected by a security enforcement system. Aggregative, throughout the datacenter the overall bandwidth required from the firewalls increases significantly.

While distributed application architectures have major advantages allowing for rapid application deployment and reuse of application components, they present unique requirements for a scalable, high-performance security solution capable of tying application flows to specific users.

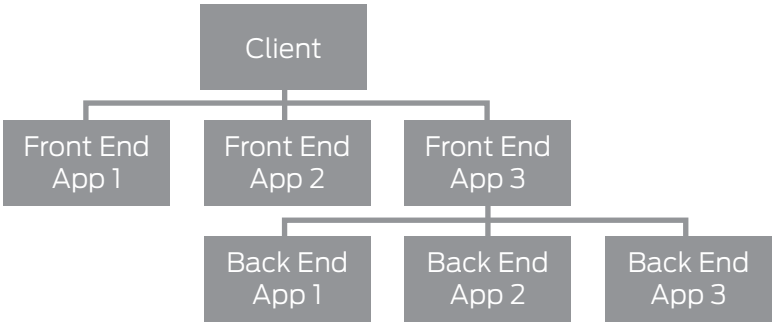

Figure 3: Mashup and SOA fan out effect

Storage over IP Networks

For economic reasons, many organizations are deploying IP-based storage technologies. Leveraging a common IP infrastructure to support data, voice, video, and storage yields economies of scale and simplifies operations. In addition, network attached storage (NAS) technologies such as Network File System (NFS) and Common Internet File System (CIFS), along with storage area network (SAN) technologies such as iSCSI, enable organizations to deploy large storage pools, which require fewer network interconnections and less management than is otherwise necessary. Organizations benefit from the ability to predict storage needs more accurately and to better manage the scaling of storage arrays.

Despite its flexibility and cost benefits, deploying storage over IP networks presents a heightened security risk. Because it is part of the overall IP infrastructure, the storage network must be protected from denial of service (DoS) and other malware attacks that can result in critical data being unavailable or corrupted, and applications not functioning properly. Data availability can also be compromised by contention on the IP network; latency- and loss-sensitive storage traffic may be crowded out by high volumes of data or video, for example.

Privilege escalation can also occur, since storage arrays are accessible to a variety of applications. As new applications are brought online, it is not uncommon for business critical and non business critical data to be stored on the same array. While storage vendors provide tools to secure storage access on the array itself (data at rest), these tools do not address the vulnerabilities of data flowing across the network (data in motion). For enterprises to achieve the benefits of consolidating storage on IP-based networks, they need the ability to ensure data availability, integrity, and confidentiality across the data center's IP infrastructure.

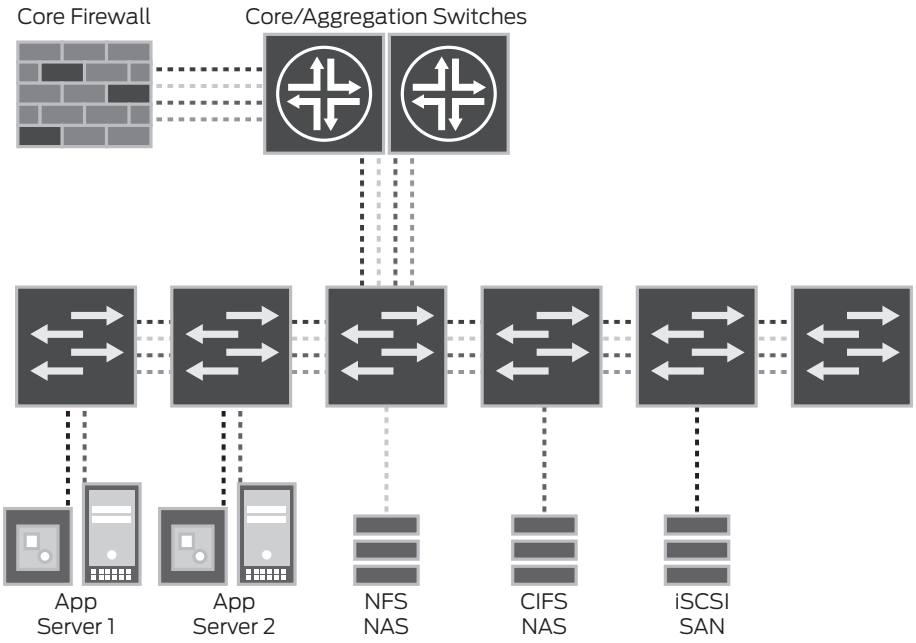

Figure 4: Storage arrays shared by multiple applications

Sophisticated DDOS attacks through botnets

The growing frequency and sophistication of cyber attacks poses a serious, ongoing security challenge. In the past, rogue hackers randomly targeted systems in order to establish credibility in the "black hat" community. Today, however, the combination of browser-based "cloud" computing, mobile data platforms, and social networking have given rise to a new breed of threat: stealthy Web-borne malware that forms into highly organized botnets that open a callback channel out of the network to expose confidential data. Browser exploits now account for 65% of all client-side attacks. Cyber criminals embed malicious code within user-generated content websites, third party ads, and high-traffic web applications to compromise servers and clients. The damage is increasingly visible as data breach disclosures become mandatory. In recent years, there has been a steady stream of disclosures from both the private and public sectors. Traditional security mechanisms do not provide a sufficient defense against a cyber criminal who attacks on multiple fronts, from OS exploits, browser attacks, and increasingly, plug-in/widget vulnerabilities.

Security Requirements for the Cloud-Ready Data Center

To effectively manage the risks resulting from new technologies in the distributed, virtualized systems of today's data center infrastructures, organizations must reevaluate their practices and implement new security solutions. Most security in data centers has been applied at the server level, by installing host-based intrusion detection, identity enforcement, antivirus, and other software agents. This approach, however, isn't scalable, doesn't encompass the range of network-attached devices, and presents major operational challenges.

Organizations need a unified security layer operating across the heterogeneous and ever changing data center infrastructure. The network is ideally suited to provide visibility into the application traffic it is carrying, and to act as an insertion point for policy enforcement devices. Managers should look for network-centric security solutions with the following characteristics.

Application Fluency

To protect data center assets, network security products such as enforcement gateways, firewalls, and monitoring systems must be intelligent enough to identify application context, and conversations. From a policy definition and security enforcement perspective, the classic TCP/IP session "5 tuple" is no longer sufficient to define or enforce business security policies.

Enterprises need application fluent security products that allow them to precisely define what actions are allowed within certain application instances. Additionally, such solutions must provide visibility into the application infrastructure, making it possible to determine application usage profiles and other valuable application-level information.

As mentioned earlier, in today's data center, multiple application instances may run on a shared host or pool of resources. To differentiate among these applications, some organizations run them on separate TCP port numbers or on separate virtual IP addresses. This approach is quite challenging operationally, since each new application instance requires a change in the network security policy in order to add the necessary TCP port number or virtual IP address.

Newer applications use application-specific contextual data such as a pre-pended value in an HTTP URL or in an SQL Bind, to direct application processing requests to the appropriate application instance. Consequently, the network and security operations teams cannot determine whether a session is for business or personal purposes simply by connecting to an application server IP address over port 80 (i.e., connecting to google.com over port 80).

Organizations need a new breed of security tools that can fluently identify application instances, application transactions, and application actions without relying solely on TCP/IP header information. The ability to identify an application based on internal characteristics such as protocol attributes is absolutely critical; without this granular visibility, it is impossible to provide the security protection that regulations dictate and business processes demand.

Identity-Based Access Enforcement

Organizations also need the ability to control application and resource access based on user identity, not just source IP address. With today's mobile, dynamic workforce connecting to application elements that reside on multiple servers within the data center, organizations can no longer assign access privileges based on a well controlled and determined user location represented by the IP address. Rather, IT must be able to provide access for any user, whether an employee, contractor, vendor, or other defined role connecting into the data center from anywhere at any time based on the user role.

User access cannot be allowed or denied based solely based on location because a given IP address may only be serving a user temporarily, for example, as a mobile user moves from one location to another. In addition, large pools of translated Network Address Translation (NAT) and proxy network addresses may serve a variety of users in various roles, and therefore are not tied to a specific user.

To control user access to dense application clusters, organizations need a network security solution that can enforce identity-based and role-based security policies. Such products should tie user identity to application access information, and make security decisions accordingly. In addition, to accommodate the growing use of collaboration tools across data centers, enterprises need a network security solution with identity federation (IF) capabilities that allow for seamless integration of services from external compute environments without compromising security.

The industry has defined a number of standard technologies to help disparate networks exchange identity and privilege information and share common notions of user identities, including the Security Assertion Markup Language (SAML), Extensible Access Control Markup Language (XACML), and Interface for Metadata Access Point (IF-MAP). These technologies make it easier for network security devices to enforce policies based on identity attributes. Enterprises should look for security solutions that support these standards and other technologies for identity store interaction.

High Capacity, Scalable Platform Design

Network-based devices designed to secure the data center must be able to handle the diversity and volume of traffic from today's applications. Data center interior firewalls must handle very high throughput rates, and be able to inspect and control high volumes of traffic crossing between different internal domains at LAN rates.

In addition, data center firewalls must support rapid session setup and teardown for very large numbers of sessions. In today's data centers, high numbers of users are accessing a targeted set of compute resources; as a result, a firewall must be capable of managing a million sessions even as tens of thousands of new and old sessions are being set up and torn down.

Given the demands on the data center—the reliability required by storage traffic and the number of TCP sessions and large packet sizes associated with SOA/mashup application architecture—firewalls must have a purpose-built, hardware-based design capable of tremendous processing power. In addition, data center security platforms must have a pay-as-you-grow design, one which allows service processing power to be added as needed without the need to physically or logically redesign the network as a consequence.

Centralized Management

Given the highly distributed, complex nature of today's data centers, it is a challenge to implement a consistent set of security policies across the entire data center infrastructure. Organizations need a comprehensive security solution that includes management consoles from which operations can centrally manage all functions, including defining a unified security policy and aggregating compliance information.

A robust, centralized management system gives the ability to define security policies that are detailed enough to apply granular controls to users, applications, and resource domains, but abstract enough that they don't need to be customized for each enforcement point. With centralized policy creation, managers are spared having to define security policies for each system within the data center, and the potential vulnerabilities created by a patchwork of policies are avoided.

Beyond supporting consolidated policy definitions, centralized management systems can significantly reduce security operational overhead in other areas. Organizations have a wide range of hardware and software deployed in data centers today, including switches, routers, server platforms, operating systems, application platforms, and application instances. Keeping all of these systems secure and up-to-date is a tremendous challenge. Organizations can significantly mitigate this challenge by deploying a management system that automatically receives updates and patches for a range of equipment, and allows IT to manually disseminate them. For example, a centralized management system can automatically locate and download the latest vulnerability protections, which operators can manually deploy to various enforcement points in a controlled fashion at their convenience.

Finally, centralized management tools with intelligent security information and event management (SIEM) capabilities can interface with network-based security appliances to provide detailed visibility into application usage and user access patterns, both in real time and historically. This consolidated information is necessary for compliance reporting, auditing, and other regulatory requirements. In addition, a centralized SIEM system gives IT the information needed to optimize data center resources and infrastructure. Organizations should look for SIEM tools with in-depth traffic analysis capabilities, including context and transaction analysis. These features enable IT to correlate detailed L7 deep inspection (DI) events with network flow information to form a real-time view of application usage trends.

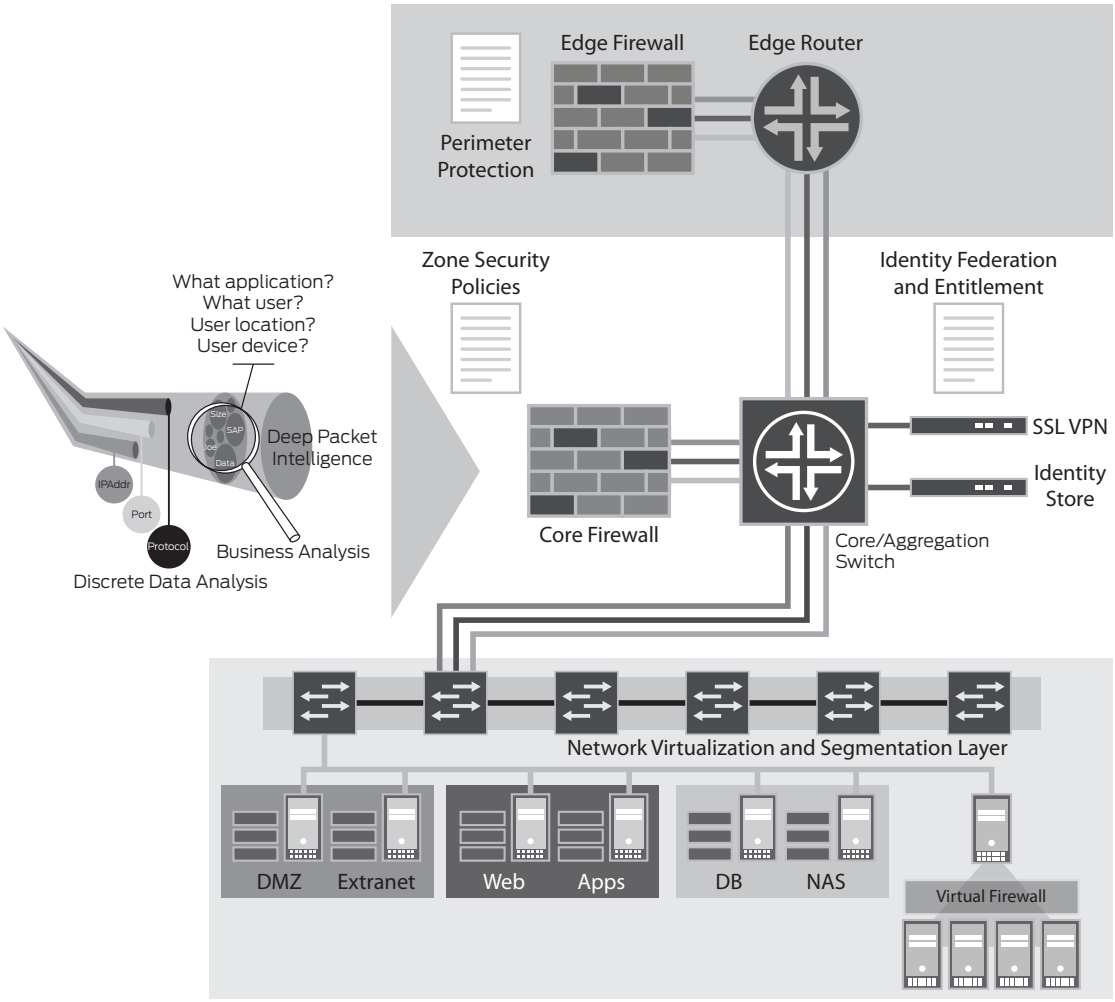

Figure 5: Data center security enforcement reference model

Conclusion—Securing the Data Center with a Network-Centric Approach

Today's organizations rely more than ever on their data centers to enable their operations. The data center continues to evolve as organizations concentrate resources and implement technologies such as server virtualization, distributed application models, and IP-based storage. As a by-product of this economic advantage, data centers are also more vulnerable now than before to an increasingly diverse mix of security risks.

In addition to cyber theft and increasing levels of malware, organizations must guard against new vulnerabilities introduced by data center technologies themselves. For example, unmanaged networks of virtual machines, the new communication patterns that result from mashups and SOA, and the use of IP infrastructure for storage networking all introduce new and unique security challenges.

To evolve and thrive with these evolutions data center teams should revisit existing information security best practices, technologies, and design approaches in light of these risks. A network-centric approach to securing the virtualized, cloud-ready data center offers a number of benefits, including scalability, unified policy definition and enforcement, and reduced operations overhead. In evaluating security solutions for the data center infrastructure, organizations should look for application fluency, identity-based policy enforcement, centralized management, and appliances with very large-scale processing.

Because the network touches every device in the data center, it is an ideal location from which to manage a holistic suite of security tools. And with advances in technology, it is possible to physically connect security devices to just a few points in the network and logically extend them to multiple network segments. Using a network-centric approach makes it easy to implement a unified security policy in a distributed environment, and to extend data center services to virtually any user in any location.

About Juniper Networks

Juniper Networks, Inc. is the leader in high-performance networking. Juniper offers a high-performance network infrastructure that creates a responsive and trusted environment for accelerating the deployment of services and applications over a single network. This fuels high-performance businesses. 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.

2000332-001-EN Oct 2009