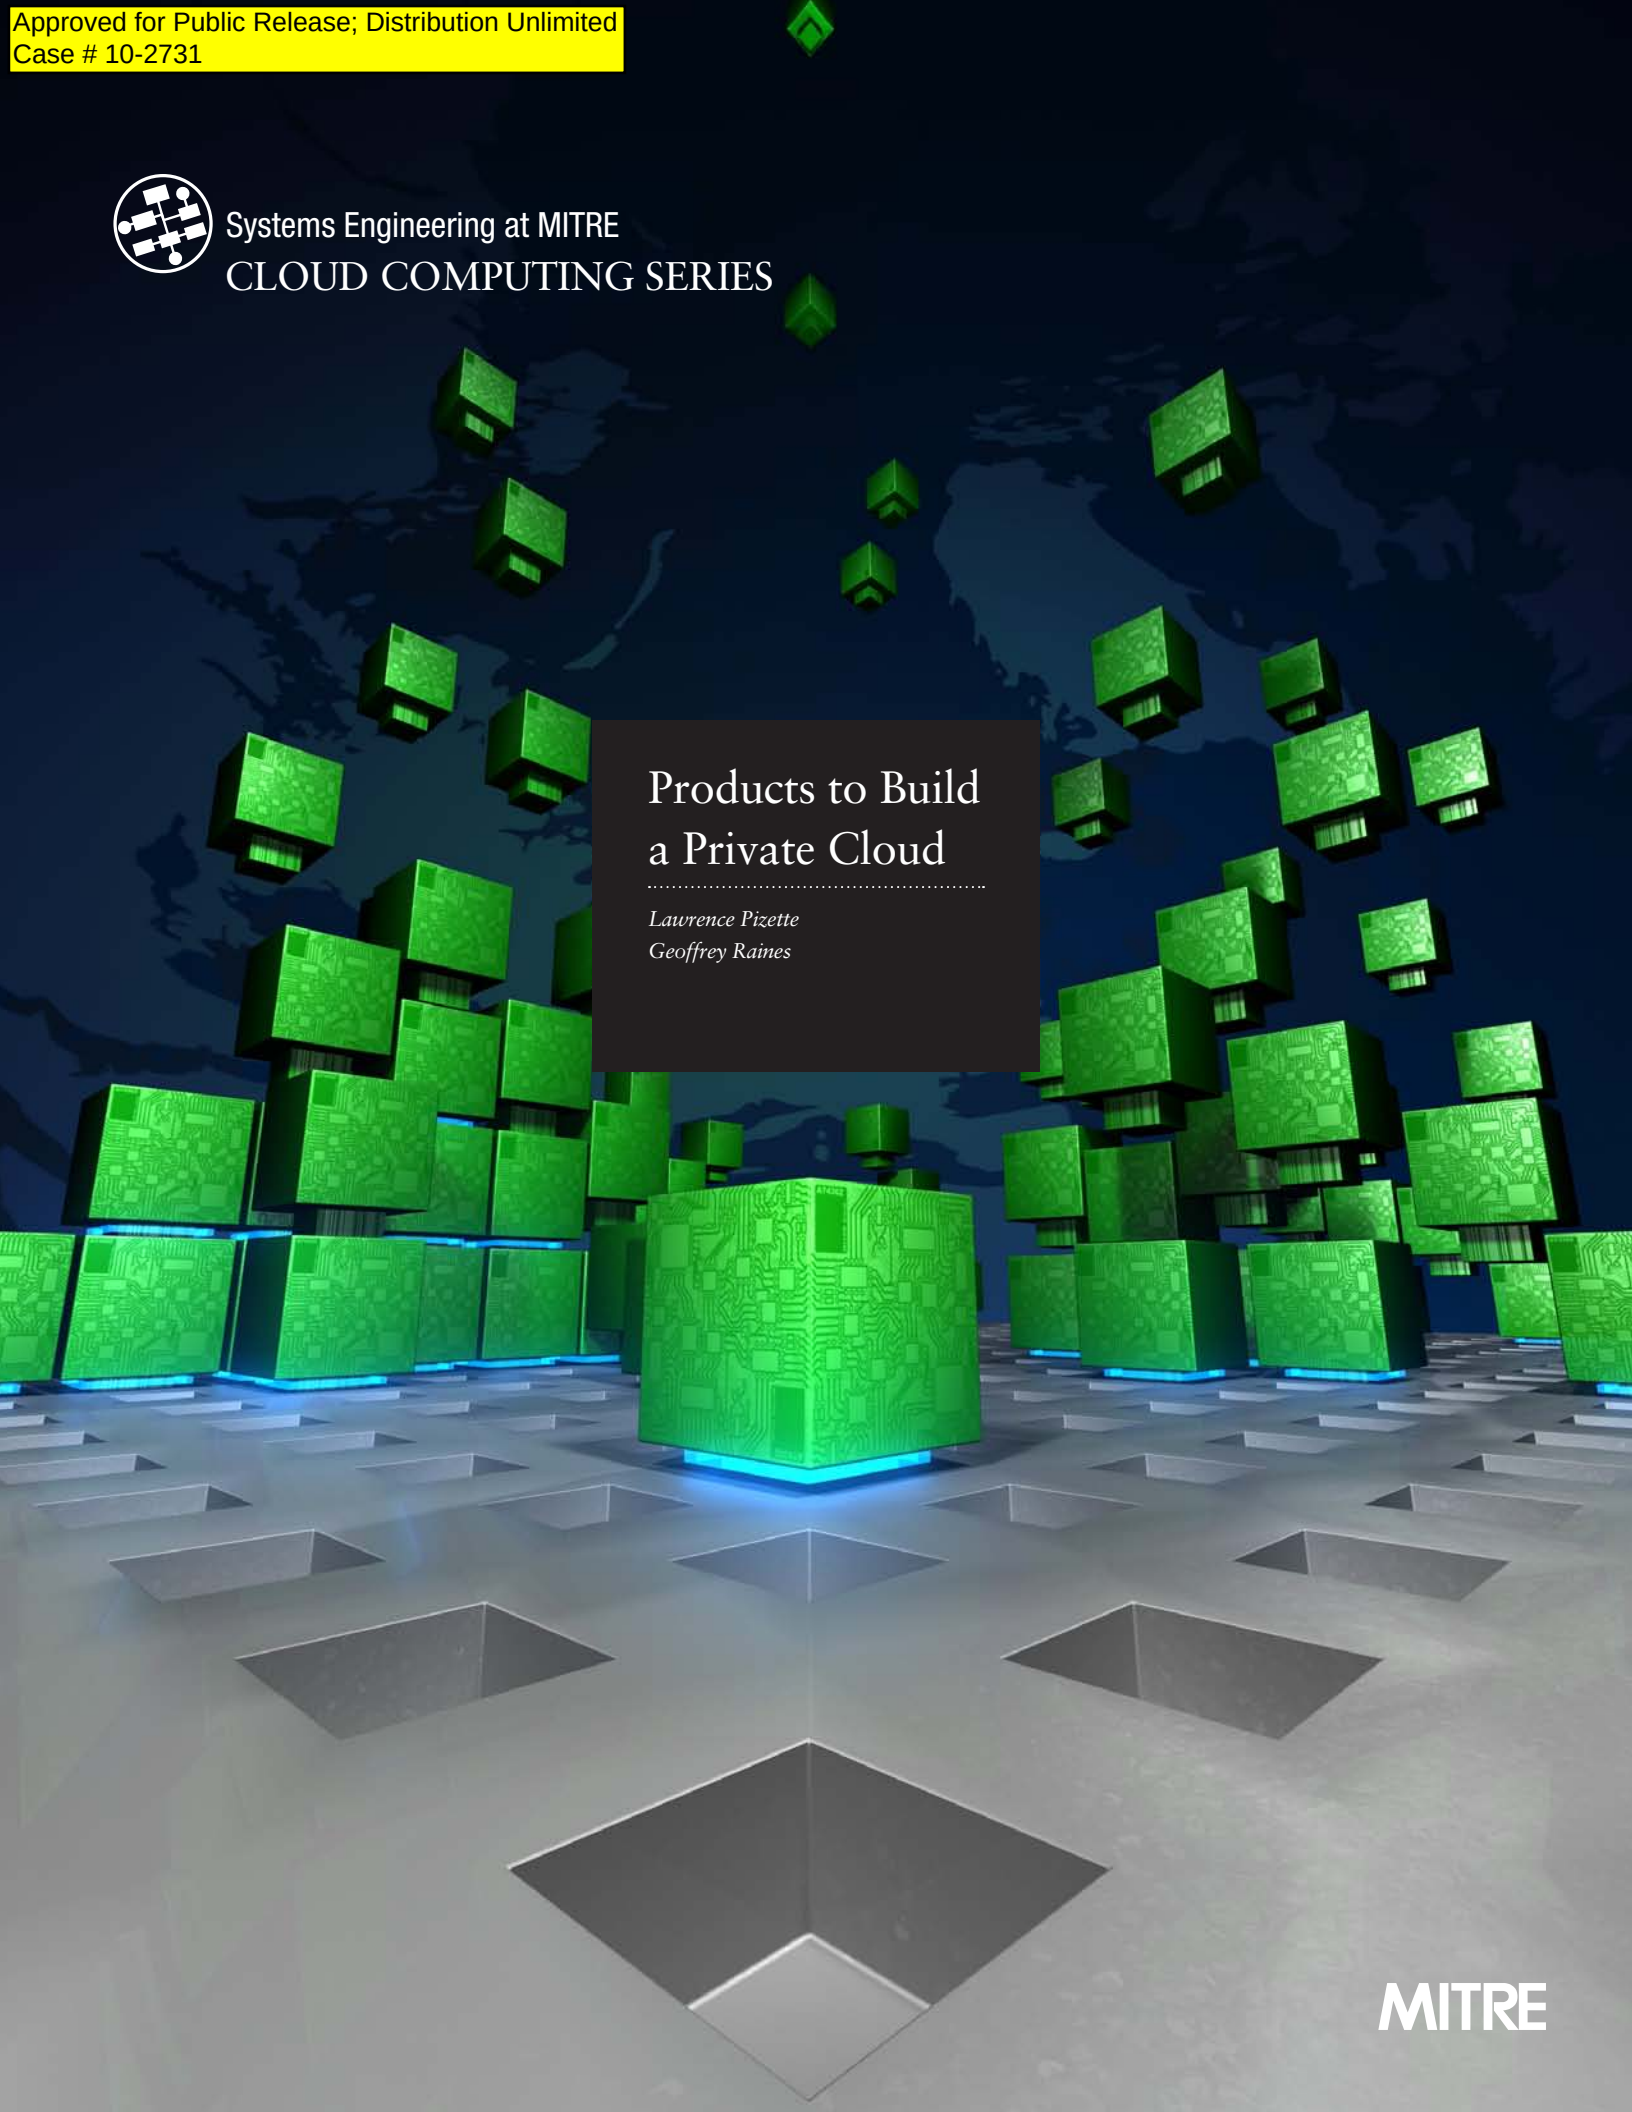

Systems Engineering at MITRE
CLOUD COMPUTING SERIES

Products to Build a Private Cloud

Lawrence Pizette
Geoffrey Raines

MITRE

July 2010

Executive Summary

Federal information technology (IT) leaders who are seeking to leverage cloud computing concepts while maintaining full control and ownership of their IT capability may elect to implement a private cloud computing environment. A private cloud capability can offer increased security and provide the flexibility to implement unique requirements along with the benefits and attributes of ownership, including data rights and control over the system. Some of the benefits of using a private cloud approach over traditional IT are: it allows for better utilization of hardware, the ability to easily provide continuity of operations (COOP), the ability to quickly provision capabilities, and the ability to provide location independent access. Organizations that pursue development of a private cloud infrastructure are investing their own capital with commensurate development schedules, risks, and operating costs. They are trading off the attributes of a private cloud for a different set of benefits and risks that they might attain if they chose a different cloud computing deployment model. In order to help these leaders to understand the private cloud computing products available in early 2010 and how they may fit into their data centers, this paper surveys the private cloud computing market.

While the benefits of a private cloud may be numerous, shared community or public cloud offerings may be available sooner at lower costs, but will likely offer a different set of attributes for mission assurance and other important characteristics.

Attributes of Cloud—While there is no single agreed upon definition for cloud computing, the National Institute for Standards and Technology (NIST) has listed several *essential characteristics* of a private cloud approach.[1] The essential characteristics that they list are on-demand self-service, broad network access, resource pooling (e.g., multi-tenancy), rapid elasticity (e.g., rapid scaling), and measured service. As aspects of these characteristics are present in other IT approaches, it is the combination that describes cloud computing.

> *"With today's virtual infrastructure solutions and a growing list of internal cloud platform technologies, it's fairly easy for an enterprise to start up a cloud-like environment within its own domain."*
>
> *– James Staten et al., Deliver Cloud Benefits Inside Your Walls, Forrester Research, April 13, 2009*

Categories of Infrastructure for a Private Cloud—The product categories listed below were selected based upon a generalization of the private cloud computing marketplace and available technologies. However, industry innovation is not constrained to neat categories; some vendors are offering suites of capabilities that span multiple categories.

- **Virtualization** technology provides the hypervisor software, which allows multiple instances of "guest" operating systems to run concurrently on the same physical server. This capability enables "multi-tenancy," which is the sharing of physical resources providing for better server utilization. The improved server utilization can reduce heating, ventilation, and air conditioning (HVAC) and electric costs, data center, size and other related infrastructure costs. It also can allow advanced users or data center operators to provision capabilities quickly, facilitating scalability, self-service provisioning, and COOP.
- **Storage** technology, such as disk arrays, storage area networks (SANs), and storage connection

technologies can provide the underlying persistent storage for a private cloud. The market for these types of systems is huge: In Q4 2009, the worldwide external disk systems market was $5.288 billion with the top five vendors, EMC, IBM, HP, NetApp, and Dell taking 69 percent of the market share.[2] Regardless of which vendor provides it, storage is essential to a private cloud as it can enable burst capacity and provide capabilities to facilitate COOP and location independent access.

- **Security** capabilities for a private cloud are essential. Products offered to the private cloud computing market provide capabilities such as identity management, logging and auditing, anti-malware software, intrusion detection systems (IDS) and intrusion prevention systems (IPS), and firewall capabilities between virtual machines. While there are many products on the market, addressing security concerns within multi-tenant, cloud computing environments is an active area for research for many organizations. Therefore, products will continue to evolve as more knowledge is learned about threats and measures to mitigate them.
- **Provisioning, management, and metering** tools are key to providing Federal IT leaders and users with many of the advantages a private cloud can offer: self-service provisioning, burst capability, service-level agreement (SLA) compliance, and, if needed, metering for "pay as you go" functionality. Tools in this category provide the ability to monitor server utilization, system resources, and other operational attributes. The ability to provision environments on-demand, independent of the underlying hardware enables the ability to scale environments in near real time and, if necessary, to instantiate environments in a different location.

Considerations—

- When developing a private cloud, a comprehensive systems engineering process should be employed that considers the Federal IT leader's high-level objectives and specific requirements. While some private clouds may be driven by the desire to reduce costs, others may be driven primarily by new capabilities that can be provided. Additionally, there can be many specific requirements in a private cloud such as performance, self-service provisioning, COOP, location independent access, and robust security, which need to be considered carefully before acquiring specific hardware and software products.
- Security should be considered from the beginning of the structured systems engineering process. Security architects should perform an analysis and consider a pilot to understand new threats and mitigating strategies. The analysis should consider how private cloud capabilities would be integrated into the rest of the data center's environment (e.g., identity management systems).
- A private cloud will require updated management and governance. For example, if virtual environments are not managed, the absence of controls coupled with the relative ease of creating virtual machines can result in a proliferation of environments and a condition informally known as "VM sprawl." This condition can offset anticipated efficiency gains and allow for security exposures due to unmaintained environments. In order to mitigate the risk, piloting and testing of the integrated environment should be considered as part of the overall systems engineering plan.

A private cloud implementation can offer significant benefit to Federal organizations seeking to realize some of the benefits of cloud computing while maximizing control over their environment. Recognizing that the benefits coupled with control can be compelling for many Government organizations, industry has provided a wide selection of products for realizing the vision. This paper provides insight into example products and considerations for their use.

Table of Contents

1.0 Introduction to Products to Build a Private Cloud	1
2.0 Virtualization and Hypervisors	2
3.0 Storage Solutions	4
4.0 Security Products	6
5.0 Provisioning, Management, and Metering	8
Appendix A—Products Mentioned	11
References	12

THE BIG PICTURE: A private cloud computing environment can offer a path forward for Federal IT leaders seeking to realize some of the benefits of cloud computing while maximizing control over their environment.

Products to Build a Private Cloud

Lawrence Pizette
Geoffrey Raines

1.0 Introduction to Private Clouds

***Summary*—**For Federal information technology (IT) leaders considering employing a private cloud approach, this paper provides a window into products that are available from the commercial marketplace. The paper is organized into the general categories of *virtualization and hypervisors, storage* (e.g., storage area networks, storage arrays), *security,* and *provisioning, management, and metering.* The product categories in this paper were selected based upon a generalization of the private cloud computing marketplace and the technologies available in early 2010. However, vendor products do not always fall directly within one category; some vendors are offering suites of capabilities that span multiple categories and other vendors are offering appliances, some of which package both hardware and software into "a cloud in a box." Similarly, other vendors are offering capabilities for hybrid clouds, which can blur the line between private cloud capabilities and capabilities that leverage public cloud offerings over the network.

Due to space constraints in the paper, we felt it was best to describe exemplars in each general category. While the marketplace may have many products in a category, we looked for offerings that would best represent the concepts. *Our selection of one vendor over another is not an endorsement of the vendor or its products.* Also, while the general considerations described in this paper are expected to endure for some time, cloud computing represents a fast moving market so that specific products will likely be replaced rather quickly.

***Cloud Deployment Model Definitions*—**The National Institute of Standards and Technology (NIST) describes four deployment models of cloud computing: private, community, public, and hybrid.[3] According to the NIST definition:

> *"Companies that want the benefits of cloud computing services without the risks are looking to create cloud-like environments in their own data centers. To do it, they'll need to add a layer of new technologies—virtualization management, cloud APIs, self-service portals, chargeback systems, and more—to existing data center systems and processes."*
>
> *– Charles Babcock, Information Week, April 11, 2009*

- **Private clouds** are "operated solely for an organization."
- **Community cloud** "infrastructure is shared by several organizations and supports a specific community that has shared concerns."
- **Public 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** "infrastructure is a composition of two or more clouds."

In this paper, hybrid cloud will refer the ability to leverage a private cloud capability along with a community or public cloud offering.

***Reasons for Building a Private Cloud*—**A private cloud capability can offer increased security over other deployment models, while still providing the flexibility to implement unique capabilities with the benefits of ownership (e.g., data rights, control of the system). Forrester Research states, "An internal cloud [is] a multitenant, dynamically provisioned

and optimized infrastructure with self-service developer deployment, hosted within the safe confines of your own data center."[4]

The benefits of using a private cloud approach (over traditional IT) may involve better utilization of hardware, the ability to easily provide continuity of operations (COOP), the ability to quickly provision capabilities, and the ability to provide location independent access or fast access to Software as a Service (SaaS) capabilities within a private organization. While these benefits of a private cloud may be numerous, organizations that pursue development of a private cloud infrastructure are investing their own capital with commensurate development schedules, risks, and operating costs.

Categories of Infrastructure for a Private Cloud—

- **Virtualization** technology provides the hypervisor software, which allows multiple instances of "guest" operating systems to run concurrently on the same physical server. This capability enables "multi-tenancy," which is the sharing of physical resources providing for better server utilization. The improved server utilization can reduce heating, ventilation and air condition (HVAC) and electric costs, data center size and other related infrastructure costs. It also can allow advanced users or data center operators to provision capabilities quickly, facilitating scalability, self-service provisioning, and COOP.
- **Storage** technology, such as storage area networks (SANs), can provide the underlying persistent storage for the data center. It can provide capabilities to scale usage to petabyte levels of storage, while facilitating COOP and location independent access of applications.
- **Security** products offered to the private cloud computing market provide capabilities such as identity management and virtual firewalls. They will continue to evolve as private cloud products mature and more knowledge is learned about threats.
- **Provisioning, management, and metering** tools provide the ability to deploy virtual operating systems, services, and applications. They may provide the ability to provision Platform as a Service (PaaS) environments with enhanced features and have self-service portals. Tools in this category provide the ability to monitor server utilization and system resources, application service-level agreements (SLAs), and other operational attributes. Some products also provide usage information for billing and accounting purposes.

Considerations—

- Before initiating a private cloud investment, Federal IT leaders should examine the attributes of a private cloud in comparison to other deployment options. The private cloud can offer complete control, but it may not provide the same economies of scale as community or public clouds.
- The goals of Federal leadership should significantly impact the technology investment. While some private clouds may be driven by the desire to reduce costs, others may be driven primarily by its other tangible benefits, such as scalability, location independent access, and the ability to support COOP. For more details on analyzing costs, MITRE will be releasing a Cloud Computing Cost and Business Case Considerations white paper in Q4 2010.
- When developing a private cloud, a comprehensive systems engineering process should be employed; there can be many specific requirements in a private cloud such as performance, self-service provisioning, and robust security, which need to be considered carefully before acquiring specific hardware and software products.
- As Federal organizations can have a large portfolio of legacy systems, IT leaders should develop a migration plan that identifies the capabilities that will be moving to the cloud, and testing and integration activities.

2.0 Virtualization and Hypervisors

Summary—Hypervisor software provides a virtual operating system (OS) environment, enabling multiple virtual OSs to be present on one physical server—the hypervisor is considered by many to be the foundational layer of a cloud computing environment and a critical component of a private cloud. The hypervisor software can reduce IT capital expenditures and operating expenses through higher utilization of processing resources; the higher utilization of servers allows for a decrease in the total number of servers, providing for lower overall data center costs through reductions in electric usage, HVAC costs, and data center floor space. Additionally, the hypervisor can facilitate

"Federal Government organizations increasingly are incorporating virtualization to consolidate IT resources, reduce space and power consumption and gain greater agility in launching new applications and services for employees and constituents."

– Barbara DePompa, Benefits of Virtualization for Government, Federal Computer Week

multi-tenant usage and the ability to provision environments rapidly.

Server virtualization products are available from multiple sources, including commercial-off-the-shelf (COTS) offerings and open source software products such as Xen and Kernel-based Virtual Machine (KVM). While there are many options, as shown in Figure 2-1, VMware and Microsoft have the greatest penetration in the fast-moving x86 server virtualization market as of Q3 2008. (IBM has a non-x86 virtualization offering for its System Z platform, which would not appear on this list.)

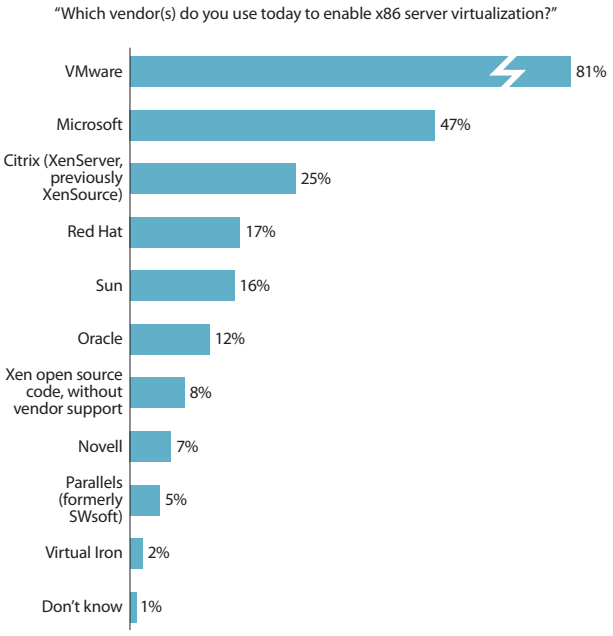

Base: 106 North American and European x86 server virtualization decision-makers at companies with 100+ employees that have already implemented x86 server virtualization or are implementing it in the next 12 months (multiple responses accepted)

Source: Enterprise and SMB Hardware Survey, North America and Europe, Q3 2008

Figure 2-1. Virtualization Market

http://www.forrester.com/rb/Research/server_trends_hypervisor_wars_heat_up/q/id/48165/t/2

There are two types of hypervisors on the market: A Type 1 "bare-metal" hypervisor runs directly on the hardware without the overhead of a host OS. A Type 2 "hosted" hypervisor runs on a host OS; all system resource requests go through the host OS.[5] As a result, the Type 1 hypervisor technology can provide better performance than the Type 2 hypervisor, but a Type 2 hypervisor "supports the broadest range of hardware configurations."[6]

Figure 2-2 shows an example Type 1 hypervisor technology stack. Examples of this type of hypervisor on the market are VMware ESX and ESXi, XenServer from Citrix, and Hyper-V from Microsoft.[7] In addition to the COTS hypervisors, there are two open source hypervisors, Xen and KVM, that are Type 1. Xen is the underlying foundation for Citrix Systems XenServer and Oracle VM.[8] KVM has been adopted as the foundation for Red Hat's Virtualization Hypervisor, which can be "deployed either as the standalone bare metal hypervisor, or as Red Hat Enterprise Linux 5.4 and later installed as a hypervisor host technology."[9]

Figure 2-3 shows a Type 2 hypervisor technology stack, which requires a host operating system to run.

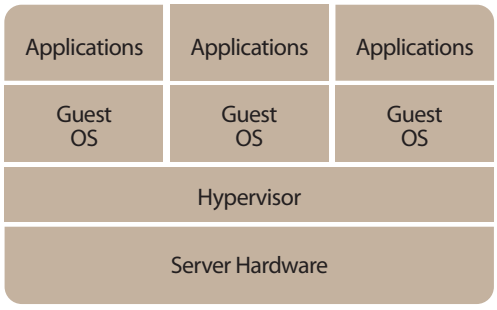

Figure 2-2. Type 1 Hypervisor

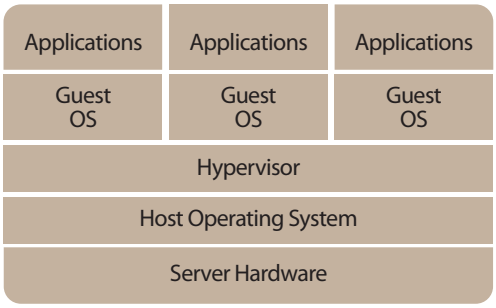

Figure 2-3. Type 2 Hypervisor

VMware Server is an example of a hosted virtualization technology. All application requests for resources need to go through both the hypervisor and the operating system to get to the hardware.

Hardware and OSs Supported—When examining the hypervisor market, there are fundamental features that must be understood, such as which processors are compatible with the hypervisor and which guest OSs are supported. For example, Intel and AMD both provide special processor features, VT and AMD-V, respectively, known as hardware-assisted virtualization, which are required by some hypervisors. As of March 2010, KVM requires Intel-based processors with virtualization technology or AMD-based processors with virtualization technology.[10] Similarly, Microsoft requires "a 64-bit environment" for the Hyper-V platform and processors which "support hardware-assisted virtualization (Intel VT or AMD-V) technology."[11]

Considerations—

- While there is a significant benefit to increasing the utilization of servers through virtualization, the hypervisor adds new software and potential complexity to the IT stack. IT organizations will need to plan for the incorporation of this new technology.
- Virtual environments require management and governance. If they are not managed, the absence of controls coupled with the relative ease of creating virtual machines can result in a proliferation of environments and a condition informally known "VM sprawl," which could offset anticipated efficiency gains.
- Virtualization products can increase software licensing costs and the complexity of maintaining software license compliance, which should be factored into the cloud computing cost-benefit analysis.
- Virtual machines need to be allocated to physical servers in an optimal manner for the data center and monitored and configured for appropriate performance. The exact number of virtual machines per server depends on the resources used in the virtual machine and other factors such as networking requirements. If the servers are under-utilized, cost savings will not accrue. If the servers are over-utilized, SLAs may not be met and performance will suffer. Getting the number of virtual machines and configuration right, based upon an analysis of usage and performance factors (including processor, memory, disk input/output, network input/output) and application performance is essential. "Some virtualization implementations fail to produce the expected savings because the cost per virtual machine (VM) is too high. This usually results from poor VM density and high hardware costs in shops that are too conservative with their deployments," states Forrester Research's Galen Schreck. As an example, he notes, "Many companies break even with physical hardware after three VMs, but they only put a total of five onto a server capable of running 15 VMs. They could have provisioned 10 more VMs for no incremental cost."[12]
- An analysis of server environments and testing will need to be planned as part of the process of migrating legacy applications to virtualized environments. Not all legacy applications can seamlessly be migrated from traditional servers to virtualized environments. Not all operating systems are supported by all hypervisors, and some applications may require dedicated physical resources. Applications that require extensive use of a graphics processing unit and systems with real-time requirements often are poor choices for a virtualized environment.

3.0 Storage Solutions

Summary—A persistent storage solution is an essential capability for a private cloud environment. Commonly used storage solutions, incorporating technologies such as disk arrays and SANs, can provide the underlying persistent data storage for high-end, private cloud offerings. Disk arrays typically are hardware units with multiple disk drives and controller units that allow for individual disk failures without loss of data (see Figure 3-1). SANs provide a network of storage resources that can be allocated to multiple servers. With a SAN, individual servers are not directly attached to individual disks. While disk arrays, SANs, and traditional data center storage solutions may be satisfactory for a private cloud implementation, cloud-focused storage features are starting to emerge in the marketplace that add features specifically targeted for the cloud architectures. As noted on SearchStorage, "Over the past year or so, vendors have come out with new products or re-worked old ones by adding multi-tenancy and security features that make the products suited to power private storage clouds in the enterprise."[13]

"Theoretically, almost any storage system and software can be used to build a private storage cloud. But over the past year or so, vendors have come out with new products or re-worked old ones by adding multitenancy and security features that make the products suited to power private storage clouds in the enterprise."

– Private Storage Cloud Software And Hardware Product Guide, February 15, 2010, SearchStorage.com

Important factors to understand with storage solutions are how the features support the overall requirements of the private cloud, such as performance, scalability, location independence, and COOP. Existing data center storage technologies may be used in today's private cloud implementations. However, in the future, more products will likely emerge in the market place with new features that are unique and compelling for private clouds.

The private cloud architect will have to choose among a number of storage connection technologies. When combining solutions, it is important to recognize that some storage solutions utilize industry standard protocols, facilitating a heterogeneous data center, while others are limited to propriety interfaces; Fibre Channel, hardware and software iSCSI and NFS are common standards-based protocols for moving data to and from storage devices. In a 2008 study, VMware compared these protocols for performance using their ESX Server 3.5 and the IOmeter for benchmark comparison.[14] They found that the Fibre Channel was superior in throughput, latency, and CPU utilization, but that did not exclude the other protocols that could meet requirements in certain circumstances with lower costs.

Storage Market—When Federal IT leaders look to acquire storage technology—which can be a significant cost driver for a new private cloud investment—they find a handful of major vendors that command a large portion of the market. In a 2006 survey, Forrester Research found that storage spending was almost 19 percent of overall hardware IT spending in North America.[15] In a more recent 2010 report, IDC indicates that the worldwide external disk storage system factory revenue in Q4 2009 was $5.288 billion.[16] The top five vendors, EMC, IBM, HP, NetApp, and Dell commanded 69 percent of the market; however, the remaining companies had a 31 percent market share that is worth $1.643 billion.

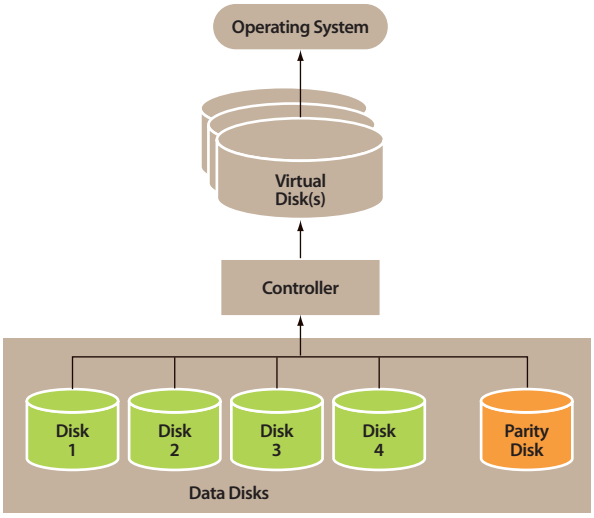

Figure 3-1. Example RAID Disk Array

http://docs.hp.com/en/B7961-90019/ch03s02.html

Scalability—As Federal IT leadership plans their private cloud investments, scalability in the underlying storage infrastructure should be considered in their technology selection. The ability for users to leverage cloud infrastructure in an "on demand" environment with burst capability is an important attribute of cloud computing; therefore, scalability in the underlying storage is essential. For example, EMC Atmos provides the ability to scale to the Petabyte level.[17] Similarly, IBM Scale Out Network Attached Storage (SONAS) offers "file-based storage, supporting over 14 PBs."[18]

Continuity of Operations and Location Independent Access—For the Federal Government, the ability to provide COOP when adverse events happen (whether the result of a natural occurrence or an adversary) is essential. Being able to access and reconstitute data within a specified time limit is an imperative for many systems. Frequently, the underlying storage system participates in the delivery of this capability. As an example of a capability that a storage product can provide, Atmos provides a GeoProtect feature. This GeoProtect feature allows data to be replicated in multiple geographic locations according to policy and the ability to recover data distributed across geographies.[19] Other storage vendors also have solutions for COOP and location independent access. For example, HP offers the "StorageWorks B-series Remote Replication Solution."[20] Similarly, IBM's SONAS provides

capabilities needed for many private clouds. They state, "In a cloud environment, applications and services are not tethered to specific hardware components. Instead, processing is handled across a distributed, globally accessible network of resources, which are dispensed on demand." [21]

Considerations—

- Given the cost of incorporating new technology into a data center, an analysis of whether existing data center storage technologies can be repurposed for a private cloud should be considered.
- System architects should weigh costs against the need for performance when selecting connection technologies. While Fibre Channel provides much better performance than iSCSI and NFS, the additional performance may not be needed for all systems.
- When investing in storage technologies, organizations should be aware of whether the products integrate with their existing infrastructure through open standards (e.g., Fibre Channel, iSCSI, NFS) or whether they have proprietary, closed interfaces.

4.0 Security Products

Summary—Federal IT leadership should acquire and employ security technology within the context of an overall security plan, system requirements, and organizational needs. Given the variation in security requirements and approaches for private clouds, the products listed in this section should be considered as examples of security products on the market rather than a comprehensive list of ingredients for securing a private cloud. Additionally, many modern data centers will have variations of these capabilities for their existing infrastructure. New capabilities may need to integrate seamlessly with existing infrastructure, which can have a significant impact on the choice of technologies and products.

To exemplify the market offerings for securing private clouds, we list products with the following capabilities: identity management; logging and auditing; firewall, intrusion detection systems (IDS) and intrusion prevention systems (IPS); and antivirus software for virtualized environments.

Identity Management—Federal IT leadership should consider the IT capabilities that are required for identity management and access management in a private cloud. While identity management concepts will be similar, there are new products to integrate (e.g., the virtualization infrastructure) and new capabilities that require controlled user access (e.g., the ability to provision virtual environments). Identity management includes authentication to validate that a requester is legitimate and authorization to determine a legitimate requester's access privileges. As this functionality is a "must" for many organizations, most existing data centers already will have an identity management capability that may be able to be extended or integrated with a private cloud implementation. For example, there are products on the market, such as Centrify Suite for Linux/Unix, that provide the capability to leverage Active Directory for identity management and single sign-on to virtualized environments (e.g., Citrix XenServer and VMware ESX Server).[22] There are many variations to identity management products ranging from simple systems to comprehensive capabilities that can provide robust authentication, facilitate operational and policy management, and federate identity capabilities. For example, RSA offers a wide range of security products, including RSA Access Manager.[23] RSA Access Manager supports authentication, authorization, single sign-on, and centralized control across an enterprise. Another example is Oracle's Identity Management and Access Management suite, which can be used for managing access in a PaaS private cloud environment.[24] Oracle's offering supports single sign-on, centralized access control policies, and identity federation. (Identity federation allows more than one system to trust in the authenticity of a user or requester of services.)

> *"[A] cloud computing model changes the way data is accessed, and private clouds must have safeguards in place to protect an organization's most critical and sensitive data."*
>
> *– IBM discussion with Phil Hochmuth, senior analyst at research firm Yankee Group, Public clouds, private clouds and your security*

Logging and Auditing—Logging of events and auditing is important for forensics, helping to identify threats and eliminating access and attack vectors. In a virtualized environment, there is a need to monitor the host environment, virtualized environments, and communication between the virtualized environments. For example, OPNET ACE Live VMon can provide information on the status of

transactions that are "flowing between VMs on the same physical host" and provide a packet-based data set.[25] Logging and auditing also may be required for regulatory and statutory compliance. For example, the Health Insurance Portability and Accountability Act (HIPAA) requires "procedures to regularly review records of information system activity, such as audit logs, access reports, and security incident tracking reports" for protected health information.[26] RSA enVision is an example tool that can provide auditing and event management for both the host environment and a VMware virtualized environment.[27] It also provides event correlation across a broad range of data center inputs to help identify threats. Oracle's offering mentioned above provides reports on authentication and authorization records, which may be needed for demonstrating reporting compliance. Other products can be used to provide audit records through simple data downloads; for example, Citrix provides the XenServer Audit Tool, which can be used to export the XenServer log file.[28]

Firewalls for Virtualized Environments—A virtual firewall is essential for allowing authorized communications and preventing unauthorized or malicious network traffic. In a traditional server, this network traffic is monitored at the boundaries of the server. However, in a virtual environment it is necessary to provide the capability between virtual machines as shown in Figure 4-1. For example, Altor Networks provides a product that can implement firewall policies at the virtual machine level and monitor inter-virtual machine communications.[29] The Altor VF Firewall leverages VMware's VMsafe suite of APIs, which provides third-party access directly into the VMware vSphere virtualized environment. Similarly, IBM Virtual Server Security for VMware can inspect traffic between VMware virtual machines.[30] In addition to other security capabilities, it can inspect the virtual environments for the installation of malicious rootkits, which is malware that can compromise the operating system.

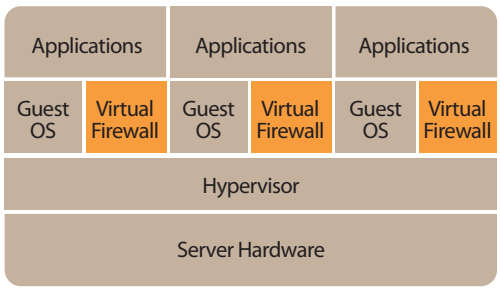

Figure 4-1. Virtual Firewall

Anti-malware Software for Virtualized Environments—Anti-malware software, which can include anti-virus, anti-rootkit, anti-Trojan horse, and anti-spyware protection, is necessary for virtual guest operating systems. This is similar to the need for anti-malware client software in a traditional server environment. There are many anti-malware software products on the market that can be used. For example, Trend Micro offers their ServerProtect product for preventing malware. They also offer a suite of tools to run in VMware ESX and ESXi environments that can leverage the VMsafe APIs.[31] Trend Micro Core Protection for Virtual Machines protects active and dormant virtual machines and provides a layer of separation between the virtual machine being scanned and the scanning agent; the scanning agent is run on a separate virtual machine.

Intrusion Detection and Intrusion Protection—IDS and IPS capabilities are needed to detect and prevent malicious access to data, applications, and systems. IDS and IPS software products can require information from multiple sources, such as operating system events and network activity, to detect and prevent intrusions. IDS systems log these events (e.g., for forensics) and IPS systems attempt to prevent them by automatically modifying configurations. Similar to the malware detection software described above, Trend Micro uses VMware's VMsafe API to provide IDS and IPS capabilities.

Considerations—

- Depending on security needs, security architects should perform an analysis and consider a pilot to understand new threats and mitigating strategies.
- The data center IT security plan, including policies and procedures, should be modified to address and mitigate new threats as they emerge.
- An overall security plan should ensure that virtual environments are appropriately isolated from each other and that communication is monitored between virtual machines, even on the same physical server (e.g., consider a virtual firewall and intrusion detection and prevention software).
- Security architects should consider the additional needs for logging and auditing capabilities within a virtualized environment. For each

physical server, there can be multiple guest operating system environments generating events, which may need to be correlated and examined.
- There may be cost savings and reduced complexity for data centers when acquiring private cloud security products that integrate with their existing data center solutions (e.g., anti-malware software and identity management software).

5.0 Provisioning, Management, and Metering

"The impact of rapid provisioning on ROI business cases can be profound."

– Open Group, Building Return on Investment from Cloud Computing

Overview—For Federal IT leaders looking to harness the benefits of cloud computing in a private environment, provisioning virtual machine environments and applications to the cloud is one of the most important aspects of a private cloud technology. Provisioning an environment is the act of creating a server or virtual server environment with all the requisite software and configurations needed for it to provide the desired capabilities. For example, a server may require an operating system, application software, security software, and multiple configuration settings to function correctly. The mixture of software and configurations can vary significantly, based upon the function of the server. The ability to dynamically provision preconfigured environments, independent of the underlying hardware, facilitates "on-demand" scalability and COOP. This capability can allow IT organizations to stand up virtual servers quickly in geographically disparate locations or scale application resources based upon increased demand. System resource utilization can be monitored through management software, which can also measure compliance with SLAs to ensure the IT infrastructure is meeting operational needs. Additional capabilities may include "metering," which is the gathering of usage date for accounting or billing purposes. As shown in Figure 5-1, these software capabilities provide many of the compelling features for Government organizations to employ private cloud technology.

Provisioning Environments—The market for provisioning in cloud computing is burgeoning with products and a variety of approaches. For example, IBM's approach to provisioning in their CloudBurst hardware-plus-software appliance is to provide an integrated tool set based upon a suite of IBM CloudBurst and Tivoli software products (e.g., Tivoli Service Automation Manager).[32] As another example, 3Tera, recently purchased by CA, is similar to IBM in that they provide an integrated solution with a provisioning and management software suite for use in private clouds; however, they differ significantly in their approach to hardware. Some vendors are choosing to bundle hardware and software together, while others are offering solutions that operate on generic commodity hardware. For example, the IBM CloudBurst appliance is comprised of hardware and software, whereas 3Tera AppLogic comes with a software suite that runs on "commodity" hardware.[33]

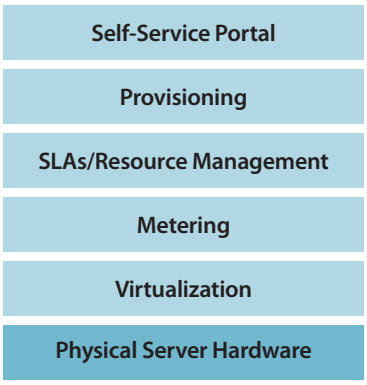

Figure 5-1. Provisioning, Billing, and Self-Service Capabilities

Self-Service Portal—Driven by the need to provide capabilities to users "on-demand" and lower costs, Federal IT leaders should consider employing self-service capabilities in their private cloud implementations. For example, IBM CloudBurst provides a self-service portal capability that provides the ability to provision and scale preconfigured environments (e.g., Linux, Apache, MySQL, applications, and environment variables). In addition, when provisioning a virtual machine environment through the portal, users can allocate system resources to the virtual machine, such as the amount of random access memory (RAM) that it can use. Similar to IBM CloudBurst, AppLogic allows comprehensive application environments to be dynamically provisioned into virtual machines.

Metering—Federal IT leaders may want to look for metering capabilities if they intend to allocate costs or charge projects based upon usage. For example,

both the IBM and AppLogic suites described above provide metering capabilities. Metering is the process of collecting usage information that cloud computing operators can feed into billing systems or other similar downstream accounting systems to allocate costs or charge users. By implementing metering, the data center may be able to offer users, projects, or other organizations a "pay as you go" model. For some private clouds, this capability may not be needed, but others may find it essential due to accounting rules or for funding the cloud environment.

Monitoring and Service-Level Agreements—In sharing and scaling resources, Federal IT leadership will want to ensure that operational needs are being met. Monitoring software can be used to maintain SLAs and ensure users have the system "up time" and responsiveness they need. For example, the AppLogic platform enables a scalable "grid" and allows for the mirroring of data across the commodity hardware, providing for increased availability and backup through the software. AppLogic allows for an N-tier environment across virtual machines to be managed as a single entity through a graphical user interface. IBM also provides monitoring software and management of SLAs in their CloudBurst product.

Hybrid Clouds—Federal IT leaders whose organizations are implementing private clouds may want to consider hybrid cloud technology in the future as the technology matures. While many organizations are pursuing private cloud approaches to maximize security and control over processing and information, there may be opportunities to pursue a hybrid approach in the future to provide burst capacity, ensure SLAs are met, and lower costs. Innovative companies are working on hybrid cloud solutions that will enable IT organizations to extend their private clouds to Government-run community cloud or public cloud offerings. For example, the Eucalyptus open source platform provides an innovative self-service provisioning environment with a different set of attributes and trade-offs from many other private cloud environments. By using the same APIs as Amazon Elastic Compute Cloud (EC2), Eucalyptus can offer its users easy portability—for both applications and tools—between a private cloud and a public cloud. By reusing the EC2 API, Eucalyptus facilitates the ability to move capabilities between a private cloud and a public offering. The enterprise edition with commercial extensions provides integration with storage hardware (iSCSI, NAS, and SANs) and works with multiple hypervisors, including VMware vSphere, VMware ESX, VMware ESXi, KVM, and Xen.[34]

In an effort to provide hybrid capabilities, some vendors are collaborating across platforms. RightScale provides integration with Eucalyptus, Amazon Web Services, and others, enabling users to manage hybrid cloud implementations through a single portal.[35] In addition, there are commercial products on the market that can directly provision a preconfigured environment and applications to either a private environment or Amazon's EC2.

In another example, Cloudswitch has a new product in beta as of March 2010 for transparently hosting VMware-based applications in a private cloud or public cloud offering with a secure networking connection (like a VPN) between the cloud provider and the organization's data center.[36] Products in this space are emerging and provide significant promise for expansion in the future. The expanded capabilities may enable portability, facilitate COOP, provide scalability, and lower costs.

Considerations—

- Federal leaders should plan to pilot and test provisioning and management software to ensure that it meets requirements and integrates with other data center capabilities as anticipated.
- Examine capabilities of products in this domain for their ability to integrate with the layers of the private cloud. For example, some products may require a particular type of hypervisor software for integration or only work with certain security products.
- If workflows need to be monitored in order to meet SLAs, ensure that the management software can meet this need. For example in multi-tier application architectures, a remedy to degraded performance may require a sophisticated response with the ability to provision multiple new instances of virtual servers with different capabilities (e.g., Web servers, application servers).
- Within an enterprise, it can be very efficient to have a self-service portal, but it also opens up new resource management and security challenges. Therefore, data center leaders should consider the security controls and policies necessary for a self-service portal. The need goes beyond malicious activity and includes efficient, prioritized use of data center resources.

If organizations do not control this access, they could deploy many virtual environments, which may be similar or not high priority, resulting in inefficient use of computing resources. These low priority, virtual environments could compete with high value processing needs—a situation that can be avoided through security, procedure, and policy.

- If a private cloud spans multiple geographic locations, Federal leaders should determine policies and procedures for provisioning applications to specific data centers (e.g., based upon performance, COOP requirements, cost efficiency, available capacity).

Conclusion—A private cloud computing approach is a possible path forward for Federal IT leaders looking to leverage cloud computing while maintaining control over their environment. A private cloud approach can be used to harness the processing power of data centers efficiently while delivering new features to users. As part of a comprehensive systems engineering process, existing capabilities can be integrated with innovative new products for features such as COOP, on-demand scalability, and location independent access. These features coupled with appropriate processes, a self-service portal, and comprehensive security, can be effective at achieving the vision of a private cloud.

Appendix A—Products Mentioned

Product	Website	Description
Altor Networks VF Firewall	http://altornetworks.com/products/vf/	Virtual firewall software
AMD-Virtualization	http://sites.amd.com/us/business/it-solutions/virtualization/Pages/amd-v.aspx	Extensions to x86 instruction set to support virtualization
Amazon Elastic Compute Cloud	http://aws.amazon.com/ec2/	Public infrastructure as a service cloud computing
Apache	http://www.apache.org/	Open source software Web server
CA 3Tera AppLogic	http://www.3tera.com/AppLogic/	Distributed, scalable private cloud infrastructure that runs on commodity hardware
Citrix Systems XenServer	http://www.citrix.com/english/ps2/products/product.asp?contentID=683148	Server virtualization platform
EMC Atmos	http://www.emc.com/products/detail/software/atmos.htm	Private cloud storage
Eucalyptus Systems Enterprise Edition	http://www.eucalyptus.com/products/eee	Private cloud infrastructure software with APIs compatible with Amazon EC2: open source software plus proprietary extensions
Eucalyptus (open source software)	http://open.eucalyptus.com/	Private cloud infrastructure software with APIs compatible with Amazon EC2: open source software
HP StorageWorks B-series Remote Replication Solution	http://h71028.www7.hp.com/enterprise/cache/512015-0-0-0-121.html	Storage replication solution
IBM Cloud Burst	http://www-01.ibm.com/software/webservers/cloudburst/	Private cloud appliance, "Cloud-in-a-box"
IBM Scale Out Network Attached Storage (SONAS)	http://www-03.ibm.com/systems/storage/network/sonas/	Network attached storage
IBM SystemZ	http://www-03.ibm.com/systems/z/	IBM mainframe computers
IBM Tivoli Service Automation Manager	http://www-01.ibm.com/software/tivoli/products/tsam-facts.html	Provisioning automation
IBM Virtual Server Security for VMware	http://www-01.ibm.com/software/tivoli/products/virtual-server-protection/	Security software for VMware environment
Intel VT	http://www.intel.com/technology/virtualization/technology.htm	Extensions to x86 instruction set to support virtualization
Kernel-based Virtual Machine	http://www.linux-kvm.org/page/Main_Page	Open source software Linux virtualization solution
Microsoft Hyper-V	http://www.microsoft.com/windowsserver2008/en/us/hyperv.aspx	Virtualization software
MySQL	http://www.mysql.com/	Open source database software
OPNET ACE Live VMon	http://www.opnet.com/whitepapers/EMA_Vmon_Reveals_Virtual_Network_Performance.pdf	Monitoring software for inter-virtual machine traffic/transactions
Oracle Identity Management and Access Management Suite	http://www.oracle.com/us/products/middleware/identity-management/index.html	Identity management tool
Oracle VM	http://www.oracle.com/us/technologies/virtualization/oraclevm/index.html	Virtualization software
RSA Access Manager	http://www.rsa.com/node.aspx?id=1186	Identity management tool
RSA enVision	http://www.rsa.com/node.aspx?id=3170	Software to collect and analyze security event information
RightScale	http://www.rightscale.com/	Cloud management environment
Trend Micro ServerProtect	http://us.trendmicro.com/us/products/enterprise/serverprotect-for-microsoft-windows/	Anti-malware software for Windows servers
Trend Micro Core Protection for Virtual Machines	http://us.trendmicro.com/us/solutions/enterprise/security-solutions/virtualization/index.html	Anti-malware software for VMware ESX/ESXi environments
VMware ESX	http://www.vmware.com/products/esx/	Bare-metal hypervisor
VMware ESXi	http://www.vmware.com/products/esx/	Bare-metal hypervisor
VMware vSphere	http://www.vmware.com/products/vsphere/	Virtualization platform
Xen	http://www.xen.org/	Open source software hypervisor

References

[1] Mell, P. and T. Grance, October 7, 2009, "The NIST Definition of Cloud Computing," http://csrc.nist.gov/groups/SNS/cloud-computing/cloud-def-v15.doc.

[2] IDC, March 5, 2010, "Total Disk Storage Systems Turn a Corner, Posting First Year-Over-Year Gain in More Than Four Quarters," http://www.idc.com/about/viewpressrelease.jsp?containerId=prUS22236010&sectionId=null&elementId=null&pageType=SYNOPSIS.

[3] Mell, P. and T. Grance, October 7, 2009, "The NIST Definition of Cloud Computing," http://csrc.nist.gov/groups/SNS/cloud-computing/cloud-def-v15.doc.

[4] Staten, J. et al., April 13, 2009, "Deliver Cloud Benefits Inside Your Walls," Forrester Research, http://www.forrester.com/rb/Research/deliver_cloud_benefits_inside_walls/q/id/54035/t/2.

[5] IBM, "Virtual Systems Overview," http://publib.boulder.ibm.com/infocenter/eserver/v1r2/index.jsp?topic=/eicay/eicayvservers.htm.

[6] VMware, "Understanding Full Virtualization, Paravirtualization, and Hardware Assist," http://www.vmware.com/files/pdf/VMware_paravirtualization.pdf.

[7] Jones, D., June 1, 2009, "Bare Metal Hypervisors: A Group Test," ZDNet Australia, http://www.zdnet.co.uk/reviews/bare-metal/2009/06/01/bare-metal-hypervisors-a-group-test-39662582/.

[8] October 13, 2009, "Xen vs. KVM Linux Virtualization Hypervisors," SearchServerVirtualization.com, http://searchservervirtualization.techtarget.com/generic/0,295582,sid94_gci1371226,00.html.

[9] "Red Hat Enterprise Virtualization for Servers," http://www.redhat.com/virtualization/rhev/server/.

[10] KVM, http://www.linux-kvm.org/page/Status.

[11] "Microsoft Virtualization with Hyper-V: FAQ," http://www.microsoft.com/windowsserver2008/en/us/hyperv-faq.aspx#TechnicalInformation.

[12] http://www.forrester.com/rb/Research/why_isnt_server_virtualization_saving_us_more/q/id/48081/t/2., accessed October 22, 2010.

[13] February 15, 2010, "Private Storage Cloud Software and Hardware Product Guide," SearchStorage.com http://searchstorage.techtarget.com/generic/0,295582,sid5_gci1381130,00.html.

[14] VMware, "Comparison of Storage Protocol Performance, ESX Server 3.5," http://www.vmware.com/files/pdf/storage_protocol_perf.pdf.

[15] Balaouras, S. and F.Gillett, April 18, 2007, "Storage Buyer Profile: 2006," http://www.forrester.com/rb/Research/storage_buyer_profile_2006/q/id/41962/t/2.

[16] IDC, March 5, 2010, "Total Disk Storage Systems Turn a Corner, Posting First Year-Over-Year Gain in More Than Four Quarters," http://www.idc.com/about/viewpressrelease.jsp?containerId=prUS22236010&sectionId=null&elementId=null&pageType=SYNOPSIS.

[17] "EMC Atmos Product Information," http://www.emc.com/products/detail/software/atmos.htm.

[18] "IBM Scale Out Network Attached Storage (SONAS)," http://www-03.ibm.com/systems/storage/news/center/sonas/index.html.

[19] "EMC Atmos Product Video," http://www.youtube.com/watch?v=yuy-HZh0FJI.

[20] "HP StorageWorks B-series Remote Replication Solution," http://h71028.www7.hp.com/enterprise/cache/512015-0-0-0-121.html.

[21] "IBM Scale Out Network Attached Storage," http://www-03.ibm.com/systems/storage/network/sonas/features.html.

[22] "Centrify Suite for Linux & UNIX Systems," http://www.centrify.com/products/active-directory-integration-linux-unix.asp.

[23] RSA, http://www.rsa.com/node.aspx?id=1155.

[24] Oracle, October 2009, "Platform-as-a-Service Private Cloud with Oracle Fusion Middleware," http://www.oracle.com/us/technologies/cloud/036500.pdf.

[25] "OPNET ACE Live VMon Reveals Virtual Network Performance," http://www.opnet.com/whitepapers/EMA_Vmon_Reveals_Virtual_Network_Performance.pdf.

[26] "HIPPA Administrative Simplification, Regulation Text," p. 40, 164.308, http://www.hhs.gov/ocr/privacy/hipaa/administrative/privacyrule/adminsimpregtext.pdf.

[27] "RSA enVision platform," http://www.rsa.com/node.aspx?id=3170.

[28] "XenServer Auditing Tool Readme," http://community.citrix.com/download/attachments/54591507/XenServerAuditingReadme.pdf?version=1.

[29] "Altor Virtual Firewall," http://altornetworks.com/products/vf/.

[30] "IBM Security Solutions: Threat Mitigation," http://www-935.ibm.com/services/us/index.wss/offering/iss/a1031810.

[31] Trend Micro, "Cloud Virtualization Environment," http://us.trendmicro.com/us/solutions/enterprise/security-solutions/virtualization/index.html.

[32] "IBM CloudBurst," http://www-01.ibm.com/software/tivoli/products/cloudburst/

[33] "CA 3Tera AppLogic - What It Can Do for You," http://www.3tera.com/AppLogic/What-it-can-do.php.

[34] "Eucalyptus Enterprise Edition Data Sheet," http://www.eucalyptus.com/themes/eucalyptus/pdf/eee_datasheet_v4_sept92009.pdf.

[35] Right Scale, "Multi-Cloud Engine," http://www.rightscale.com/products/features/multi-cloud-engine.php.

[36] Brodkin, J., March 2, 2010, "CloudSwitch Moves Virtual Apps to Amazon Cloud," TechWorld, http://news.techworld.com/data-centre/3213977/cloudswitch-moves-virtual-apps-to-amazon-cloud/.

MITRE

www.mitre.org

MITRE

www.mitre.org