



Reference Architecture

Managing Risk in the Cloud

Mike Scanlin, NetApp, with support from the NetApp Field Centers for Innovation
October 2010 | RA-0006



EXECUTIVE SUMMARY

More often than not, IT as a service (ITaaS) is used in an abstract sense to describe one or more cloud computing service models (IaaS, PaaS, SaaS) with an emphasis on the “IT” rather than on the service that “IT” is expected to provide. Designed for multi-tenant environments that demand secure separation, NetApp and partner secure multi-tenancy (SMT) technologies deliver concrete business value by enabling customers to reap the cost savings and efficiency benefits of a shared infrastructure without relinquishing control or compromising security.

TABLE OF CONTENTS

1	INTRODUCTION	3
2	FRAMING THE CLOUD DISCUSSION	4
	WHAT IS CLOUD COMPUTING?	4
	WHAT IS MULTI-TENANCY AND WHY IS IT IMPORTANT?	6
3	THREATS, RISKS, AND BENEFITS OF CLOUD COMPUTING	7
	WHAT ARE THE THREATS IN THE CLOUD?	7
	CAN YOU AFFORD A BREAK IN THE CLOUD?	7
4	HELPING TO ASSURE THE CLOUD WITH SMT	9
	THE FOUR PILLARS OF SMT	9
	SMT: HELPING TO ASSURE INFORMATION AT THE NETWORK, COMPUTE, AND STORAGE LAYERS	10
5	APPLYING THE RISK MANAGEMENT FRAMEWORK	14
	HOW CONTROLS PROVIDE A MORE SECURE JOURNEY IN THE CLOUD	15
	NEED MORE ASSURANCE?	16
	ENABLING LEGAL, REGULATORY, AND STANDARDS COMPLIANCE IN THE CLOUD	17
6	CONCLUSION	20
	WHAT CUSTOMERS SAY ABOUT SMT TECHNOLOGIES	20
	WHAT'S NEXT	20
	WHY NETAPP?	21
7	REFERENCES	23
	NETAPP SECURITY GUIDELINES AND BEST PRACTICES	23
8	ACKNOWLEDGEMENTS	24
	ENDNOTES	24

LIST OF TABLES

Table 1)	Potential impact definitions for security objectives	18
----------	--	----

LIST OF FIGURES

Figure 1)	Cloud reference model (from Cloud Security Alliance Guidance Version 2.1 [2009])	5
Figure 2)	SMT architecture overview (graphic supplied by Cisco Systems)	9
Figure 3)	Risk management framework (from NIST Special Publication 800-37)	15
Figure 4)	Mapping the cloud model to the security control and compliance model	19



1 INTRODUCTION

IT decision makers often find themselves of two minds when considering a move to a shared, or cloud-based, infrastructure. While the promise of cost efficiencies, scalability, and rapid provisioning compel a move to a cloud-based computing model, anxiety over one factor—security—often delays the move. Respondents to a September 2009 cloud user survey by the International Data Corporation (IDC) identified security as the leading challenge and issue inherent to the cloud model.¹ An April 2009 survey conducted by the European Network and the Information Security Agency (ENISA) cites concerns over data and service confidentiality, privacy, integrity, and availability as show-stoppers to cloud adoption.²

This paper examines information security concerns in the cloud and challenges the misperception that a shared infrastructure introduces risk that can be neither quantified nor controlled. Information assurance (IA) in the cloud is achievable through a holistic approach to information security that integrates people, processes, and technology into solutions that are aligned with an organization's business needs and risk threshold.

Through partnerships with industry leaders and built on best-of-breed network, compute, and storage technologies, NetApp and partner secure multi-tenancy technologies can alleviate customer anxiety and provide a secure journey to and through the cloud. At the storage layer, NetApp employs MultiStore[®] technologies that were rigorously evaluated as part of an independent security analysis conducted by Matasano Security in 2008. Simulating the efforts of "...a well-funded, highly motivated team of attackers willing to research and develop new exploits for previously unknown vulnerabilities," Matasano Security found no vulnerabilities that compromise the security model of the MultiStore feature.³



2 FRAMING THE CLOUD DISCUSSION

WHAT IS CLOUD COMPUTING?

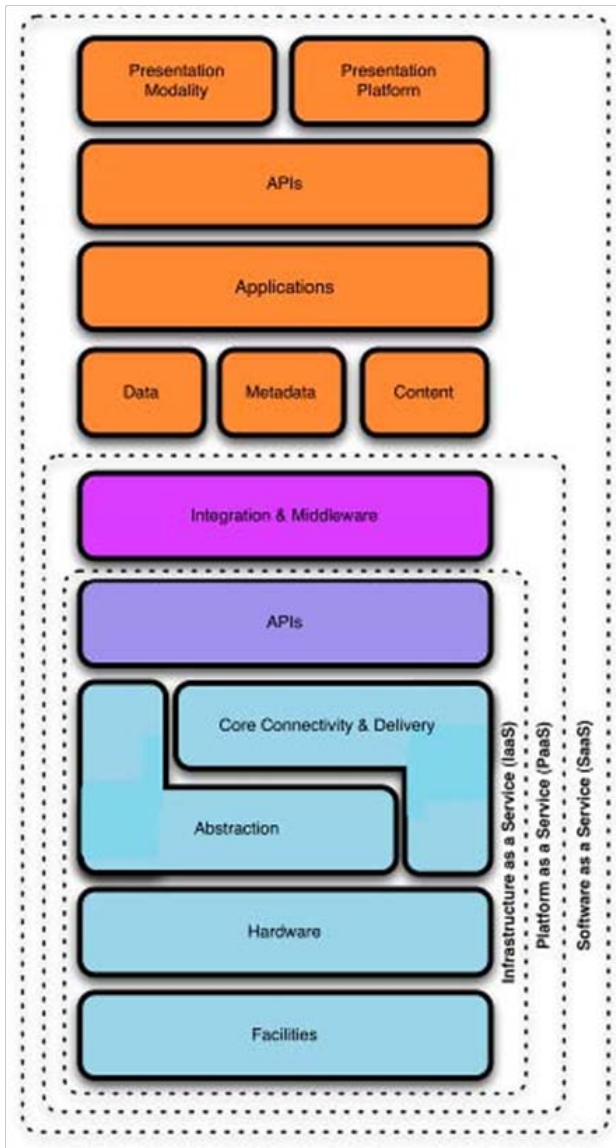
While myriad definitions and descriptions abound, the framework established by the National Institute of Standards and Technology (NIST) is widely accepted in both public and private sectors as a basis for emerging standards and as an acceptable starting point for cloud dialogue. NIST defines cloud computing as "...a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (for example, networks, servers, storage, applications, and services) that can be provisioned rapidly and released with minimal management effort or service provider interaction."⁴ The NIST cloud definition framework is composed of five essential characteristics, three service models, and four deployment models.

ESSENTIAL CHARACTERISTICS

- **On-demand self-service.** A consumer can unilaterally and automatically provision computing capabilities, such as server time and network storage, as needed without requiring human interaction with each service's provider.
- **Broad network access.** Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (for example, mobile phones, laptops, and PDAs).
- **Resource pooling.** The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (for example, country, state, or data center). Examples of resources include storage, processing, memory, network bandwidth, and virtual machines.
- **Rapid elasticity.** Capabilities can be rapidly and elastically provisioned, in some cases automatically, to quickly scale out, and can be rapidly released to quickly scale in. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.
- **Measured service.** Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (for example, storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the service.

SERVICE MODELS

Figure 1) Cloud reference model (from Cloud Security Alliance Guidance Version 2.1 [2009]).



The cloud reference model identifies the common and unique characteristics of these service models.⁵

- **Cloud software as a service (SaaS).** The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface such as a Web browser (for example, Web-based e-mail). The consumer does not manage or control the underlying cloud infrastructure, including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

- **Cloud platform as a service (PaaS).** The capability provided to the consumer is to deploy consumer-created or acquired applications, created using programming languages and tools supported by the provider, onto the cloud infrastructure. The consumer does not manage or control the underlying cloud infrastructure, including network, servers, operating systems, or storage, but has control of the deployed applications and possibly the application hosting environment configurations.
- **Cloud infrastructure as a service (IaaS).** The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, and deployed applications, and possibly limited control of select networking components (for example, host firewalls).

DEPLOYMENT MODELS

- **Private cloud.** The cloud infrastructure is operated solely for an organization. It can be managed by the organization or a third party and can exist on premise or off premise.
- **Community cloud.** The cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (for example, mission, security requirements, policy, and compliance considerations). It can be managed by the organizations or a third party and can exist on premise or off premise.
- **Public cloud.** The cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.
- **Hybrid cloud.** The cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (for example, cloud bursting for load-balancing between clouds).

WHAT IS MULTI-TENANCY AND WHY IS IT IMPORTANT?

Multi-tenancy is an architectural approach for sharing one or more infrastructures, databases, or application-level resources among multiple consumers, or tenants. In a virtualized environment, a tenant is any entity (application, individual, or organization), inside or outside the enterprise, that requires a secure and exclusive virtual computing environment. Although a one-size-fits-all definition is suitable for conceptual discussions, a more granular understanding of how each cloud service model (IaaS, PaaS, SaaS) affects multi-tenancy is important for identifying and managing risk consistent with an organization's business needs and risk threshold.

The three cloud computing service models are distinguished by which network, compute, and storage resources are shared and controlled by tenants, as illustrated in Figure 1. In an IaaS model, tenants control (subject to policy and design) processing power, an operating system, storage, deployed applications, and networking components without control of the underlying physical infrastructure. In a PaaS model, tenants control the applications that run in a shared hosting environment but not the underlying operating system, hardware, or network infrastructure. In a SaaS model, tenants share all or part of a core application but control none of the underlying platform or infrastructure.

Multi-tenancy is fundamentally important to both providers and consumers of cloud services. It is the foundation for achieving highly scalable, highly agile, cost-efficient, and computationally efficient architectures. These rewards are not without risk. Multi-tenancy, especially in a virtualized environment, employs layered technologies that require comprehensive security controls to provide the confidentiality, integrity, and availability of information. Aligned with the right people and processes, NetApp and partner SMT technologies provide this reliability with solutions that protect, detect, correlate, and react so that every journey into the cloud is risk-managed consistent with the business needs and the risk threshold.



3 THREATS, RISKS, AND BENEFITS OF CLOUD COMPUTING

Threat and risk are often used interchangeably; however, the distinction between threat and risk, whether inside or outside the cloud, is an important one. While threat describes an adverse scenario or outcome potentially applicable to a broad environment or population, risk measures the likelihood and impact that a given threat poses to a specific individual or organization. For example, while influenza poses a worldwide threat, the risk (probability and impact) of developing the flu is significantly less for individuals who receive flu vaccinations. Knowing what threats are out there is an important first step in differentiating acceptable risks from those that require mitigation.

WHAT ARE THE THREATS IN THE CLOUD?

The Cloud Security Alliance identifies seven top threats to cloud computing.⁶

- **Abuse and nefarious use of cloud computing.** Criminals leverage new technologies to improve reach, avoid detection, and improve the effectiveness of their activities. Cloud providers who employ weak registration systems and monitoring capabilities that facilitate anonymity are at risk.
- **Insecure interfaces and APIs.** Ignorance is not bliss. Consumers must know how their cloud provider integrates security with their service models and the security implications related to usage, management, orchestration, and monitoring of cloud services. Weak interfaces and APIs expose organizations to a variety of security issues related to confidentiality, integrity, availability, and accountability.
- **Malicious insiders.** Beware the enemy within. Broad access to internal resources enables malicious insiders to inflict grave damage to an organization's brand, finances, and productivity. Cloud consumers must know what measures are in place to protect against the insider threat.
- **Shared technology issues.** If security requirements are not baked in during the design phase, shared infrastructure technologies (for example, disk partitions, CPU caches, GPUs, and so on) can provide attack vectors that allow attackers to gain unauthorized access to data and impact the operations of other cloud customers.
- **Data loss or leakage.** Over 20TB of data, including sensitive data, has been exfiltrated from the Department of Defense, the defense industrial base, and civilian government organizations in recent years.⁷ Beyond financial, compliance, and legal implications, data loss or leakage can erode an organization's reputation as well as the morale and trust of employees, partners, and customers.
- **Account or service hijacking.** Stolen credentials, through techniques such as phishing and social engineering, can allow attackers to access critical areas of cloud infrastructure and services and compromise the confidentiality, integrity, and availability of those services.
- **Unknown risk profile.** Is your cloud provider transparent about the existence and effectiveness of security controls? What information will it disclose in the event of a security incident? Arguably, the greatest threat to a cloud consumer is the unknowing acceptance of risk due to a provider's veiled disclosure of internal security procedures, configuration hardening, patching, auditing, and logging.

CAN YOU AFFORD A BREAK IN THE CLOUD?

While the following examples of compromised data are not specific to cloud computing environments, they highlight the potential costs when any component of information assurance, including people, process, or technology, is the weakest link.

A 2008 breach at Heartland Payment Systems, one of the nation's largest providers of bankcard payment processing services, resulted in the theft of names, credit and debit card numbers, and expiration dates from an estimated 130 million bankcards.⁸ An annual study published by the Ponemon Institute, a research center dedicated to privacy, data protection, and information security policy, reports that the average total cost of a data breach rose to \$6.65 million in 2009.⁹ The following 2010 midyear statistics compiled by the Privacy Rights Clearinghouse show no evidence that this trend is slowing:¹⁰

- **1.2 million.** In January 2010, Lincoln National Corporation notified state authorities of a vulnerability uncovered within its portfolio management system. Users were sharing passwords and thereby exposing the personal information of about 1.2 million customers.
- **1 million.** BlueCross BlueShield of Tennessee reported that the theft of 57 hard drives from a training facility in October 2009 exposed the private information of 1 million customers in multiple states. The hard drives contained customers' personal data and protected health information, which was encoded but not encrypted.
- **208,000.** AvMed Health Plans announced in February 2010 that the theft of two company laptops from its corporate offices may have compromised the personal information of some current and former subscribers.
- **3.3 million.** Educational Credit Management Corporation, a guarantor of federal student loans, reported in March 2010 that portable media with personally identifiable information on 3.3 million borrowers was stolen from its headquarters.
- **409,262.** In April 2010, Affinity Health Plan, a New York managed care service, notified more than 400,000 current employees, former employees, and others that their personal data might have been exposed through the loss of a digital copier hard drive that had not been erased.

What about the costs that can't be measured in dollars alone? In 2008, classified military computer networks at the U.S. Department of Defense were compromised when an infected flash drive was inserted into a U.S. military laptop at a base in the Middle East. Embedded by a foreign intelligence agency, the malicious computer code uploaded itself onto a U.S. Central Command network, spread undetected throughout classified and unclassified systems, and established "... a digital beachhead from which data could be transferred to servers under foreign control."¹¹ The intrusion realized the worst fears of network administrators and military commanders alike: "...a rogue program operating silently, poised to deliver operational plans into the hands of an unknown adversary."¹²

In 2009, the Transportation Security Administration accidentally posted an airport screening procedures manual onto a public Web site. The procedures included "...details for screening passengers; checking for explosive devices; special rules for handling the CIA, diplomats, and law enforcement officials; and the technical settings and tolerances used by metal and explosive detectors used at airports."¹³ In another 2009 incident, a document marked by the President as "Highly Confidential Safeguards Sensitive" was posted on the public Web site of the Government Printing Office. The documents contained "...detailed information on hundreds of civilian nuclear sites...including those storing enriched uranium.... [along with] details on programs at nuclear weapons research labs at Los Alamos, Livermore, and Sandia."¹⁴

In these incidents, dollar costs are comparatively inconsequential to costs measured in lives lost, damage to national security, or harm to our nation's critical infrastructure.



4 HELPING TO ASSURE THE CLOUD WITH SMT

Secure Multi-Tenancy (SMT) technologies enable a cloud-computing architecture consistent with the IaaS model as defined in the NIST Cloud Definition Framework. SMT provides security capabilities not commonly found in the IaaS model by delivering integrated and effective security measures that scale with the cloud.

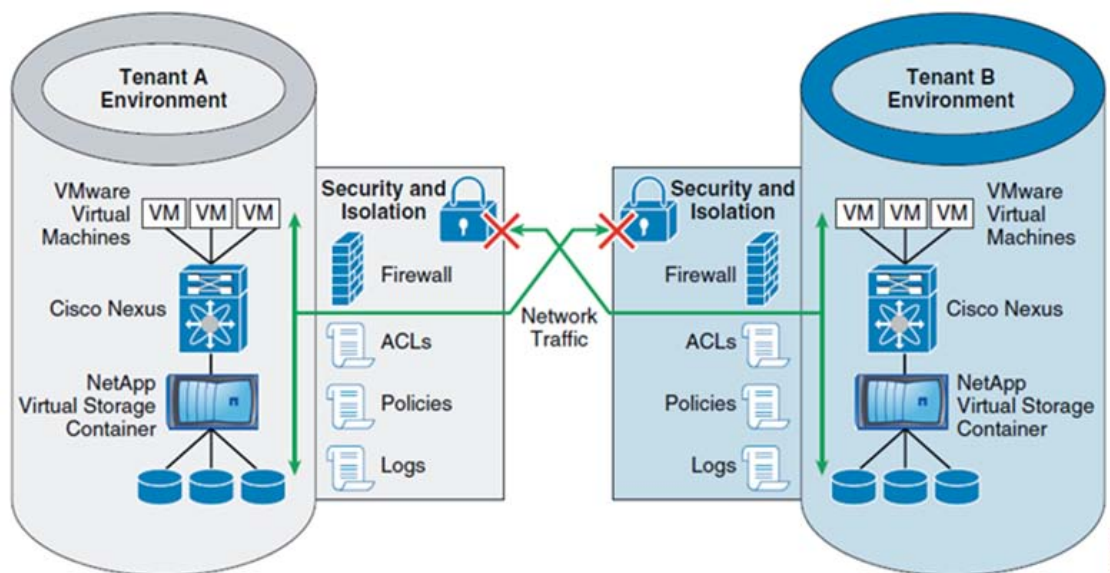
Beyond logical separation—a prerequisite for all virtualized multi-tenant environments—SMT technologies provide a framework for supplemental security hardening commensurate with an organization's business needs, security objectives, and risk tolerance. SMT sets the foundation for a secure cloud architecture that supports the:

- Logical separation of tenants at all layers (network, compute, and storage)
- Flexible integration of supplemental security controls (management, operational, and technical) that protect the confidentiality, integrity, and availability of the system and its information
- Attainment of information system Certification & Accreditation (C&A) with respect to a prescribed compliance model (for example, FedRAMP, DIACAP, PCI-DSS, HIPAA).

SMT technologies enable customers to reap the cost savings and efficiency benefits of a shared infrastructure without relinquishing control or compromising security. SMT is designed specifically for customers who demand inviolable methods for partitioning resources at the network, compute, and storage layers of multi-tenant environments.

Figure 2 illustrates the SMT architecture overview.

Figure 2) SMT architecture overview (graphic supplied by Cisco Systems).



THE FOUR PILLARS OF SMT

The key to developing a robust design is to clearly define the requirements and to apply proven methodology and design principles. The following four pillars are the foundation for SMT:

- **Secure Separation** makes sure that one tenant does not have access to another tenant's resources, such as the virtual machine, network bandwidth, and storage. Each tenant must be securely separated using techniques such as access control, Virtual LAN (VLAN) segmentation, and virtual storage controllers. Also, each layer has its own means of enforcing policies that help reinforce the policies of the adjacent layers.

- **Service Assurance** provides isolated compute, network, and storage performance during both steady state and nonsteady state. For example, the network can provide each tenant with a certain bandwidth guarantee using quality of service (QoS). Resource pools within VMware® help balance and guarantee CPU and memory resources, while FlexShare® can balance resource contention across storage volumes.
- **Service Availability** allows the infrastructure to meet the expectation of compute, network, and storage to always be available even in the event of failure. Like the secure separation pillar, each layer has its own manner of providing a high-availability configuration that works seamlessly with adjacent layers. Security and availability are best deployed from a layered approach.
- **Service Manageability** is required to rapidly provision and manage resources and view resource availability. In its current form, each layer is managed by vCenter™, the Unified Computing System® (UCS®) Manager, DC Network Manager, and NetApp Operations Manager, respectively.

SMT: HELPING TO ASSURE INFORMATION AT THE NETWORK, COMPUTE, AND STORAGE LAYERS

The Committee on National Security Systems (CNSS) Instruction No. 4009, National Information Assurance Glossary, defines IA as “[m]easures that protect and defend information and information systems by ensuring their availability, integrity, authentication, confidentiality, and non-repudiation...[to include] providing for restoration of information systems by incorporating protection, detection, and reaction capabilities.”¹⁵

IA is a holistic approach to information security that seeks to balance and integrate people, processes, and technology through a defense-in-depth strategy based on organizational business needs and risk thresholds, defense in multiple places, layered defenses, persistent data protection, and infrastructures that protect, detect, correlate, and react to threats. The three essential elements of IA are confidentiality, integrity, and availability.

CONFIDENTIALITY

Confidentiality refers to preserving authorized restrictions on information access and disclosure, including means for protecting personal privacy and proprietary information [44 USC 3542 (b)(1)(B)].

At the network layer, SMT helps to assure confidentiality through:

- **VLANs** that logically segment and secure switched network traffic according to organizational policy. VLAN segmentation allows the network to be reconfigured according to the needs of individual tenant groups through logical methods instead of physically moving network hardware and connections. VLAN technology logically segments the switched infrastructure into separate Layer-2 broadcast domains, enabling service providers to assign one or more discrete network address spaces to each tenant and to manage them as an organizational entity.
- **Nexus 1000V** switches that provide security services to the individual virtual machine as well as policy portability, so the proper network and security policies follow every virtual machine as it moves around the data center. This isolation allows individual access restrictions to be placed on each virtual machine as if it were a completely separate physical machine.
- **Virtual Device Contexts (VDCs)** that secure and isolate tenants by enabling the partitioning of a single physical switch into four virtual switches. Each VDC operates similar to a stand-alone switch with a distinct configuration file, a complement of physical ports, and separate instances of Layer-2 and Layer-3 services.

At the compute layer, SMT helps to assure confidentiality through:

- **Virtual routing and forwarding (VRF)** that provides the capability to direct traffic independently through multiple VDCs, providing more granular Layer-3 isolation within a VDC for the cloud tenants and their independent application stacks.

- **IP and MAC-based access controls (ACLs)** that provide stateless traffic filtering based on an IP address or the Media Access Control (MAC) physical hardware address. This feature enables control and monitoring of virtual machine traffic flows within the virtual and physical cloud infrastructure.
- **vShield Zones** that provide data center administrators with greater visibility and enforcement of network activity to provide the necessary level of isolation and means for compliance for each tenant. Consistent with effective ports, protocols, and services management (DODI 8551.1), vShield Zones provide granular virtual firewall capabilities that can detect rogue services, prohibit virtual machine communication, and serve as a regulatory compliance visualization tool.
- **VMsafe** technology that helps to protect the virtual infrastructure in ways previously not possible with physical machines. VMsafe provides a unique capability for virtualized environments through an application program interface (API)—sharing program that enables partners to develop security products for VMware environments. This allows direct integration of firewalls, intrusion detection, data leakage protection, and other security capabilities.

At the storage layer, SMT helps to assure confidentiality through:

- **NetApp MultiStore technology** that enables cloud providers to quickly and easily create separate and completely private logical partitions on a single NetApp storage system as discrete administrative domains called vFiler™ units. Each vFiler unit encapsulates both the data and the administrative functions for a given tenant and is restricted to the VLANs associated with that tenant. Therefore, even the tenant administrator (who has root privilege on his or her vFiler unit) cannot connect to another tenant's vFiler unit, let alone access the data managed by it.
- **NetApp DataFort™ appliances and Brocade Encryption Switches (BES)** that combine secure access controls, authentication, hardware-based encryption, and secure logging to help protect stored data. These appliances can be transparently deployed inline to gain security advantages without impacting user workflow. NetApp DataFort fits seamlessly into existing storage environments, supporting wire-speed encryption and compression for CIFS, NFS, IP SAN, FC disk, FC tape, and iSCSI protocols. NetApp DataFort protects data both at rest and in flight with encryption certified by the NIST.
- **NetApp SnapLock® technology** that allows organizations to meet the strictest data retention regulations and internal IT governance policies. SnapLock aids in compliance with records retention regulations that require archival of e-mails, documents, audit information, and other data in an unalterable state for years. It is especially valuable for retrieving unregulated yet crucial reference data that has not been changed or deleted but that must be accessed quickly. SnapLock creates nonrewritable, nonerasable volumes to prevent files from being altered or deleted until a set retention date. NetApp allows an administrator to back up this Write Once Read Many (WORM) data to disk or tape for an additional level of data protection.

INTEGRITY

Integrity refers to guarding against improper information modification or destruction, and includes assuring information nonrepudiation and authenticity [44 USC 3542 (b)(1)(A)].

At the network layer, SMT helps to assure integrity through:

- **Control Plane Policing (CoPP)**, a Cisco® feature that allows users to manage the flow of traffic handled by the route processor of their network devices. CoPP is designed to prevent unnecessary traffic from overwhelming the route processor that, if left unabated, could affect system performance. This can protect all tenants of the cloud environment by maintaining the integrity of the cloud infrastructure.
- **TrustSec and SAN fabric encryption** technologies that create a cloud of trusted fabric devices. The components of the Cisco cloud architecture authenticate themselves to one another, allowing encrypted transport between devices, packet classification, and access control among other services. Individually, these capabilities support the security requirements of cloud computing. Combined, they can be leveraged across a cloud architecture to provide an extremely robust and tailored security solution safeguarding the assets of each tenant.

At the storage layer, NetApp can further help to provide integrity with:

- **NetApp RAID-DP[®]** that significantly increases the fault tolerance from failed disk drives over traditional single-parity RAID. When all relevant numbers are plugged into the standard mean time to data loss (MTTDL) formula for RAID-DP versus single-parity RAID, RAID-DP is thousands of times more reliable on the same underlying disk drives. With this reliability, RAID-DP exceeds even RAID 10 mirroring for fault tolerance, but at RAID 4 pricing. RAID-DP offers businesses the most compelling total cost of ownership storage option without putting their data at an increased risk.
- **SnapLock technology** enables organizations to meet the strictest data retention regulations and internal IT governance policies. SnapLock aids in compliance with records retention regulations that require archival of e-mails, documents, audit information, and other data in an unalterable state for years. It is especially valuable for retrieving unregulated yet crucial reference data that has not been changed or deleted but must be accessed quickly. SnapLock creates nonrewritable, nonerasable volumes to prevent files from being altered or deleted until a set retention date. NetApp allows an administrator to back up this WORM data to disk or tape for an additional level of data protection.
- **Checksumming, mirroring, and cyclical redundancy checks (CRCs)**. Checksumming is built into the WAFL[®] file system. All data written is checksummed at a disk level. Data integrity is maintained by matching all read data to checksums. CRCs are part of the Fibre Channel protocol. CRC values are generated for all packets sent over the Fibre Channel network. These values are then sent with the packets over the network to be checked at the far end for integrity. Data within volumes on a storage system can be mirrored to another storage system for backup and disaster recovery purposes to provide data integrity.

AVAILABILITY

Availability refers to the timely and reliable access to, and use of, information [44 USC 3542 (b)(1)(C)].

At the network layer, SMT helps to assure availability through:

- **EtherChannel** (dual-Gigabit Ethernet [10GbE]) that connects the UCS 6120 Fabric Interconnects and NetApp FAS storage controllers with the Nexus 5000 access switch. Multiple active paths provide increased redundancy and bandwidth with a lower required port count.
- **virtual PortChannel (vPC)** technology that enables two separate devices to appear and act as a single logical PortChannel endpoint. The benefits of hardware redundancy are combined with the benefits of PortChannel loop management.
- **Device/Link redundancy** that provides continued availability in the event of a single network device or link failure.
- **Active/Passive Virtual Supervisor Module (VSM) redundancy** that is supported by the Nexus 1000V and the Nexus 1010 Virtual Services Appliance. If a primary VSM goes offline, a standby VSM is ready to take over.

At the compute layer, SMT helps to assure availability through:

- **VMware high availability** that allows continuous monitoring of a virtual environment to detect operating system and hardware failures and automatic restart of virtual machines onto another physical host if a problem is detected. With functionality built directly into the hypervisor, this feature limits downtime due to software error or hardware failure and enables cloud providers to provide strong uptime service-level agreements (SLAs).
- **Fault tolerance (FT)** that keeps hardware failures from having an impact on the virtual machine. By running two synchronized instances of a virtual machine on separate physical hosts, the failure of a primary virtual machine host results in a seamless and instant failover to the secondary virtual machine.

At the storage layer, SMT helps to assure availability through:

- **High availability system configuration** that enables each NetApp storage controller to service its own workload while monitoring the other controller. If a controller fails, the other one seamlessly takes

over the additional workload. This failover is transparent to users and applications. A manual failover can also be triggered to perform scheduled maintenance and upgrades. This allows a cloud provider to stay up to date and compliant while continuing to provide superior storage availability.

- **RAID-DP** technology that immunizes against downtime or data loss and enables each tenant to use the infrastructure with confidence that all application and underlying operating system data remains consistent and protected from physical storage failures. More reliable than RAID 1 mirroring and statistically over 10,000 times more reliable than a single-parity solution, RAID-DP technology can sustain two simultaneous disk drive failures.
- **NetApp Snapshot[®] copies** that protect against accidental deletion or corruption by allowing a consistent view of the file system to be captured instantaneously with virtually no storage overhead. These copies can be created manually or automatically on a regular schedule set by the administrator. Because multiple Snapshot copies can be taken and saved per volume, tenants can schedule hourly, daily, weekly, and monthly Snapshot copies to create a schedule that best suits their individual data protection needs. Files can be recovered by accessing the snapshot directly without burdening the storage administrator or by invoking a NetApp SnapRestore[®] operation to roll back part or all of the active file system.
- **NetApp SnapMirror[®]** data replication features that efficiently mirror data, applications, and their underlying operating systems between two NetApp storage controllers to provide an off-site disaster recovery infrastructure and maintain continuity in the event of a site-wide failure.



5 APPLYING THE RISK MANAGEMENT FRAMEWORK

Beyond managing technology, information assurance is about managing risk. NIST Special Publication (SP) 800-53, *Recommended Security Controls for Federal Information Systems and Organizations*, defines risk management as “[t]he process of managing risks to organizational operations (including mission, functions, image, reputation), organizational assets, individuals, other organizations, and the Nation, resulting from the operation of an information system, and includes: (1) the conduct of a risk assessment; (2) the implementation of a risk mitigation strategy; and (3) employment of techniques and procedures for the continuous monitoring of the security state of the information system.”¹⁶

In a cloud environment, the greatest risks await those who make the transition to a shared infrastructure without realizing how the selection of a cloud service model, a deployment model, and a service provider impacts risk. Effective risk management goes beyond the “known knowns” and examines the “known unknowns”—known risks not quantified in terms of probability of occurrence or impact to business. At a minimum, a shared infrastructure risk assessment must seek to quantify the impact or harm to an enterprise if:

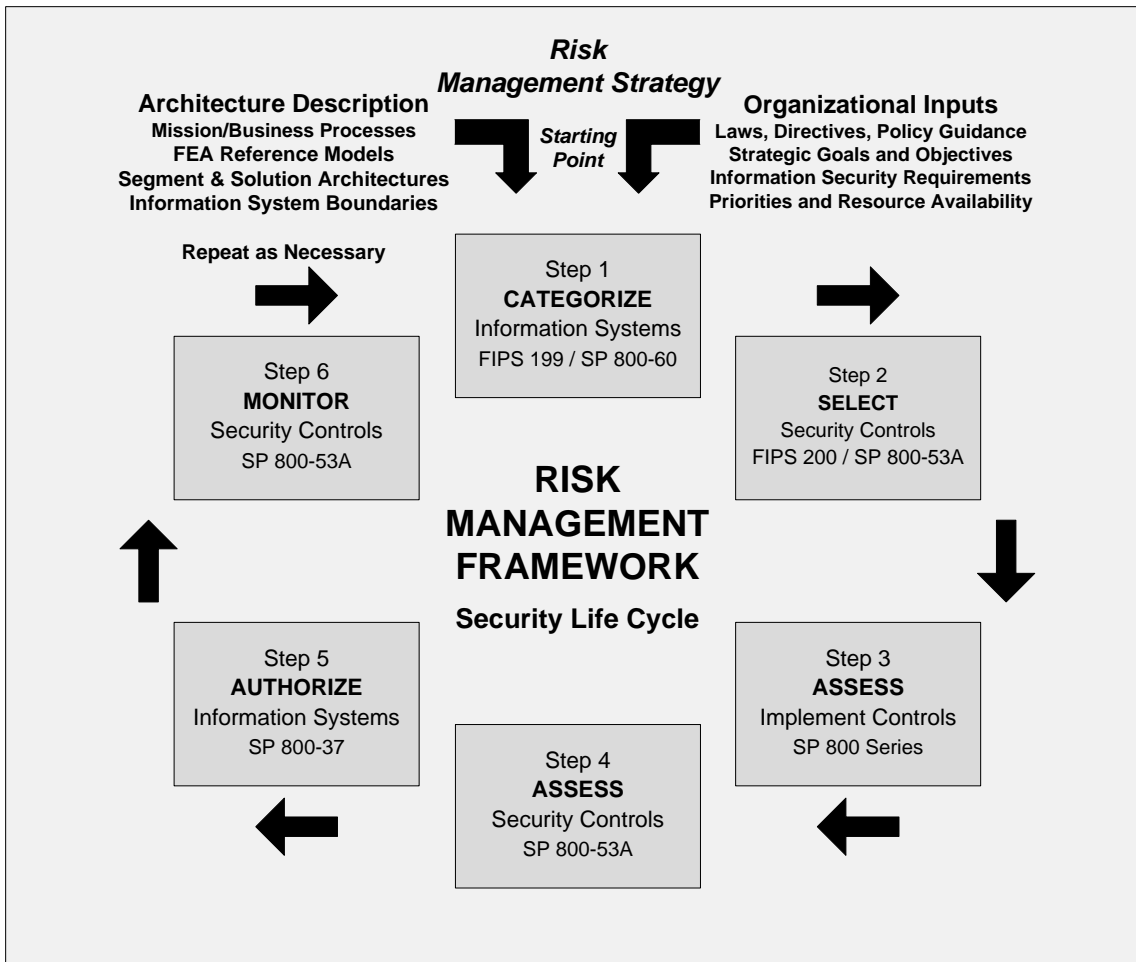
- An information asset becomes widely public and widely distributed.
- An employee of the cloud provider accesses the asset.
- The process or function is manipulated by an outsider.
- The process or function fails to provide expected results.
- The information and data are unexpectedly changed.
- The asset is unavailable for a period of time.

The NIST “Guide for Applying the Risk Management Framework (RMF) to Federal Information Systems” provides an effective construct for this examination. It allows both public and private sector organizations to rigorously examine their information security posture and determine if the level of residual risk is acceptable and consistent with the business needs and risk threshold.

The RMF is a departure from historical methods that assessed information system security (ISS) based on point-in-time snapshots. The RMF’s lifecycle approach to security maintains that constantly evolving threats and threat actors demand a dynamic versus static approach to ISS risk management based on continuous monitoring. The RMF is composed of six lifecycle activities:

- **Categorize** the information system and the information processed, stored, and transmitted by that system based on an impact analysis.
- **Select** an initial set of baseline security controls for the information system based on the security categorization and tailoring. Supplement the security control baseline as needed based on an organizational assessment of risk and local conditions.
- **Implement** the security controls and describe how the controls are employed within the information system and its environment of operation.
- **Assess** the security controls using appropriate assessment procedures to determine the extent to which the controls are implemented correctly, operating as intended, and producing the desired outcome with respect to meeting the security requirements for the system.
- **Authorize** information system operation based on a determination of the risk to organizational operations and assets, individuals, other organizations, and the nation resulting from the operation of the information system and the decision that this risk is acceptable.
- **Monitor** the security controls in the information system on an ongoing basis including assessing control effectiveness, documenting changes to the system or its environment of operation, conducting security impact analyses of the associated changes, and reporting the security state of the system to designated organizational officials.

Figure 3) Risk management framework (from NIST Special Publication 800-37).



HOW CONTROLS PROVIDE A MORE SECURE JOURNEY IN THE CLOUD

While no journey through the real or virtual clouds is without risk, appropriately applied security controls permit the mitigation of risk to acceptable levels.

On average, over 87,000 flights transit the airspace over the United States on any given day.¹⁷ Despite this volume of traffic and the complexities of airspace deconfliction, security controls inherent to aircraft design and enhanced by the Air Traffic Control (ATC) services, controlled airspace, and navigational aids of the National Airspace System (NAS) mitigate potential risks and make commercial airline travel one of the world’s safest modes of transportation.

Appropriate security controls are no less essential for providing a safe and secure journey through a virtual cloud environment. The “2010 Data Breach Investigations Report” published by Verizon’s Business RISK Team in cooperation with the U.S. Secret Service (USSS) concluded that 96% of the data breaches examined in 2009 were avoidable through simple or intermediate controls.¹⁸

The primary purpose of information system security controls is to enable and monitor compliance with industry-specific standards. In the public sector, security controls are designed to assure compliance with Federal Information Processing Standards (FIPS) in accordance with the Federal Information Security Act (FISMA). In the commercial sector, the Payment Card Industry Data Security Standard (PCI DSS) is an example of a compliance model derived from collaboration between Visa and MasterCard to create

common industry security requirements. Regardless of the sector or compliance model in question, the same questions apply:

- What security controls are needed to adequately mitigate the risk incurred by the use of information and information systems in the execution of organizational missions and business functions?
- Have the selected security controls been implemented or is there a realistic plan for their implementation?
- What is the desired or required level of assurance, that is, grounds for confidence, to which the selected security controls, as implemented, are effective in their application?

Implementation, assessment, and monitoring of information system security controls in accordance with prescribed objectives for confidentiality, integrity, and availability allow an organization to determine if residual risk posture is aligned with business risk tolerance. Per NIST SP 800-53, continuous monitoring informs an organization when a reassessment of security controls is warranted based on trigger events such as:

- An incident that results in a breach of the information system, producing a loss of confidence by the organization in the confidentiality, integrity, or availability of information processed, stored, or transmitted by the system;
- A newly identified, credible, information system–related threat to organizational operations and assets, individuals, other organizations, or the nation is identified based on intelligence information, law enforcement information, or other credible sources of information;
- Significant changes to the configuration of the information system through the removal or addition of new or upgraded hardware, software, or firmware or changes in the operational environment potentially degrade the security state of the system; or
- Significant changes to the organizational risk management strategy, information security policy, supported missions and/or business functions, or information being processed, stored, or transmitted by the information system.¹⁹

NEED MORE ASSURANCE?

In 2008, Matasano Security conducted an independent security analysis of NetApp MultiStore, the foundation for SMT secure separation at the storage layer, and concluded the following:

“As a result of the technology we were testing (storage protocols) and the circumstances of the test (mission-critical isolation features), our assessment committed extraordinary efforts to the attempt to break the security of the MultiStore feature. Rather than relying on well-known vulnerabilities and off-the-shelf testing tools, our test simulated the efforts of a well-funded, highly motivated team of attackers willing to research and develop new exploits for previously unknown vulnerabilities. Our testing spanned all the major storage protocols supported by NetApp MultiStore, including CIFS, NFS, and iSCSI. At the conclusion of the test, our team can report that we know of no vulnerabilities that compromise the security model of the MultiStore feature:

- We know of no software flaws in the NetApp implementations of CIFS, NFS, or iSCSI that would allow attackers to exploit common C-code flaws like buffer overflows, integer overflows, or race conditions to execute code remotely in a FAS Storage System.
- We know of no architectural flaws in the storage protocols supported by MultiStore or their management interfaces that would allow attackers to use access to a vFiler unit to reconfigure the FAS Storage System itself or any other associated vFiler units.
- We know of no protocol vulnerabilities in CIFS, NFS, or iSCSI that would allow an attacker to use a connection to their own vFiler unit to gain access to storage resources on other vFiler units, such as iSCSI LUNs or CIFS shares.
- We know of no vulnerabilities in the TCP/IP stack of the FAS Storage System that would allow attackers to bridge traffic from untrusted networks to trusted networks.”²⁰

ENABLING LEGAL, REGULATORY, AND STANDARDS COMPLIANCE IN THE CLOUD

In both public and commercial sectors, legal and regulatory mandates drive the need for enhanced information security. The following key legislation influences how highly regulated entities move to the cloud:

- **Federal Information Security Management Act.** FISMA requires each federal agency to “...develop, document, and implement an agency-wide program to provide information security for the information and information systems that support the operations and assets of the agency, including those provided or managed by another agency, contractor, or other source.”
- **Federal Information Processing Standard 200.** FIPS 200 is a mandatory federal standard developed by NIST in response to FISMA. Compliance requires that organizations first determine the security category of their information system in accordance with FIPS 199, *Standards for Security Categorization of Federal Information and Information Systems*, derive the information system impact level from the security category in accordance with FIPS 200, and then apply the appropriately tailored set of baseline security controls in NIST SP 800-53, *Security Controls for Federal Information Systems and Organizations*.
- **Payment Card Industry Data Security Standard.** The PCI DSS was developed by the founding payment brands of the PCI Security Standards Council, including American Express, Discover Financial Services, JCB International, MasterCard Worldwide, and Visa Inc. to help facilitate global adoption of consistent data security measures and to help organizations proactively protect customer account data.²¹
- **Health Insurance Portability and Accountability Act of 1996 (HIPAA).** Technical safeguards outlined in HIPAA section 164.312(a)(1) require that health providers “Implement technical policies and procedures for electronic information systems that maintain electronic protected health information to allow access only to those persons or software programs that have been granted access rights...”²²
- **Financial Modernization Act of 1999, also known as the “Gramm-Leach-Bliley Act” (GLBA).** GLBA provides the protection of nonpublic customer information and affords customers “... informational and nondisclosure rights with respect to the sharing of customer information by financial services organizations, requires the federal banking and securities agencies to adopt customer privacy regulations, and prohibits ‘pretext calling’ [i.e., social engineering practices] and other actions to obtain customer information from a financial institution under false pretenses.”²³
- **Sarbanes Oxley Act of 2002 (SOX).** SOX mandated a number of reforms to enhance corporate responsibility, enhance financial disclosures, and combat corporate and accounting fraud. Section 404 focuses on internal controls: those processes “...effected by an entity’s board of directors, management, and other personnel, designed to provide reasonable assurance regarding the... effectiveness and efficiency of operations, reliability of financial reporting, and compliance with applicable laws and regulations.”²⁴
- **Statement on Auditing Standards (SAS) No. 70.** Developed by the American Institute of Certified Public Accountants (AICPA), SAS 70 defines the standards an auditor must employ to assess the internal controls of a service organization. SAS 70 enables service organizations to demonstrate—through an in-depth audit of their control objectives and control activities—adequate safeguards for hosting and processing customer data.²⁵

In conjunction with the Federal Information Processing Standards Publication 199–based impact analysis (refer to Table 1) performed during Step 1 (Categorize) of the RMF, these regulations and standards establish the minimum levels of confidentiality, integrity, and availability required for compliance.

Table 1) Potential impact definitions for security objectives.²⁶

Security Objective	Potential Impact		
	Low	Moderate	High
<p>Confidentiality Preserving authorized restrictions on information access and disclosure, including means for protecting personal privacy and proprietary information. [44 U.S.C., SEC. 3542]</p>	<p>The unauthorized disclosure of information could be expected to have a limited adverse effect on organizational operations, organizational assets, or individuals.</p>	<p>The unauthorized disclosure of information could be expected to have a serious adverse effect on organizational operations, organizational assets, or individuals.</p>	<p>The unauthorized disclosure of information could be expected to have a severe or catastrophic adverse effect on organizational operations, organizational assets, or individuals.</p>
<p>Integrity Guarding against improper information modification, or destruction, and includes ensuring information non-repudiation and authenticity. [44 U.S.C., SEC. 3542]</p>	<p>The unauthorized modification or destruction of information could be expected to have a limited adverse effect on organizational operations, organizational assets, or individuals.</p>	<p>The unauthorized modification or destruction of information could be expected to have a serious adverse effect on organizational operations, organizational assets, or individuals.</p>	<p>The unauthorized modification or destruction of information could be expected to have a severe or catastrophic adverse effect on organizational operations, organizational assets, or individuals.</p>
<p>Availability Ensuring timely and reliable access to and use of information. [44 U.S.C., SEC. 3542]</p>	<p>The disruption of access to or use of information or an information system could be expected to have a limited adverse effect on organizational operations, organizational assets, or individuals.</p>	<p>The disruption of access to or use of information or an information system could be expected to have a serious adverse effect on organizational operations, organizational assets, or individuals.</p>	<p>The disruption of access to or use of information or an information system could be expected to have a severe or catastrophic adverse effect on organizational operations, organizational assets, or individuals.</p>

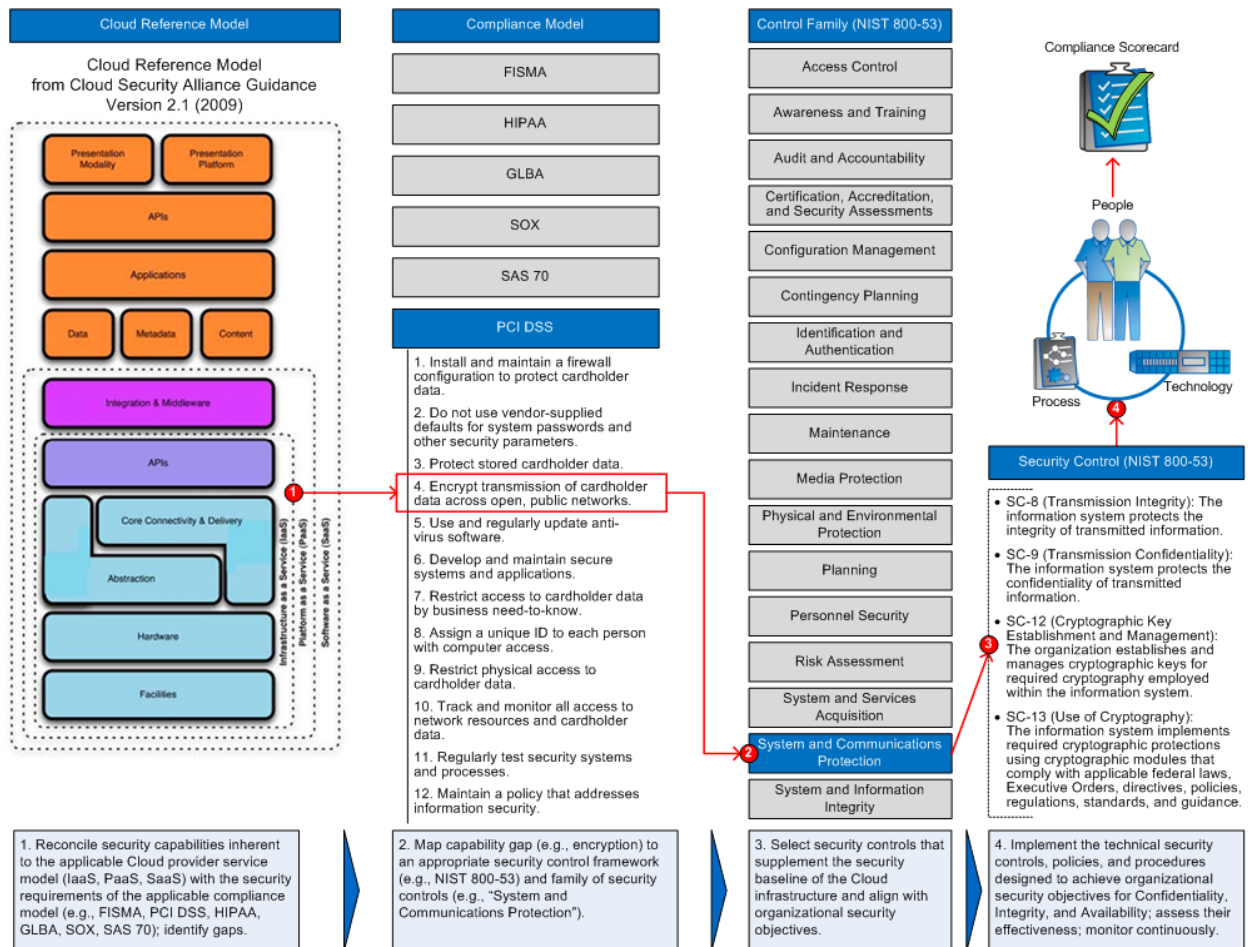
NIST SP 800-53 covers the spectrum of potential impacts (low, moderate, high) for each security objective (confidentiality, integrity, and availability) and provides a comprehensive framework that maps effectively to key compliance models (FISMA, PCI DSS, HIPAA, GLBA, SOX, and so forth) both within and outside the public sector.

Categorizing the information to be processed, stored, and transmitted by a system is an essential prerequisite to selecting (Step 2 – RMF) and implementing (Step 3 – RMF) adequate security controls for that system. In the context of a cloud computing environment, an understanding of the baseline security controls inherent to the applicable cloud service model is equally important.

As shown in Figure 4, to achieve regulatory and standards compliance in the cloud, organizations must:

1. Reconcile security capabilities inherent to the applicable cloud provider service model (IaaS, PaaS, SaaS) with the applicable compliance model (for example, FISMA, PCI DSS, HIPAA) security requirements to identify gaps that impose risk.
2. Map capability gaps (for example, encryption) to an appropriate security control framework (for example, NIST SP 800-53) and family of security controls (for example, "System and Communications Protection").
3. Select security controls that supplement the security baseline of the cloud infrastructure and align with organizational security objectives.
4. Implement the technical security controls, policies, and procedures designed to achieve organizational security objectives for confidentiality, integrity, and availability; assess their effectiveness; and monitor them continuously.

Figure 4) Mapping the cloud model to the security control and compliance model.





6 CONCLUSION

WHAT CUSTOMERS SAY ABOUT SMT TECHNOLOGIES

Beyond robust security controls that provide high levels of confidentiality, integrity, and availability, NetApp and its partners deliver measurable business value to data center cloud environments. Here is what our customers say:

“The Cisco-NetApp-VMware infrastructure is a game-changer. We can run our entire business on an integrated and verified platform that we deployed in our data center in just a matter of hours. This infrastructure is where everybody’s going to be in 2012, but we’re already here.”

– Brian Denton, CTO, ExamWorks

“The big enablers of [T-Systems] Dynamic Services are VMware on the server side and NetApp on the storage side. NetApp technology contributes to each of our competitive advantages: lower cost, greater flexibility, and higher SLAs.”

– Dr. Stefan Bucher, Global Delivery Manager, T-Systems

“NetApp technology, combined with low TCO, helps us offer advantages that our competitors don’t have or can’t deliver as efficiently. NetApp deduplication, for example, lets us use 50% less capacity for virtual machines, with corresponding savings in real estate, power, and cooling. This lowers our cost of goods sold. Passing our savings along makes CartikaCloud services accessible to more businesses.”

– Andrew Rouchotas, Chief Executive Officer, Cartika, Inc.

“We’d have 300 physical servers today instead of 80 if not for virtualization. By supporting our virtual infrastructure the NetApp solution is helping us save at least \$2 million over a three- to five-year period.”

– Chris Rima, Supervisor of Infrastructure Systems, Tucson Electric Power

“With NetApp MultiStore, our customers can have their own private storage island provisioned in about an hour with all the features of an individual, dedicated SAN: DR, scalability, and more.”

– Jared Wray, Principal, Tier 3

“Our clients realize that with NetApp we’re giving them not just lip service but a real DR solution that can recall their data within an hour. It virtually guarantees their business continuity.”

– Remy O’Keefe, Business Manager, Tier 3

WHAT’S NEXT

In addition to currently available storage security systems that deliver nondisruptive, comprehensive protection for sensitive data across your enterprise, NetApp is preparing to launch NetApp Storage Encryption (NSE). NSE embeds encryption in the storage device itself, a practice known as full disk encryption (FDE). FDE provides security to “data at rest” and the added benefits of simplified key management (based on the Key Management Interoperability Protocol [KMIP]), negligible performance impact, and encryption capabilities that automatically scale as storage is added to the system.

NetApp and its partners plan to publish a security hardening guide for select SMT architectures designed to pass an independent third-party audit based on the SysAdmin, Audit, Network, Security (SANS) Institute’s “Top 20” critical controls. This guide outlines best-practice configurations at the network, storage, and compute levels that establish a strong security controls baseline for agencies and organizations pursuing certification and accreditation of a cloud computing solution based on FIPS199 security objectives and NIST SP 800-53 security controls.

WHY NETAPP?

The “Best Practices for Storage Networks” white paper, published by the Systems and Network Analysis Center at the National Security Agency (NSA), states that “the greatest security benefits can be provided by components that have achieved the highest levels of certification....”²⁷ NetApp agrees and has a body of evidence that proves our commitment to information security.

CERTIFICATIONS

- Internet Protocol Version 6 (IPv6)1
- Data ONTAP® 7.3.1 operating system and FAS, V-Series, SA, R200, and IBM N series family of data storage controller systems
 - Tested on DoD DISR IPv6 Approved Products List (APL) for the DoD IPv6 Standards Profiles for IPv6 Capable Products, Version 4 (March 2009)
- Data ONTAP 7.3.3 operating system and FAS, V-Series, SA, R200, and IBM N series family of data storage controller systems
 - Tested as conformant and interoperable with IPv6 Ready Logo Phase 2 host requirements and on the Approved Products List (July 2010)
 - Tested and declared conformant to the IPv6 host profile in NIST, “A Profile for IPv6 in the U.S. Government - Version 1.0” (August 2010)

Common Criteria Evaluation Assurance Level (EAL)

- EAL 4+ for NetApp DataFort FC520 v2, LKM 2.5.1
- EAL 3+ for Data ONTAP 7.3.1.1
- EAL 2+ for Data ONTAP 7.2.5.1
- EAL 2 for Data ONTAP 7.0.3 and 7.0.4
- EAL 2 for Data ONTAP 6.5.2R1

FIPS 140-2, Level 33

- NetApp DataFort Storage Encryption Processor (SEP)
- NetApp DataFort SAN SEP v2.0
- NetApp DataFort NAS SEP v1.0
- NetApp DataFort LKM SEP v1.0
- NetApp DataFort SCSI SEP v1.0

ISO/IEC 27001 CERTIFICATION

W. E. Deming said, “If you can't describe what you are doing as a process, you don't know what you're doing.” NetApp knows information security and has the ISO 27001 certification to prove it.

ISO/IEC 27001, the global standard for Information Security Management Systems (ISMS), provides NetApp customers, partners, vendors, and employees with the confidence and assurance that information in the hands of NetApp is protected information.

NetApp security processes covered by the ISO 27001 certification include:

- 3rd Party Audit
- 3rd Party Info Handling & Sharing
- Antivirus
- Application Security
- Business Continuity
- Content Monitoring
- Firewall Management
- Forensic Investigations
- Identity Management
- Incident Response Mgmt
- Intrusion Prevention
- Monitoring
- Security Architecture and Design
- Security Audit and Assessment
- Awareness, Education, & Training
- Security Management
- Security Operations
- Security Policy Management

- Contracts
- Extranet Management
- Public Key Infrastructure
- Risk Management
- Threat & Vulnerability Management

Information assurance in the cloud should not be a shell game in which consumers are left to guess what measures are in place to protect their information assets. Trusted solutions are built on a foundation of trusted and certified products that embed security best practices throughout the design and implementation lifecycle. Cloud solutions are not secure solutions unless they capably demonstrate:

- Reinforcing security controls at the network, compute, and storage layers of the architecture
- Processes and technologies that protect, detect, correlate, and react
- The ability to mitigate risk consistent with business needs and risk threshold
- The ability to achieve compliance with laws, regulations, and standards

In “Cloud Computing: Business Benefits with Security, Governance, and Assurance Perspectives,” the Information Systems Audit and Control Association (ISACA) states that “Cloud computing represents a rare opportunity to rework security and IT controls for a better tomorrow.”²⁸

At NetApp, we’re not waiting for tomorrow. Together with our partners we’re delivering the people, processes, and technologies that help to assure a better “today.”



7 REFERENCES

NETAPP SECURITY GUIDELINES AND BEST PRACTICES

- TR-3868: “PCI-DSS and Data ONTAP Checklist”
<http://media.netapp.com/documents/PCI-DSS+Checklist+for+Data+ONTAP.pdf>
- TR-3834: “Security Guidelines for Data ONTAP 8.0 7-Mode”
<http://media.netapp.com/documents/tr-3834.pdf>
- TR-3649: “Best Practices for Secure Configuration of Data ONTAP 7G”
<http://media.netapp.com/documents/tr-3649.pdf>
- TR-3107: “Antivirus Scanning Best Practices Guide”
<http://media.netapp.com/documents/tr-3107.pdf>
- TR-3580: “NFSv4 Enhancements and Best Practices Guide—Data ONTAP Implementation”
<http://media.netapp.com/documents/tr-3580.pdf>
- TR-3367: “NetApp Storage Systems in a Microsoft Windows Environment”
<http://media.netapp.com/documents/tr-3367.pdf>
- TR-3749: “NetApp and VMware vSphere Storage Best Practices”
<http://media.netapp.com/documents/tr-3749.pdf>
- TR-3458: “Unified Windows and UNIX Authorization Using Microsoft Active Directory LDAP as a Directory Store”
<http://media.netapp.com/documents/tr-3458.pdf>
- TR-3481: “Kerberos NFS in a NetApp Storage System Using a UNIX-Based Kerberos Authentication Server”
<http://media.netapp.com/documents/tr-3481.pdf>



8 ACKNOWLEDGEMENTS

This report was developed in concert with the NetApp Field Centers for Innovation.

ENDNOTES

¹ Frank Gens, et al., “Cloud Computing 2010 – An IDC Update,” 29 September 2009 [IDC Executive Telebriefing], 11; available from <http://www.slideshare.net/JorFigOr/cloud-computing-2010-an-idc-update>; Internet; accessed 15 July 2010.

² European Network and Information Security Agency (ENISA), “An SME Perspective on Cloud Computing,” 20 November 2009 [ENISA survey], 15; available from http://www.enisa.europa.eu/act/rm/files/deliverables/cloud-computing-sme-survey/at_download/fullReport; Internet; accessed 20 July 2010.

³ Thomas Ptacek and Eric Monti, “NetApp MultiStore: An Independent Security Analysis,” 18 June 2008, 12; available from <http://www.silicon.com/white-papers/view/management/netapp-multistore-an-independent-security-analysis-60915621/> [registration required for download]; Internet; accessed 11 June 2010.

⁴ “NIST Definition of Cloud Computing,” v15, 7 October 2009; available from <http://csrc.nist.gov/groups/SNS/cloud-computing/cloud-def-v15.doc>; Internet; accessed 2 July 2010.

⁵ Cloud Security Alliance (CSA), “Security Guidance for Critical Areas of Focus in Cloud Computing,” version 2.1, December 2009, 18; available from <http://www.cloudsecurityalliance.org/csaguide.pdf>; Internet; accessed 2 July 2010.

⁶ Cloud Security Alliance, “Top Threats to Cloud Computing,” version 1.0, March 2010, 6–14; available from <http://www.cloudsecurityalliance.org/topthreats/csathreats.v1.0.pdf>; Internet; accessed 26 August 2010.

⁷ Cloud Security Alliance, “Top Threats to Cloud Computing,” version 1.0, March 2010, 6–14; available from <http://www.cloudsecurityalliance.org/topthreats/csathreats.v1.0.pdf>; Internet; accessed 26 August 2010.

⁸ Jaikumar Vijayan, “The 2009 Data Breach Hall of Shame,” *Computerworld*, 17 December 2009 [online magazine], 1; Rpt in *Network World* [online magazine]; available from <http://www.networkworld.com/news/2009/121709-the-2009-data-breach-hall.html>; Internet; accessed 10 August 2010.

⁹ Ellen Messmer, “Data Breach Costs Top \$200 per Customer Record,” *Network World*, 25 January 2010 [online magazine]; available from <http://www.networkworld.com/news/2010/012510-data-breach-costs.html>; Internet; accessed 10 August 2010.

¹⁰ “Midyear Breach Report,” *Information Security*, July/August 2010 [online magazine]; available from http://viewer.media.bitpipe.com/1152629439_931/1279750495_63/0710_ISM_updated_072010.pdf, 15; Internet; accessed 12 August 2010.

¹¹ William J. Lynn III, “Defending a New Domain - The Pentagon's Cyberstrategy,” *Foreign Affairs*, September/October 2010 [online magazine]; available from <http://www.foreignaffairs.com/articles/66552/william-j-lynn-iii/defending-a-new-domain?>; 1; Internet; accessed 17 September 2010

¹² Ibid.

¹³ Vijayan, *Computerworld*, 1; available from <http://www.networkworld.com/news/2009/121709-the-2009-data-breach-hall.html>; Internet; accessed 10 August 2010.

¹⁴ Vijayan, *Computerworld*, 2; available from <http://www.networkworld.com/news/2009/121709-the-2009-data-breach-hall.html>; Internet; accessed 10 August 2010.

¹⁵ Committee on National Security Systems (CNSS) Instruction No. 4009, "Guideline for Identifying an Information System as a National Security System," 26 April 2010, 35; available from http://www.cnss.gov/Assets/pdf/cnssi_4009.pdf; Internet; accessed 27 September 2010.

¹⁶ NIST, Special Publication 800-53 Rev 3, "Recommended Security Controls for Federal Information Systems and Organizations," August 2009, B-11; available from http://csrc.nist.gov/publications/nistpubs/800-53-Rev3/sp800-53-rev3-final_updated-errata_05-01-2010.pdf; Internet; accessed 27 September 2010.

¹⁷ National Air Traffic Controllers Association (NATCA), "Air Traffic Control: By the Numbers," available from <http://www.natca.org/mediacenter/bythenumbers.msp>; Internet; accessed 24 August 2010.

¹⁸ Wade Baker et al., "2010 Data Breach Investigations Report," 21 July 2010, 2-3; available from http://www.verizonbusiness.com/resources/reports/rp_2010-data-breach-report_en_xg.pdf; Internet; accessed 28 July 2010.

¹⁹ NIST, Special Publication 800-53 Rev 3, "Recommended Security Controls for Federal Information Systems and Organizations," August 2009, 27; available from http://csrc.nist.gov/publications/nistpubs/800-53-Rev3/sp800-53-rev3-final_updated-errata_05-01-2010.pdf; Internet; accessed 4 May 2010.

²⁰ Ptacek and Monti, 12.

²¹ PCI Security Standards Council, "About the PCI Data Security Standard (PCI DSS)," available from https://www.pcisecuritystandards.org/security_standards/pci_dss.shtml; Internet; accessed 7 September 2010.

²² *Health Insurance Portability and Accountability Act of 1996 (HIPAA), Technical Safeguards*, sec. 164.312 (2002) 737; Internet; available from http://edocket.access.gpo.gov/cfr_2004/octqtr/pdf/45cfr164.312.pdf; Internet; accessed 7 September 2010.

²³ Mayer-Brown, "The Gramm-Leach-Bliley Act Executive Summary," 17 December 1999, 12; available from <http://www.securitization.net/pdf/ExecSummary.PDF>; Internet; accessed 7 September 2010.

²⁴ The Institute of Internal Auditors, "Sarbanes-Oxley Section 404: A Guide for Management by Internal Controls Practitioners," 2nd Ed., January 2008, 11; available from <http://www.theiia.org/download.cfm?file=31866>; Internet; accessed 7 September 2010.

²⁵ "SAS 70 Overview," 2010, linked from the SAS70.com homepage at "About SAS 70," available from http://sas70.com/sas70_overview.html; Internet; accessed 7 September 2010.

²⁶ NIST, Federal Information Processing Standards Publication (FIPS) 199, "Standards for Security Categorization of Federal Information and Information Systems," February 2004, 6; available from <http://csrc.nist.gov/publications/fips/fips199/FIPS-PUB-199-final.pdf>; Internet; accessed 3 September 2010.

²⁷ National Security Agency, "Best Practices for Storage Networks," 18 October 2007; available from http://www.nsa.gov/ia/_files/vtechrep/l732-012R-2007.pdf; Internet; accessed 10 April 2010.

²⁸ Information Systems Audit and Control Association (ISACA), 29 October 2009, 9; available from <http://www.isaca.org/Knowledge-Center/Research/ResearchDeliverables/Pages/Cloud-Computing-Business-Benefits-With-Security-Governance-and-Assurance-Perspective.aspx>; Internet; accessed 26 August 2010.

NetApp provides no representations or warranties regarding the accuracy, reliability or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

© Copyright 2010 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, Data ONTAP, FlexShare, MultiStore, NetApp DataFort, RAID-DP, SnapLock, SnapMirror, SnapRestore, Snapshot, vFiler, and WAFL are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Cisco, Unified Computing System, and UCS are registered trademarks of Cisco, Inc. VMware is a registered trademark and vCenter and vSphere are trademarks of VMware, Inc. Windows, Microsoft, and Active Directory are registered trademarks of Microsoft Corporation. UNIX is a registered trademark of The Open Group. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. RA-0006



www.netapp.com