paloalto NETWORKS

the network security company™

Embracing the Federal Cloud First Initiative

August 2012

Table of Contents

Executive Summary 3

Cloud Computing Brings Significant Benefits to Federal Agencies 3

Federal Cloud Deployment and Service Models Overview 5

Security and Privacy Implications of the Cloud 6

Governance and Compliance 6

Data Privacy 7

Secure Enablement of Applications 7

Threat Protection in the Cloud 7

New Security Attack Surface 8

Client-Side Protection 8

Security Architecture Flexibility 8

Management 8

Cloud Orchestration 9

Multi-Tenancy 9

Selecting the Right Security Solution for Your Cloud 9

Mapping Cloud Security and Privacy Requirements to Solutions 11

Governance and Compliance 11

Data Privacy 12

Safe Application Enablement 12

Hypervisor Protection 13

Flexible Network Architecture 13

Consistent, Secure Access to Cloud Resources 13

Multi-Tenancy 13

Centralized Management 14

Cloud-Ready Integration 14

Next-Generation Firewalls In Action 14

Summary 16

Executive Summary

In December 2010, the Federal Chief Information Officer (CIO) at the time Vivek Kundra instituted a new cloud computing initiative called the Cloud First policy. His vision was to enable government agencies to be more efficient, agile and innovative by leveraging cloud infrastructures to provision applications rapidly, paying only for IT resources consumed, and to increase or decrease usage depending on requirements and budget constraints. Kundra also believed that the adoption of cloud computing initiatives would allow federal agencies to be able to serve the American public better, and enable innovative services that would support that mission.

As part of this initiative, the United States Federal CIO laid out a series of guidelines and recommendations for moving towards cloud computing, and introduced government bodies such as the National Institute of Standards and Technology (NIST) and General Service Administration (GSA) to provide cloud computing standards and develop cloud-based application. One of the key guidelines driven by these bodies is on security. Security challenges such as data privacy, access control, multi-tenant environments and many CIA (confidentiality, integrity, availability) tenets are the biggest barriers to the adoption of cloud computing, regardless of the cloud computing architecture.

In fact, in a survey conducted by the Ponemon Institute to federal agencies[1] on the state of cloud (September 2011), 35% of respondents cited concerns about safety and security as reasons for not deploying cloud services, while 45% of respondents were not confident about data protection and security features of their current or prospective cloud service provider.

Therefore, a security solution that effectively solves security challenges in a cloud computing environment and addresses federal standards is needed for federal agencies to confidently embrace the Cloud First initiative. Cloud computing environments can be secure with the right security solution. This whitepaper details the security challenges in a cloud computing environment, provides clarity on the many federal standards, and how to address them in a pragmatic approach with a credible, effective and flexible network security solution.

Cloud Computing Brings Significant Benefits to Federal Agencies

Interest in cloud computing has accelerated in the last couple of years due to the ability to deliver flexibility and availability of resources at lower cost. The U.S. National Institute of Standards and Technology (NIST) defines cloud computing in Special Publication (SP) 800-145 as "a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (such as networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction."[2]

Cloud computing brings significant benefits to a current federal IT infrastructure that is outdated, expensive, complex, inefficient and sometimes proprietary because this shared pool of computing resources can be conveniently utilized and provisioned when necessary. Similar to a public utility service like electric power, the cloud economies of scale enable agencies and groups of agencies to centralize and standardize infrastructure which in turn lowers investment costs, enables on-demand service and improves scalability. As originally envisioned by the Office of the Chief Information Officer (CIO), a federal agency can leverage a cloud infrastructure to more efficiently serve millions of users quickly, and spend less time managing complex IT resources.

The Cloud First initiative introduced by the U.S. CIO is expected to bring the following benefits:

- **Efficient Usage of Infrastructure:** Cloud computing will enable federal IT organizations to move away from complex, heterogenous environments towards pooled IT infrastructure resources that can be shared across large numbers of organizations. The lower infrastructure expenditure brings significant cost savings that can be directed towards other objectives. At the same time, cloud computing also delivers better utilization of existing assets, simplifies IT and provides consistent resource management, eliminating the need for every agency to overprovision and maintain reserve capacity to meet periodic or unexpected demand.

- **Responsive and Scalable Federal Services:** Cloud computing will also enable agencies to more quickly respond to changing needs. Because the pool of computing resources can be rapidly provisioned, federal programs and applications that exhibit unpredictable service demands within a short deployment window can be efficiently handled using a cloud computing approach. With a larger pool of resources to rely on, federal agencies can also mitigate service outages by drawing from underutilized infrastructure or rapidly increasing capacity to meet demands.

- **Technology Innovation:** Cloud computing also enables innovation by allowing agencies to deliver innovative services that are not technically or economically feasible. For example, the ability to enable federal information to be accessed from any type of mobile devices, or to enable federal employees to work from remote locations can be facilitated with cloud computing. The ability to conceive, develop and test projects quickly in a cloud platform development environment allows projects to be quickly evaluated for promise, reduces the need for huge upfront investment and provides an experimental platform for innovative ideas. By consolidating IT infrastructure and leveraging common mission processes across multiple agencies, many innovative services can become a reality.

Cloud computing services are currently being used or considered by governments worldwide:

- **Canadian Government:** In April 2010, the Canadian Government's CTO of Public Works Government Services, Jirka Danek presented the Government of Canada's cloud computing roadmap that aims to improve the efficiency of the government network. According to Danek, there are 325,000 employees in the Canadian federal government, 140 departments (all with their own CIO), 124 networks and 144 data centers across the country. In this network, 120,000 Wintel and Unix servers use less than 10 per cent of their capacity. As a result, the Government of Canada's plans to leverage cloud computing for pay, pension, Campus-Direct, GC Intranet and Canada.gc.ca applications. Short-term goals are to use SaaS for internal collaboration (such as GCPedia, GCConnex and GCForum), PaaS for commoditized Web hosting and IaaS for virtual storage and computing services.

- **United Kingdom G Cloud:** The UK Government has announced the establishment of a UK onshore, private Government Cloud Computing Infrastructure called G-Cloud in 2011. The G-Cloud plan calls for 50% of new government IT spending to move to cloud computing services by 2015. The initial focus is on introducing cloud services into government departments, local authorities and the wider public sector. The program will include Infrastructure-as-a-Service (IaaS), Middleware/Platform-as-a-Service (PaaS) and Software-as-a-Service (SaaS). In relation to SaaS the government would establish a Government Application Store that offers "off the shelf" IT services such as email, word processing, enterprise resource planning and electronic records management that meet government standards.

- **Japan Kasumigaseki Cloud:** In 2009, Japan's Ministry of Internal Affairs and Communications (MIC) released a report outlining the Digital Japan Creation Project (ICT Hatoyama Plan) which includes a nation-wide Cloud Computing infrastructure tentatively called the Kasumigaseki Cloud. The Kasumigaseki Cloud will be deployed in stages through 2015 and aims to cut development and operating costs while also improving performance. The cloud will require Japan to create new platforms for shared services and consolidate hardware, including the possible building of new data centers.

- **European Union:** The Seventh Framework Programme (FP7) bundles all research-related EU initiatives together under a common roof. The objective of this program is to play a crucial role in reaching the goals of growth, competitiveness and employment in the European Union. The FP7 is funding several projects on cloud computing and has also compiled a group of experts to outline the future direction of Cloud Computing research.

- **Research Organization Testbed:** The Open Cirrus™ project provides systems researchers with a testbed of distributed datacenters they can use for systems-level cloud computing research. The project is a joint initiative sponsored by HP, Intel, and Yahoo!, in collaboration with the National Science Foundation (NSF), the University of Illinois (UIUC), Karlsruhe Institute of Technology, Russian Academy of Sciences, MIMOS a Malaysian research and development organization and the Infocomm Development Authority of Singapore. 80 research projects are currently running on this test bed including a stem cell research project by Carnegie Mellon.

Federal Cloud Deployment and Service Models Overview

NIST defines cloud computing with five key characteristics—on demand self-service, broad network access, resource pooling, rapid elasticity and measured service. The first three characteristics define the ability to easily provision these "pool" of services over the network without requiring human interaction, while the last two characteristics define the ability for the service to scale and optimized for a particular level of service.

Federal cloud applications can be deployed in a variety of different deployment modes—public, private, community, and hybrid. A public cloud is owned and operated by an external cloud provider, and infrastructure and computational resources are made available over the Internet. A private cloud is operated exclusively for an organization, although management may be provided within the organization or by a third party. A community cloud is similar to a private cloud except it extends to a group of like-minded organizations with common industry, privacy and security considerations. Hybrid clouds can be composed of two or more clouds—private, public or community.

The service model describes how applications are delivered to agencies and the scope and control over the computational and application development environment.

- **Software as a Service (SaaS):** Agencies are provided access to an application running on a cloud infrastructure. The application is accessible to the end user from various client devices and interfaces, but the agency does not manage or control the underlying cloud infrastructure. Examples of SaaS include Salesforce.com and Google Apps.

- **Platform as a Service (PaaS):** Agencies can deploy and control supported applications on the provider's cloud infrastructure using programming languages and tools supported by the provider. This means developers can develop applications without the cost and complexity of buying and managing the hardware and software. Examples of PaaS include Salesforce.com's Force.com, Google's App Engine and Microsoft's Azure.

- **Infrastructure as a Service (IaaS):** Agencies can provision processing, storage, networks, and other computing resources, and deploy and run operating systems and applications. The customer does not manage or control the underlying cloud infrastructure, but controls the specific platform virtualization environment being used—servers, operating systems, storage, and deployed applications. Examples of IaaS include Rackspace and Amazon's Elastic Compute Cloud (EC2).

Key Cloud Characteristics

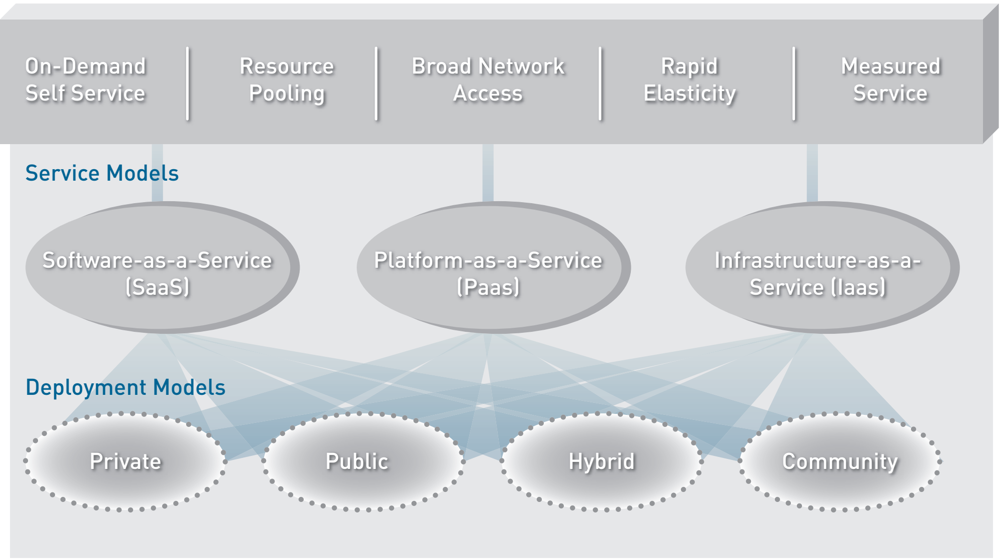

Security and Privacy Implications of the Cloud

There are a variety of different security risks that need to be considered with cloud deployments. Whether the deployments being evaluated include public, private, community or hybrid clouds, IT now essentially operates a service delivery model, in a dynamic environment, with other "tenants" on the same infrastructure.

As agencies evaluate cloud computing environments, many of the same security challenges in a data center environment continue to be applicable in the cloud. In addition, the unique architecture used to deliver cloud services such as virtualization and orchestration bring additional complexities to network security. Security and privacy implications of cloud computing include:

Governance and Compliance

Because cloud computing services are so easily deployed, one of the challenges is rogue, sprawling cloud applications being deployed without being governed by appropriate federal operational policies. Privacy, security and oversight could be overlooked, and vulnerable systems could be deployed as a result. Governance or the control and oversight over policies, procedures and standards for application development, the design, implementation, test and monitoring of services, and security audits must be extended to the cloud environment.

There are specific laws and regulations that need to be met for cloud environments. Most revolve around the data that is stored in the cloud, for example, regulatory requirements on disclosure of security breaches of personal information, privacy of personal information, data location and transfer, and electronic discovery of information.

Contracted or external service providers are subject to the same regulations as federal agencies. For example, FISMA and OMB policies require that external providers handling federal information or infrastructure and information systems on behalf of the federal government, meet the same security requirements as federal agencies. Under the Federal Records Act and National Archives and Records Administration (NARA), contractors must manage federal records in accordance with the corresponding laws and regulations.

The following laws and regulations may need to be considered in a cloud environment by agencies and cloud providers:

- **Clinger-Cohen Act:** The Clinger-Cohen Act assigns responsibilities for the efficiency, security and privacy of computer systems within the federal government.
- **FedRAMP:** FedRamp requirements are based on NIST 800-53 rev 3, and govern the security requirements for a cloud environment for agencies and cloud service providers. Cloud service providers must implement the FedRAMP security requirements and hire FedRAMP approved third-party assessment organization to provide a security assessment of the cloud environment.
- **Federal Information Security Management Act (FISMA):** FISMA requires federal agencies to adequately protect their information and information systems against unauthorized access, use, disclosure, disruption, modification, or destruction.
- **National Archives and Records Administration (NARA):** Under the Federal Records Act and NARA regulations, agencies are responsible for managing federal records effectively throughout their lifecycle, including records in electronic information systems and in contracted environments.
- **Office of Management and Budget (OMB):**
  - **Circular A-130 Privacy Act:** Circular A-130 establishes policy for the management of Federal information resources. Appendix III or A-130 requires that adequate security is provided for all agency information that is collected, processed, transmitted stored or disseminated in general support systems and major applications.
  - **M-03-22, OMB Guidance for Implementing the Privacy Provisions of the E-Government Act of 2002:** This regulation provides direction to agencies on conducting PIAs. A PIA is a structured review of an information system to identify and mitigate privacy risks, including risks to confidentiality, at every stage of the system lifecycle. It can also serve as a tool for individuals working on a program or accessing a system to understand how to best integrate privacy protections.
- **Privacy Act:** The Privacy Act governs the collection, maintenance, use and dissemination of information about individuals that is maintained in systems of records by federal agencies and can be retrieved by a personal identifier. It requires each agency to publish notice of its system of records in the Federal Register and to allow individuals to request access to and correction of their records and information.
- **Industry Requirements:** Requirements such as the Health Insurance Portability and Accountability Act (HIPAA) and the Payment Card Industry Data Security Standard (PCI DSS), may apply to a particular organization. For example, if a particular healthcare service is being offered to a Federal agency, HIPAA requires both technical and physical safeguards for controlling access to protected health information, which may create compliance issues for some cloud providers.

Data Privacy

There are aspects of data storage, location and leakage that need to be considered in cloud computing environments:

- **Data Location:** In a cloud environment, data is stored redundantly in multiple physical locations, and therefore it may be hard to determine whether safeguards are in place and legal/regulatory requirements are being met. When data crosses boundaries, the jurisdiction and legal or regulatory requirements of the data flow, handing and processing of the data, and the privacy of the data may need to be considered. In addition, the electronic investigation of inappropriate or illegal activities for forensic or legal reasons, and trying to conduct them in a dynamic cloud environment is a challenge.

- **Data Retention and Ownership:** The agency's ownership of the data needs to be clearly established. This includes data that is collected by the cloud provider about customer-related activities in the cloud, as well as if the cloud provider service is terminated and the infrastructure shut down. The ability to retain ownership of the data and transfer it from the current cloud provider to a new one would be needed. If business records need to be archived, then the cloud provider or cloud-computing environment must support an archival infrastructure that is robust in the storage and retrieval of data, and meet requirements by NARA.

- **Data Leakage:** Data leakage in cloud environments is a problem because of the sensitivity of the data that could potentially fall in the wrong hands. Data leakage monitoring is a key requirement, in addition to the establishment of a response and notification policy in the event of data leakage. In addition, high-risk data that falls under a specific regulatory requirement should be encrypted while in transit and at rest.

Secure Enablement of Applications

Securely enabling access to information resources in the cloud is a key security consideration, in particular since access can comprise not only current employees or external business partners that have received access to an agency's cloud network, but also cloud provider insiders.

The ability to securely enable applications in the cloud begins with the identification of applications, but is complicated by several trends. The ease of application creation and delivery in the cloud makes it challenging to keep up with application proliferation. Application developers have been known to implement applications on any port that is convenient or bypass security controls altogether. Many applications are also designed from the outset to circumvent traditional security solutions by evasive tactics such as port-hopping, hiding within SSL encryption, tunneling within commonly used services. Many of the evasive tactics are used to hide threats within applications.

Traditional security solutions that depend on ports as a basis to understand applications, or cannot understand the techniques that applications use to hide and circumvent controls are ineffective in this new application landscape. To effectively and secure enable applications in the cloud, agencies must be able to identify, control and enable the applications at a granular level, by application characteristics, users, and content.

Threat Protection in the Cloud

The cloud computing environment is no different from any on-premise data center environment - it must be protected from threats. Security solutions in the cloud must consider the evolving threat landscape and how to effectively secure against a new set of modern attacks. These modern attacks are much more innovative in their attack strategies, taking advantage of new communication methods like social media to propagate infected links, and adopting a variety of additional techniques like proxies, circumventors, the use of non-standard ports, and tunneling to avoid security solutions. As malware has become more advanced, it has also become more targeted and customized for a particular network, thus helping it to avoid traditional signature-based anti-malware and network security products.

These techniques allow threats to run covertly on networks and systems, quietly collecting sensitive or personal data and going undetected for as long as possible. This approach enables repeated use of the same exploits and attack vectors, and the ability to cross the security "perimeter" repeatedly. To effectively secure a cloud environment requires a multi-pronged approach. It requires an understanding of the allowed applications in the cloud environment to reduce the scope of attack, followed by an integrated threat protection system that can address multiple threat vectors to control known and unknown threats.

New Security Attack Surface

In a cloud environment, in order to more efficiently deliver services, the infrastructure is typically built on a virtualized platform. Virtualization technology partitions a single physical server into multiple operating systems and applications that are deployed by the multiple tenants that are being serviced. Through the dynamic nature of virtual machine provisioning, applications can also be delivered quicker or as mission demands dictate to further improve the agility of the cloud provider. However, this new "cloud infrastructure", is made up of many different components—from hypervisor to guest operating system and application—each of these components is vulnerable in a cloud environment.

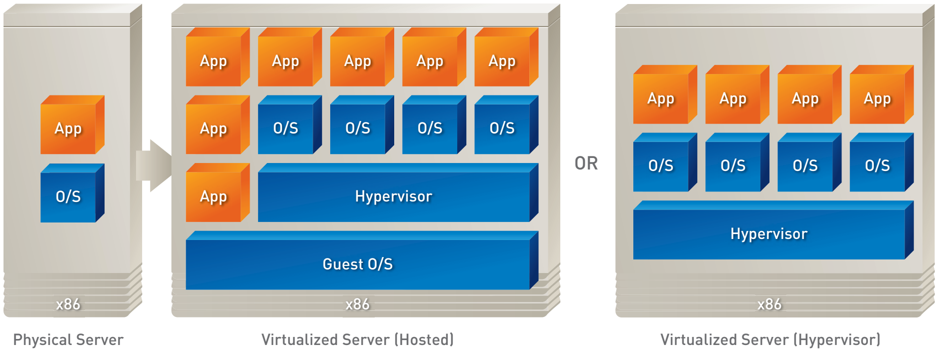

The hypervisor, a software layer that sits between the hardware and the "virtual" operating system and applications, is what allocates memory and processing resources to the "virtual" machines. Hypervisor attacks range from vulnerabilities that cause hypervisors to crash to complex "breakout" exploits that cause a guest VM system to escape and infiltrate its own host system. Because data from hundreds or thousands of companies can be stored on large cloud servers, attackers can theoretically gain control of huge stores of information through a single hypervisor attack.

In a cloud computing environment, the implications of combining virtualization workloads with different trust levels on the same server should be considered. Most cloud computing deployments will allow virtualized workloads and live migration of VMs to those with the same trust levels in the initial phase, but over time accommodate a mix of trust levels. As trust levels are integrated, additional security requirements of intra-host VM inspection will be required.

Client-Side Protection

Most of the focus with cloud protection is on the server side. The reality however is that users on a variety of different devices are going to be accessing these virtual applications. Therefore, client-side protection, ensuring that all types of devices have secure access to the cloud applications is critical.

Security Architecture Flexibility

The architecture used to deliver cloud services can vary widely, which can lead to challenges when integrating security. Enabling implementation depends not only on supporting a wide range of networking features and options, such as 802.1Q and port-based VLANs, but also the ability to integrate at layer 1, layer 2 or layer 3. Network security that is simple enough to integrate into any type of cloud environment, and flexible enough to adapt when the threat and application landscape changes is needed.

Management

The cloud computing should include the ability for federal agencies to manage security solutions in a centralized manner. Agencies should also have the ability to control and modify aspects of the security solutions to accommodate changing needs. For example, the ability to observe and change security metric details observed via dashboards and reports, the ability to modify the threshold for alerts and notifications, and the ability to change the schedule of reports/logs.

Cloud Orchestration

Cloud environments typically automate the tasks and processes using workflows that help IT teams execute change with greater speed, quality, and consistency. Deployment of security capabilities typically lags orchestration software provisioning for cloud environments, leading to security risks and considerable integration challenges.

For example, a virtual machine may be provisioned in seconds, while the provisioning of a network feature to support the VM may take minutes. A traditional security process where change control tickets to modify security policies take weeks for approvals will significantly slow down the very benefits of on-demand cloud computing services. An automated way to provision network security in line with the pace of orchestration of the elements of the cloud computing environment is needed.

Multi-Tenancy

Cloud computing environments support large numbers of customers and platforms to achieve cost benefits and economies of scale from cloud services. Multi-tenancy from multiple agencies sharing the use of a common infrastructure or platform can lead to higher security risks. For example, security breaches with one agency's network can expose other agency's network to the same attack, if the cloud computing environment is not properly segmented.

Selecting the Right Security Solution for Your Cloud

As described in the "Security and Privacy Implications of Cloud" section, there are a number of issues that need to be considered to secure the cloud. In addition to the fundamental requirement to securely enable applications and defend against threats, the security solution for the cloud must be flexible enough to support diverse cloud computing architectures, have the ability to be managed centrally, and support a variety of cloud integration features.

The security requirements for the cloud can thus be summarized as follows:

- Identify, control and safely enable all applications within the cloud computing environment.
- Protect against all data center threats all the time, without compromising the performance of the cloud computing environment.
- Flexibility to adapt to new architectures and changes in security posture.
- Centrally managed with a single unified policy.
- Cloud-readiness for integration into orchestration and automation framework.

The NIST 800-53 Rev.3[3] states that "the security challenges cloud computing represents are formidable." As acknowledged by FedRAMP, "The decision to embrace cloud computing technology is a risk-based decision, not a technology-based decision".

It is clear that multiple federal agencies recognize the importance of a strong cybersecurity posture to handle the challenges of cloud computing. These new standards and the focus on security are important in a current landscape where attacks are not only on the rise but distinguish themselves with their sophistication and covert nature.

Traditional security solutions are no longer effective in dealing with this new threat landscape. Traditional security solutions all use a port-based traffic classification engine that allow too much traffic through, in particular applications that use evasive tactics like non-standard ports, SSL or port hopping. The lack of visibility into traffic also translates into an inability to detect threats as threats use that very same behavior to circumvent security. More importantly, the ability to granularly enable safe application usage is not possible with these legacy solutions.

Next-generation firewalls have emerged as the solution of choice for security conscious enterprises tackling this new application and threat landscape, and the same level of protection is what is needed in a cloud computing environment. It is only a matter of time before the nation states and criminal organizations responsible for sophisticated enterprise attacks transition their focus to cloud computing environments. As cited earlier, in a survey commissioned by SafeGov.org, and conducted by the Ponemon Institute to federal agencies on the state of cloud in September 2011:

- 35% of respondents were not deploying cloud services because of concerns about safety and security.
- 45% of respondents were not confident about data protection and security features of their current or prospective cloud service provider.
- 61% of IT managers in this survey believe that one or more agencies other than their own will suffer a cloud-related breach in the next 12 months.

Next-generation firewalls are credible security solutions that have tackled modern threats in many enterprise environments. The next-generation firewall is needed to secure cloud computing environments and instill confidence in federal agencies that these environments can be safe. Next-generation firewalls provide the all-important tools needed to safely enable applications and deployments, by providing strong and granular segmentation options based on application, user and content. Specific users or user groups can be allowed specific access to data, applications and even particular features to ensure that applications and data are accessed appropriately. The content can be monitored for threats as well as data leakage in both directions to ensure that cloud computing does not create security vulnerabilities.

Federal agencies can address multiple federal standards that recommend threat technologies like Intrusion Prevention Systems (IPS) and malware prevention with a single integrated appliance. More importantly, with this integrated threat framework, agencies will be able to understand the interconnection of applications, exploits, malware, URLs, anomalous network behaviors and targeted malware in a uniform context. This context gets to the important conclusion faster, streamlines management and reporting and ensures predictable performance by analyzing traffic once instead of progressive scanning in multiple engines.

Before federal agencies can take advantage of the efficiency and cost savings of cloud-based applications, cloud deployments must be properly secured with next-generation firewalls. Agencies are ultimately responsible for the protection of cloud applications, and next-generation firewalls can help achieve security beyond what is recommended by standards to allow agencies to future-proof and fully protect its cloud computing environment. Palo Alto Networks™ next-generation firewalls, from the architecture to the security technologies and the strong network foundation have been designed specifically to address the security challenges in a cloud computing environment.

Mapping Cloud Security and Privacy Requirements to Solutions

Cloud computing requires a fundamental shift in how IT teams operate. Instead of just focusing on managing assets, IT teams will need to shift the focus towards services—how to deliver the most effective services, what types of meaningful services align to mission objectives, and how to measure the effectiveness of these services. This applies to security as well. Federal agencies will need to evaluate periodically security metrics such as performance, and high-availability as well as monitor how well the cloud security infrastructure is stacking up to attacks.

The selection of the cloud service model determines the level of responsibilities on the agency versus the cloud provider. In an IaaS deployment, the cloud provider owns the cloud computing infrastructure, but the agency owns the entire application stack. Therefore, while hypervisor and server security may be the domain of the cloud provider, application level security needs to be addressed by the agency. In this case, a next-generation firewall can be deployed in the IaaS environment. In a PaaS model, because the cloud service providers are providing the platform and proprietary software suite to write applications, the cloud provider must provide the necessary security solution to secure the platform and application stack. In a Saas environment, the cloud provider owns the application and infrastructure stack. The agency will need to manage security policies for access to the SaaS applications.

This section maps the security and privacy issues outlined earlier to practical solutions, leveraging Palo Alto Networks next-generation firewalls as a strategic tool in securing cloud computing environments. Depending on the choice of service model for cloud, some of these guidelines may be applicable:

Governance and Compliance

Agencies must identify laws and regulations that are applicable in a cloud environment. If deployments are to be hosted in a cloud provider environment, review and access cloud provider's ability to meet them. Leveraging next-generation firewall technologies to deliver meaningful segmentation based on application, user and content, and produce usage reports on applications will provide compliance validation to auditors.

Data Privacy

It is critical to understand what types of data is suitable to be hosted in a cloud computing environment, controlling access to the data, and securing the data while in transit and at rest. If agencies are planning to outsource their cloud computing environment, data privacy issues need to be considered with cloud provider partners to ensure that service provide are protecting the use, processing and communication of personal and personally identifiable information in the cloud.

The data filtering option on next-generation firewalls can track when confidential information flows out of a particular segment of the cloud computing environment. The data filtering option provides a means to track unauthorized file and data transfer in a cloud environment. Palo Alto Networks data filtering option is able to control flows by looking deep within the payload to identify the file type (as opposed to looking only at the file extension) to determine if the transfer of the file is allowed by policy. File blocking by type can be implemented on a per application basis. Feature level control over file transfer represents another policy option that may help balance application use with policy control. Policies can be established to allow the use of web application but deny the use of the related file transfer function. Rounding out the filtering features is the ability to identify and control the transfer of sensitive data patterns such as credit card numbers, social security numbers or custom data patterns in application content or attachments.

Safe Application Enablement

Palo Alto Networks next-generation firewalls allow agencies to have complete visibility into the cloud environment and safely enable applications. Safe enablement of applications means the ability to transform your traditional allow or deny firewall policy into meaningful mission –relevant elements such as the application identity, who is using the application and the type of content or threat to control access.

This can translate to policies such as:

- Allow only the IT group to use a fixed set of remote management applications (e.g., SSH, RDP, Telnet) across their standard ports but block their use for all other users. Limiting user privileges, and controlling rogue admins is a good best practice for security.

- Allow all groups to use Microsoft Sharepoint features such as blog, and calendar, but only users with the right security clearance can access documents.

- Explicitly block all P2P, circumventors, non-VPN related encrypted tunnels, and external proxies, regardless of port, protocol or evasive tactic.

It can also deliver more meaningful segmentation by application, user and content - Databases, applications and entire zones of the cloud computing environment can easily be segmented and controlled to ensure agencies share the information that need to be shared, yet protect the information it needs to protect. For example, an agency can segment off the virtual workload containing cardholder data, only permit access to that segment to finance users employing the payments application and monitor any cardholder data that leaves the segment. This contains and limits access, and delivers individual accountability. Having that level of control, and perhaps most importantly, auditability, is indispensable for federal agencies.

Safe application enablement is accomplished in three ways:

- **Secure enablement begins with identifying the application:** App-ID™ provides accurate traffic classification regardless of ports, protocol, evasive tactics and SSL, and can proactively reduce the attack surface by blocking rogue, misconfigured applications, unauthorized management applications and peer-to-peer file sharing. App-ID continually monitors application state, checking to see if specific features, such as file transfer, or posting are enabled and when that change in state is detected, an appropriate decision can be made based on the security policy. Additional data points that can be used to help you make a more informed decision on how to treat the application include an application description, how the application behaves, which ports it may use, and how it is categorized.

- **Tying users, not IP addresses to application:** Government agencies often need to support a wide variety of users with very different trust levels including contractors, civilian employees and even other agencies. User-ID™ gives you the control to provide the level of access that is appropriate to a user's role instead of IP addresses.

- **Application-specific threat prevention:** The threat footprint is first reduced through an explicit deny policy for unwanted applications such as external proxies, circumventors, and P2P file-sharing. Then, as you enable specific applications and associated features, you enable virus, vulnerability exploit, spyware and modern malware protection features to extend the application-specific context into threat prevention. For example, you can allow Oracle on its standard port only for finance and operations and protect against SQL injection attacks and Oracle specific vulnerability exploits.

The integrated threat protection capabilities include not only intrusion prevention features to block network and application-layer vulnerability exploits, but also denial-of-service features to control various types of traffic floods. In addition, WildFire™ provides the ability to identify malicious behaviors in executable files by running them in a virtual environment and observing their behaviors. This enables Palo Alto Networks next-generation firewalls to identify malware quickly and accurately, even if the particular sample of malware has never been seen in the wild before.

Because safe application enablement and threat protection must be delivered without impacting the performance requirements of the cloud, Palo Alto Networks hardware and software architecture was designed from the ground up specifically for performance. The multi-core architecture incorporates dedicated, specialized processing for networking, security, and content scanning so that the full suite of next-generation features can be enabled with high throughput and reliability. This optimized hardware architecture complements a single pass software architecture that processes functions in a single pass to reduce latency.

Hypervisor Protection

The foundation of hypervisor security is pretty basic. It starts off with hardened software. The organization's security policy for software operating systems on physical servers should be extended to virtualized servers. This includes how often operating systems are updated, and bugs patched.

Isolation between components within the server and controlled access to the server is critical. For example, the guest operating system should never be allowed to access the hypervisor. For the same reason that direct access to management of physical servers is restricted, access to the hypervisor management system in a virtual environment must also be protected and restricted to ensure that management functions are only available to approved IT administrators. This means locking down the plethora of management options for virtualized platforms, from browser and client-based console access to scripts, command line interfaces to centralized management tools.

Because certain hypervisor exploits can only be addressed with solutions outside the server, for example external firewalls or IPS systems, external next-generation firewalls should be deployed to front-end vulnerable virtual systems. These firewalls provide two sets of benefits – they not only address exploits, but also enforce appropriate access to critical virtualized systems.

Flexible Network Architecture

Palo Alto Networks firewalls deliver integration at Layer 1, Layer 2, Layer 3, and tap modes (or a mixture of all on the same appliance) and couple that with powerful networking capabilities for integration (VLAN trunking, link aggregation) and high availability (separation of data and control planes, active/active and active/passive deployment options). This accommodates any cloud architecture, and provides the foundation to easily integrate into the cloud computing environment in order to leverage the next-generation firewall technologies.

In addition, because Palo Alto Network's deliver next-generation firewall features on the same appliance, enabled via software licenses, federal agencies can easily adapt their cybersecurity posture to changing threat and application landscape requirements without requiring a fundamental change in their cloud computing architecture.

Consistent, Secure Access to Cloud Resources

Whether cloud resources are being accessed by government users, civilians, or contractors, security policies must be enforced for these endpoints, and access to the cloud must be secured. GlobalProtect addresses this by automatically establishing either an SSL or IPSec-based VPN connection to the Palo Alto Networks firewall. The remote access connection is authenticated through one of several mechanisms (local DB, RADIUS, LDAP, Active Directory and Smartcards), and enforces the same secure application enablement policies regardless of the user's location of access.

Multi-Tenancy

Virtual systems are unique and distinct next-generation firewall instances within a single Palo Alto Networks firewall. Rather than deploy many individual firewalls, you can deploy a single pair of firewalls (high availability) and enable a series of virtual firewall instances or virtual systems in the cloud. Each virtual system is a self-contained, fully operational Palo Alto Networks firewall, complete with separate management interfaces which ensures that other customer or departmental virtual systems will only see or modify their own policies. Within each virtual system, role-based administrative access control allows organizations to delegate feature level administrative access (enabled, read-only, or disabled and hidden from view) to different staff members.

Centralized Management

Panorama is a centralized security management system that provides global control over a network of Palo Alto Networks next-generation firewalls. Panorama allows administrators to control all aspects of the devices and/or virtual systems under management (security, NAT, QoS, policy based forwarding, decryption, application override, captive portal, and DoS protection). Using pre- and post-rules, Panorama administrators can enforce shared policies while allowing local policy flexibility. Rules in between the pre- and post-rules can be edited locally or by a Panorama administrator who has switched to the local firewall context. Software updates such as dynamic content updates (Applications, Threats and antivirus), and software licenses can also be managed centrally on Panorama.

Panorama provides the ability to view logs and run reports across dynamic or locally queried data aggregated from managed devices. Distributed reporting can be done without a need to forward logs from firewalls to Panorama. Aggregate user activity reports can be run for mobile users that access cloud resources from different locations. This will report users' activity regardless of where they are currently located globally.

Cloud-Ready Integration

The management and orchestration of a cloud computing environment is complicated, and can involve multiple phases. The first phase includes the provisioning of the virtual server, followed by the provisioning of the networking elements in the virtual environment such as virtual switches. Finally, the last phase is the security provisioning. In order to properly scale, orchestration software is needed to automate many of these processes.

Palo Alto Networks offers a powerful XML management API that enables external orchestration software to connect over an encrypted SSL link to manage and configure Palo Alto Networks firewalls. The exhaustive and fully-documented REST-based API allows configuration parameters to be seen, set and modified as needed. This management API enables the proper integration with data center orchestration software so that the security features within the next-generation firewall become part of the data center workflow. Turnkey service templating where existing security service definitions are defined can be enabled for cloud providers.

Next-Generation Firewalls In Action

In this section, we explore a validated design of a deployment of Palo Alto Networks next-generation firewalls in a cloud computing environment, and how the features described earlier play a role in securing the environment. The diagram below shows a cloud computing environment, where multiple virtual machines are being instantiated for a community cloud. In this example, Microsoft Sharepoint, MS SQL and Active Directory applications have been instantiated on virtual machines. The cloud infrastructure is built on a highly secure and scalable networking architecture.

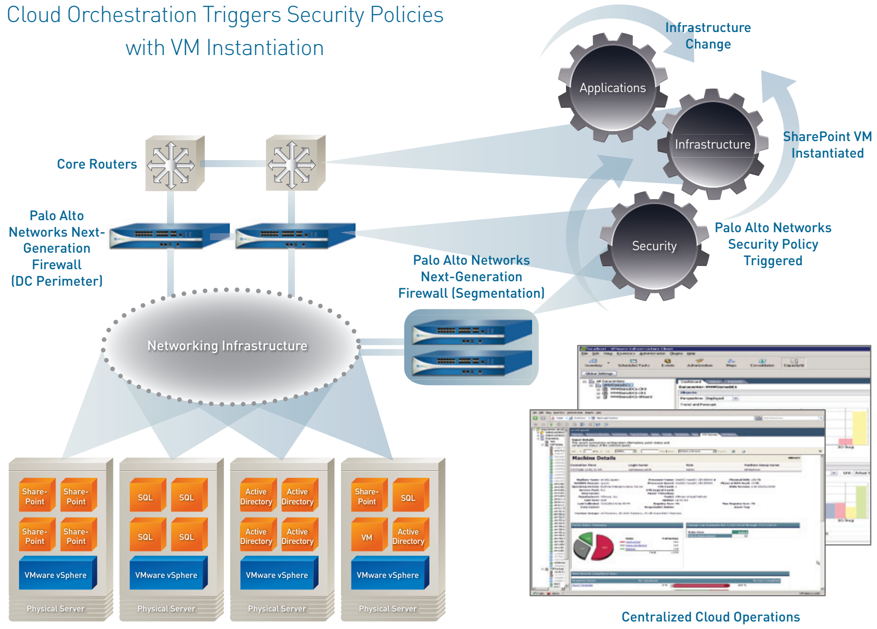

Cloud Orchestration Triggers Security Policies with VM Instantiation

Two sets of Palo Alto Network firewalls are deployed in this environment:

- **Cloud Border Firewall:** Next-generation firewalls are deployed at the border of the cloud to:
  - **Protect against threats and unauthorized access into the cloud computing environment:** Depending on the cloud environment, a variety of threat protection features can be enabled including IPS, anti-virus and denial-of-service.
  - **Terminate VPN tunnels:** Depending on the topology of the cloud environment, this same set of next-generation firewalls may also serve as VPN gateway (not shown in diagram). Access to agency-specific cloud resources is provided by an encrypted GlobalProtect VPN tunnel using two-factor authentication, and with host-level protection for diverse endpoints.
  - **Access control is strictly enforced:** Access to data center resources must be enforced based on the employee role and responsibilities.
- **Segmentation Firewall:** The internal next-generation firewall serves as a segmentation firewall in layer 3 to securely enable applications running in the cloud environment.
  - **Security Zone-based policies:** Virtualized servers are segmented appropriately into security zones based on similar risk factors and security classification. The segmentation firewall restricts access to specific zones. Management segmentation via virtual systems isolates one agency's network from another
  - **Safe application enablement:** Secure application enablement identifies and safely enables Sharepoint and MS SQL applications. Usage of application functions are strictly restricted based on user, and content is inspected to ensure that it is compliant to cloud security policies
- **Both sets of firewalls can be managed by Panorama:** Specific device-level configuration can be enabled based on agency IT roles. All traffic is inspected and logged centrally on Panorama.
- **Orchestration software enables multiple elements in the cloud including security to be provisioned when the application is instantiated:** For example, when the virtual machine for Sharepoint is instantiated, infrastructure elements ensure that proper placement in the right VLANs and appropriate networking features are enabled. At the same time, Palo Alto Networks security policies for Sharepoint are triggered to ensure that agency policies are strictly enforced.

Summary

The Federal Cloud First initiative will enable a fundamental shift in how federal agencies optimize and standardize their infrastructure, reduce IT costs and serve the American public. Many transformational and innovative services for the American people will be created because of the benefits of cloud computing technologies.

However, before federal agencies can leverage the benefits of cloud computing, security challenges such as safe application enablement, threat protection and data privacy must be considered. Palo Alto Networks next-generation firewall capabilities, with the ability to safely enable applications, protect against threats and integrate flexibly into any cloud infrastructure enable federal agencies to address security standards and secure any cloud computing environment.

Citations:

[1] NIST SP 800-145 - The NIST Definition of Cloud Computing

[2] "State of Cloud Services in the U.S. Federal Government"
http://safegov.org/media/23573/ponemon_survey_raw_results.pdf

[3] NIST 800-53 Rev.3 – Recommended Security Controls for Federal Information Systems and Organizations

Additional reading:

- NIST Special Publication 800-144 – Guidelines on Security and Privacy in Public Cloud Computing
- NIST Special Publication 800-146 – Cloud Computing Synopsis and Recommendations
- Federal Cloud Computing Strategy – Vivek Kundra, Office of the CIO, February 2011
- Federal Authorization Management Program (FedRAMP) Centralized Cloud Cmputing Assessment and authorization - http://www.fedramp.gov

the network security company™

3300 Olcott Street
Santa Clara, CA 95054

Main: +1.408.573.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

www.paloaltonetworks.com