DoD ESI White Paper

IT Virtualization Technology and its Impact on Software Contract Terms

Contractual protections to consider before taking advantage of popular virtualization technology solutions.

About DoD ESI

The DoD ESI was formed in 1998 by Chief Information Officers at the DoD. To save time and money on commercial software, a joint team of experts was formed to consolidate requirements and negotiate with commercial software companies, resulting in a unified contracting and vendor management strategy across the entire department. Today, DoD ESI's mission extends across the entire commercial IT life-cycle to include IT hardware products and services. DoD ESI has established DoD-wide agreements for thousands of products and services.
www.esi.mil

Executive Summary

As virtualization has re-emerged as an important technology tool in the IT world, especially in connection with the explosive growth of cloud computing, it has become important for government software procurement organizations to understand the basics of the technology and its impact on software contract terms and conditions.

This White Paper begins by defining key terminology, exploring a brief history of virtualization, discussing current government activity in virtualization technology and the related field of cloud computing, and discussing the relationship between virtualization and cloud computing.

The latter portion addresses how virtualization in a hosted environment impacts some key software contract terms and conditions—namely license pricing and license grants, license scope (including maintenance and support, service level agreements (SLAs), and third-party software licenses), plus data ownership, storage, and security.

License pricing and grants in a hosted virtual environment differ from traditional perpetual licenses in several ways, most notably in the length of the license term, software asset possession, payment methods, and ability to create software customizations.

The multiple software and hardware resources required to create any IT environment—including off premise environments—along with other differences in contract arrangements for Software as a Service (SaaS), require additional attention when addressing the scope of the license in a software contract.

Maintenance and Support obligations need to be defined as part of the license fee and must be considered when comparing the costs of various licensing options. SLAs take on even more significance than in a standard license arrangement, because there is nearly total dependence on vendor performance. With the vendor providing all infrastructure hardware and infrastructure software in addition to the application software, more protection is required against hidden fees or licenses for third-party software.

Data storage and security are always concerns, but in a virtual hosted environment, the potential data issues and security vulnerabilities require an even higher level of diligence. Terms and conditions should ensure data ownership is clear and adequate security measures are in place. A data escrow provision (similar to source code escrow), requiring sensitive data to be deposited with an escrow agent on a regular basis (perhaps even daily), can protect against data loss or corruption.

Table of Contents

Introduction and Background 4
- **Virtualization Defined** 4
- **Other Useful Terms to Understand** 5
- **The Major Areas of Virtualization** 6
- **A Brief History of Virtualization** 6
- **Virtualization in the Federal Government** 7
- **Virtualization and Cloud Computing** 8

The Impact of Virtualization on Contract Terms 8
- **License Pricing and License Grants** 8
- **Data Ownership and Data Security** 13
- **Other Virtualization Considerations** 14

Conclusion 15

Introduction and Background

Virtualization Defined

In the last few years, virtualization has become a major force in the IT world. Before we discuss how and why this is happening—and the implications for contract terms—we will define "virtualization" and some related terms you may encounter when considering virtual hosted environments and SaaS offerings.

The term virtualization as used in the IT world can be described in many ways. One simple, conceptual way to define virtualization is to think of it as a "logical" view of something versus a "physical" view—for example, the ability to make a single physical server appear to be, and to function as, several independent servers, each with its own operating system (OS). The virtual (logical) view seen by users hides or masks the view of the actual physical hardware and its boundaries.

One high-level technical definition of virtualization is the separation of a computer operating system's service request from the underlying physical delivery of that service by the hardware—for example, a service request to access the hard drive. This separation is achieved by "abstracting" the operating system and applications.

Abstraction can be achieved by creating images of physical servers or other aspects of an IT infrastructure (networks, storage devices, etc.) through the deployment of software capable of generating and managing such images. (See the hypervisor explanation below.)

It is not possible to turn a computer with four processors into a group of virtual machines with more total computing power than what is available from the original four processors. But creating multiple virtual machines within that computer does allow for optimal use of the processors. It also allows for greater flexibility, because each virtual machine is capable of running different operating systems and applications independent of the others.

Virtualization also makes the operating system see a group of physical or virtual servers as a single pool of computing resources. The ability to move computing requirements, dynamically and automatically, to available resources within the pool at any point in time improves utilization and efficiency, thereby reducing cost. According to techtarget.com, "Virtualization can be viewed as part of an overall trend in enterprise IT that includes autonomic computing…in which the IT environment will be able to manage itself based on perceived activity, and utility computing, in which computer processing power is seen as a utility that clients can pay for only as needed. The usual goal of virtualization is to centralize administrative tasks while improving scalability and workloads." [1]

From these definitions we see virtualization has at least four important characteristics: 1) logical versus physical resources; 2) resource pooling; 3) centralized management of resources; and 4) scalability. A common thread throughout these four characteristics is lower cost. For the purposes of this paper, we will accept the business value of virtualization as a given. A more detailed treatment of the elements and proof points of the business value can be found in the DoD ESI White Paper "Best Practices for Negotiating Cloud-based Software Contracts", available from the ESI Web site.[2]

Other Useful Terms to Understand

A topic as broad as virtualization has its own lexicon. To appreciate the concept fully, including its potential impact on contract terms, it is important to have a basic understanding of some of the key terms used by the IT community. The reader should be aware that virtualization, like most technology, is an ever-evolving topic. The information presented is not intended to be an exhaustive report of the latest developments, but rather an overview of the subject designed to give the government procurement professional a better grasp of its potential impact on contract terms.

The **hypervisor** (or virtual machine monitor—VMM) is the core virtualization component. It is a piece of software that separates operating systems and other applications from their physical resources. It is also the management console for generating, designating, and managing images.

Paravirtualization is a technique for improving operating performance and overall system performance in a virtual environment by using the hypervisor to facilitate execution of instructions to and from the guest (virtual) OS instead of having the guest OS communicating directly with the host (physical) machine. Unlike full virtualization, where the OS is unmodified, paravirtualization requires modification of the OS kernel.

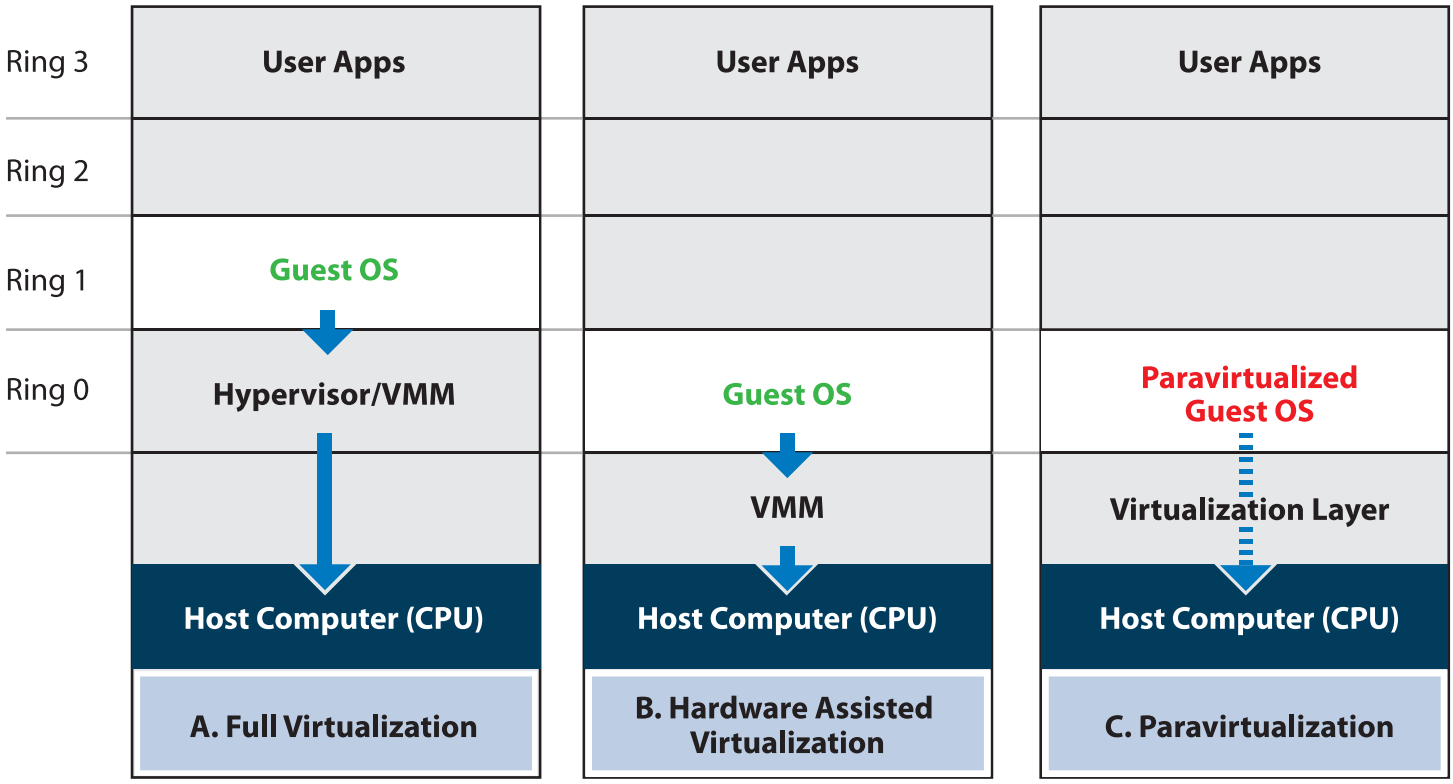

How an IT department achieves virtualization in an x86 environment can affect critical aspects of operating system (OS) software communication with the central processing unit (CPU)—even to the point of requiring modifications. Both **full virtualization** (A.) and **hardware-assisted virtualization** (B.) accommodate communication between a standard OS and the CPU. **Paravirtualization** (also called OS-assisted virtualization) (C.) requires the purchase of OS software modifications in order to communicate through the virtualization layer hypervisor, to the CPU.

A **virtual appliance (VA)** facilitates rapid deployment of SaaS offerings by providing a preconfigured application and operating system on a virtual machine. Keep in mind this is a logical view and not a physical appliance.

A **virtual machine (VM)** is an image of a real machine—a stand-alone software environment (known as a guest) that works with, but is independent of, the host operating system.

Virtualization in the application layer is a method of wrapping or encapsulating application software, to isolate it from its underlying operating system and hardware. This technique is used to improve portability of applications.

Xen is an open-source hypervisor for x86 computers—the series of computing instruction architectures based on the Intel 8086 CPU. Xen is considered paravirtualization technology because it runs on a host operating system.

The Major Areas of Virtualization

As previously discussed, virtualization can occur at various points in the IT environment. Some of the more important areas of IT infrastructure where virtualization is applied are briefly described here.

Network virtualization is a method of splitting available network bandwidth into separate channels, each of which can be assigned to a particular server or device.

Operating system virtualization is the use of virtualization software to allow a single physical piece of hardware to run multiple operating system images at the same time.

Storage virtualization is the pooling of physical storage from multiple network storage devices into what appears to be a single storage device that is managed from a central console.

Server virtualization is the replication of physical devices (hosts) into any number of virtual machines (guests) through the use of hypervisor or virtual-machine-monitor (VMM) software.

A Brief History of Virtualization

IBM is credited with the initial foray into the world of virtualization when it partitioned mainframe computers logically into separate virtual machines, allowing them to run several different applications and processes simultaneously. Since mainframes were expensive and their large capacity was often underutilized, the higher utilization afforded by multi-tasking helped increase the return-on-investment in the hardware.

Virtualization fell out of favor during the 1980s and 1990s with the decline of mainframes and the emergence of client-server architectures and applications. Small, portable, inexpensive x86 servers and desktops made distributed computing a reality.

With the explosive growth in x86 server and desktop deployments, however, some of the old mainframe inefficiencies re-emerged, along with a few new inefficiencies unique to the distributed model. These challenges included:

1) **Low utilization**—resulting from a "one application per server" mentality.
2) **Increasing physical infrastructure costs**—proliferation of relatively inexpensive servers ultimately took up more space and caused higher facility costs.
3) **Increasing IT management costs**—although the components of the client-server environment were relatively inexpensive, highly specialized talent was required for each component.
4) **High maintenance of end-user desktops**—in the distributed-computing world, managing the software and set-up of enterprise desktops became much more complicated and expensive.

In 1999, VMware solved certain technical issues regarding X86 operating system instructions that previously made X86 virtualization difficult. The new capability for X86 virtual configurations transformed large, inefficient environments into flexible, shared and efficient ones.[3]

At about the same time, one of the first movers in cloud computing—Salesforce.com—introduced the concept of delivering enterprise applications via the Internet.

Since 1999—and particularly in the last five years—virtualization and cloud computing have taken off. Companies have used virtualization to reduce costs, increase efficiencies, and improve scalability for both on-premise and off-premise environments.

The desire to outsource infrastructure and software management has been facilitated by the evolution of a business model based on using software applications more as a utility than as an on-site asset—namely, cloud computing.

Virtualization in the Federal Government

The federal government has become very active in seeking cost reductions by taking advantage of the latest developments in both virtualization and cloud computing.

According to the article "Virtualization's Next Steps", appearing in Government Computer News, "During the past year, server virtualization, which is the ability to run multiple instances of operating systems concurrently on a single hardware system, gathered momentum in the government sector. But increasingly, federal and state agencies are expanding—or at least thinking about expanding—beyond servers to apply virtualization to applications, desktop PCs and network infrastructures. The Defense Information Systems Agency, for one, is taking virtualization into the cloud. DISA recently deployed the Rapid Action Computing Environment, a cloud-computing infrastructure that lets Defense Department personnel quickly provision virtual machines so they can test and develop applications before putting them to real use."[4]

Virtualization and Cloud Computing

As we have already seen from our discussion, it is hard to discuss virtualization without mentioning cloud computing. There is a clear symbiotic relationship between the re-emergence of virtualization as a technology and the rise of cloud computing as a business model. Both help deliver optimized resources and on-demand utilization, but each approaches the matter from a different perspective. Cloud computing started outside the enterprise as a hosted, managed software model delivered via the Internet. Virtualization started inside the enterprise as an efficiency and scalability tool and moved out to hosted environments.

John McCarthy, the father of artificial intelligence and an early advocate of timesharing on mainframes, made an early reference to computing as a ***public utility*** in a 1961 speech commemorating MIT's centennial.[5] This public utility concept is important to understanding how and why the convergence of technology (Internet, virtualization, hosting) with a business model (subscription licensing, self-provisioning, shared software) has resulted in today's SaaS offerings available on the "public" cloud.

The Impact of Virtualization on Contract Terms

The focus of this section will be the contract terms impacted by off-site virtualization (versus on-site virtualization), unless specifically noted in the discussion of each contract term.

License Pricing and License Grants

When we look at the various "technology software" required to establish an IT environment for running application software, it is easy to see how complicated that landscape can become. At a minimum, the operating system, database, and application software all need to be accounted for (along with many other software packages used to create and manage a robust IT environment). The user needs to understand the nuances of each software product's licensing considerations, particularly when it comes to calculating the most economical price, regardless of the metric used to establish that pricing—named users, processors, or some other metric.

This complexity is amplified by the introduction of virtualization, not only because of adding another software layer (the virtualization software) to the mix, but also because of the multiplicity of software and user configurations created by having several virtual images of each type of software found in the IT environment.

Publishers have created appropriate mechanisms to account for named users or processors in a virtual environment, depending on the type of software and the preferred unit of measure. (For one interesting example, see the blog entry on Oracle database licensing in a virtual environment.)[6]

If the government is buying software to use in an on-site virtual environment, these pricing mechanisms and rules must be thoroughly understood in order to make the most economical purchase.

For purchases involving off-site virtual environments, the user is shielded from the complicated licensing formulas applied to the various software products needed to create the virtual environment. The hosting company (or cloud provider, if the virtual environment is in a cloud) manages the individual licensing matters. Even in the case of purchasing off-site services, however, the pricing mechanisms are still critical to appreciating how the monthly hosting fees and subscription license fees (or combined SaaS license fees, if a cloud is involved) are constructed, so that valid comparisons can be made to on-site alternatives. In other words, how else would we know if the proposed SaaS fees are appropriate and fair?

To continue our discussion, we will assume the government is contracting for a SaaS license in a virtual cloud environment that would include all required infrastructure and application capabilities.

The key SaaS license grant contract term for the government procurement organization to consider focuses on two primary issues:

1) **The type of license granted ("Right to Use" License):** When compared with the typical perpetual license granted by a traditional software license, the "right to use" license grant employed in the SaaS model does not contemplate possession or extensive customization of the software by the licensee. Rather, the license grants a right to use the software for a limited period of time, typically defined as a number of months or years, and requires a periodic payment of a subscription fee. Please refer to "Best Practices for Negotiating Cloud-Based Software Contracts"[7] for additional details.

2) **The scope of license grants:** The scope of the "right to use" license rights refers to several items:
   - Support and Maintenance Services.
   - Service Level Agreements (SLAs).
   - Inclusion of all third-party software in fees and other covenants.

Maintenance and Support Defined. It is vital for the government to specify the types of maintenance and support services to be provided by the vendor as part of the license fee. Whether these items are priced separately or are blended into the SaaS fees, they need to be accounted for in the cost comparison analysis referred to above.

Maintenance typically refers to the ongoing provision of fixes and patches to address specific software faults as well as the provision of new software releases. It is important to know whether the vendor will need to take the system off-line to apply fixes, patches, and new releases, and to know the frequency of such occurrences. The time required to take the system off-line is referred to as "planned downtime."

Another important aspect of upgrading to new releases depends on the "purity" of the SaaS model under consideration. A true SaaS deployment implies multi-tenancy, meaning all customers are running on the same instance of the (application) software. Probably the best example is salesforce.com, which hosts about 43,000 customers on eight instances—each hosted in one of eight IT environments located around the world. When a new release of the software is applied, all customers are upgraded simultaneously.

Since this impacts customer flexibility for whether (or when) to upgrade, it is important to know if the application being licensed requires simultaneous upgrades for all customers or if there is some customer discretion. It is also important to know whether the customer or the vendor is responsible for ensuring all customizations work in the new release (regression testing).

Support refers to the process of a customer reporting an apparent software fault or bug to the vendor for diagnosis and correction. Support is usually divided into three levels indicating the progression of the problem through the diagnosis and correction process (not to be confused with problem severity which has to do with SLAs for problem response and correction).

First level support simply means the initial reporting of a problem. In many cases, the customer is responsible for first level support, since many problems are easily identified and resolved as user education issues. In SaaS licenses, vendors will sometimes offer first level support. Since every service has a cost, the government needs to understand whether this service can be provided within its user community or should be provided by the vendor.

Second level support is used when the first level support is unable to resolve the problem. In SaaS licenses, the vendor typically provides second level support. Second level support resources are more experienced with the software and the underlying technology. Most problems are resolved at second level support.

When the issue is not resolved at the second level, the problem is referred to third level support, which is usually the Publisher's software development organization (sometimes called R&D). Some Publishers identify the person who wrote the allegedly defective code and assign him/her the task of resolving the problem. Referral to third level support indicates a relatively severe code problem requiring software experts to diagnose and fix the problem.

In all cases, the customer has a vested interest in making sure their software faults are promptly received, diagnosed, and fixed by the vendor. The timeliness and adequacy of this performance becomes the subject of SLAs with the vendor.

Service Level Agreements (SLAs). As stated in the previously mentioned DoD ESI White Paper, "The key components of any solid SLA should contain specific, measurable, and enforceable terms and conditions that the SaaS provider must adhere to for each component of the service provided. If the provider fails to meet an obligation under the SLA, the SLA must have the 'teeth' to help mitigate such failure from happening again. The 'teeth' to SLAs are the specific remedies that apply when the provider's obligations are not met. (In common commercial practice,) the remedies usually take the form of monetary damages that are pre-agreed between the parties for specific failures, and/or credits for future services."[8]

There are three basic types of SLAs in SaaS licenses—one related to the performance of the hosted environment, and two related to maintenance and/or support:

1) System availability (i.e. the performance of the hosted environment);

2) Response times to reports of software faults (i.e. support);

3) Response times for providing fixes to actual software faults (i.e. maintenance and support).

The following table provides an example for calculating system availability.

Criteria	Measurements	Comments
Minutes in a 90 day period	129,600 minutes	
Planned down time (assume 18 hours)	1080 minutes	*This is a standard amount of time for system maintenance*
Remaining minutes for scheduled up-time	128,520 minutes	
SLA	99.9%	*This is a moderate standard; 5 nines (99.999%) is very high*
Minutes of expected up time	128,391.5 minutes	
Allowable minutes unplanned downtime	128.52 minutes ~ 2.1 hours over 90 days!	*Little time for unplanned down time*
Penalties	Varies	*Usually a credit is given for missing the SLA*

> Please note this example is based on a three-month period, assuming planned downtime of 18 hours for system maintenance and upgrades. Scheduled uptime is the time remaining after subtracting planned downtime from the total number of minutes available in a three-month period. The specified service level is expressed as a percentage of scheduled uptime (in this case, 99.9%).

The next table shows a sample of response times and fix times for various levels of reported software Issues.

Issue Severity	Response Time to Acknowledge Issue	Response Time to Fix Issue
Level 1 (LOW) Does not have significant impact on users	Return call or email within 8 hours	Provide fix within 30 days
Level 2 (MODERATE) Causes some user issues, but most processes are functional	Return call or email within 4 hours	Provide fix within 5 days
Level 3 (HIGH) Significant impact on system use	Return call or email within 1 hour	Provide fix ASAP (24 hours or less)

As previously stated, SLAs need to have "teeth", usually in the form of financial penalties for failing to meet contracted service levels. Additionally, the government should seek to reserve the right to terminate a SaaS agreement for chronic failure to perform. This means, if the vendor repeatedly misses any of the three SLAs, it might not be enough to simply impose financial penalties. The government should be able to terminate the contract since it is not receiving the value of the software it bargained for. The parties need to define clearly the parameters of "chronic failure to meet the SLA" or "repeatedly misses the SLA." This often involves defining the frequency and degree of missing the SLA over any rolling monthly or other agreed-upon period. The government should have a back-up plan or exit strategy on how services will be continued or provided if the contract is terminated.

Third-Party Software. Many Publishers have incorporated various third-party software components in their proprietary products. One common example found today is the growing use of Open Source software embedded in or working alongside a Publisher's product as part of its proprietary software. Open Source software is generally created by independent software developers (or teams of developers) for use by individuals or companies (even software Publishers) with limited use restrictions. Popular examples include Apache HTTP Server software, Mozilla Firefox browser, the MySQL enterprise database software, and OpenOffice text, spreadsheet, database, and presentation software for desktop users.

Whether the government is procuring a standard perpetual license for a Publisher's software product or is entering into a SaaS agreement, the license or SaaS agreement should include a provision stating the Publisher or service provider has authority to use the third-party software with the applications, while also specifying that no additional fees are due from the government for the use of the third-party software. The license or contract provision should take the form of a covenant accompanied by a provision indemnifying the government from any potential infringement of third-party intellectual property or other unlawful, unauthorized use.

Additionally, many Publishers' products—particularly business applications—require other commercial software to enable them to operate. The best example is the need for most applications to use a third party's database software.

When procuring a standard perpetual license, it is common practice for the Publisher to require the buyer to procure a database license from the database Publisher. Some application Publishers have authorization from database Publishers to sell "limited use" or "run time" database licenses with their applications.

In the perpetual license scenario, it is important for the government to receive full disclosure from the application Publisher for all supplementary software products required to operate the application, as well as those products compatible with the application (often referred to as "supported by" the Publisher's products). In addition to this full disclosure, the license should specify which supplementary products are included with the application license fee and which ones must be procured separately.

When entering into SaaS agreements, the database licenses and all other licenses required to operate the application should be included in the SaaS fee. The covenant authorizing their use and the indemnification against infringement should be included in those agreements.

Data Ownership and Data Security

Data ownership is a concern any time the government's data is used or stored outside the physical control of the government. The very nature of an off-premise virtual environment, "public" cloud, or SaaS agreement implies off-site storage of the government's data. Although it may seem to be a given, the fact that the data is out of the government's physical possession and control makes it important for the agreement to include provisions acknowledging that:

- the government always retains legal ownership of its data regardless of the data's location,
- the data must be provided to the government, in an agreed upon format and timeframe, upon the occurrence of certain events, including a demand from the government for its data, (for example, the agreement should specify that the data should be provided within 10 days of receipt of written notice to the vendor)
- and that, upon such occurrence, the service provider agrees to retain no copies of the government's data.

In addition to key considerations regarding data ownership, data security and overall system security in virtual environments are also major concerns. The proliferation of virtual machines and their operating systems can lead to more targets—and additional potential vulnerabilities—for security breaches. For comprehensive discussions of the topic, see "When Virtual is Harder than Real: Security Challenges in Virtual Machine Based Computing Environments"[9] by Tal Garfinkel & Mendel Rosenblum as well as "Server Virtualization: Top Five Security Concerns"[10] by Kevin Fogarty as it appeared at cio.com.

In January 2011, the National Institute of Standards and Technology (NIST) published its "Guide to Security for Full Virtualization Technologies"[11], Special Publication 800-125, by Karen Scarfone, Murughia Souppaya and Paul Hoffman (linked from http://www.nist.gov/itl/csd/virtual-020111.cfm). The document focuses its recommendations on security measures for the hypervisor, the guest OS, virtualized infrastructure (networks and storage), and desktop virtualization.

Aside from those general security concerns, potential vulnerabilities regarding data confidentiality, data back-up provisions, and the risk of losing data in transit or at rest are all heightened in a virtual environment.

In all environments, data security is typically governed by system settings and security procedures that make storage resources available only to authorized users and networks. Those settings and procedures are particularly important when data storage or database software is virtualized. The government's data might be stored where other data has resided or, conversely, other data might be stored where government data once resided. Remnants of sensitive data, passwords, and encryption keys could be left behind. Other security risks exist, including latent viruses and incomplete or missing security patches designed to combat them.

All of these issues impact a government procurement organization's need to ensure adequate contract provisions exist for guaranteeing appropriate data security measures are in place in the vendor's environment.

An article entitled "PCI Security Standards Council Releases Guidelines for Virtual Environments",[12] posted on databreachlegalwatch.com by Nicolai Schurko, Esq. provides the following recommendations to address these and other potential issues when entities use virtual environments to store sensitive data:

- "'Hardening' (securing) the hypervisor;
- Implementation of appropriate physical access controls;
- Implementation of a 'defense-in-depth approach' that encompasses preventive, detective, and responsive controls to secure data and other assets;
- Using multiple methods to secure administrative access, such as implementing two-factor authentication or establishing dual or split-control of administrative passwords between multiple administrators; and
- Ensuring administrative, process, and technical segmentation to isolate each hosted entity's environment from the environment of other entities."

Other Virtualization Considerations

While the recommendations above might sound more like technical actions or processes to be undertaken when the environment is on site, they can also be the subject of contract terms and conditions to ensure the vendor is doing everything reasonably possible to protect the government's data.

This excerpt of an article on simplysecurity.com underscores the importance of being vigilant. "Virtualization has been lauded as an important piece of the future of enterprise computing, as companies seek out innovative means for dealing with the data boom of recent years. With the technology, companies maximize their data storage space to ensure they retain all the information necessary for short- and long-term operations. However, as a new report can attest, it is imperative that IT departments keep data security considerations in mind when deploying virtual machines. Failing to do so may result in data loss, or leave the company vulnerable to an outside attack. According [to] the respondents of a recent poll from data management solutions provider Kroll Ontrack, data loss within virtual environments has spiked 140 percent this year compared to 2010. That was evidenced by the 65 percent of respondents who said they frequently suffer a data security incident within their virtual environment. Among respondents, 53 percent said they have suffered five data loss incidents pertaining to virtualization in the past year. Another 12 percent have experienced more than five, the report found."[13]

For particularly sensitive data, the government might want to specify in its SaaS agreements that its data will reside only on dedicated physical storage devices—in the hosted environment, or possibly even on-site at a government data center. Dedicated storage resources would change the virtual environment to a partially virtual solution, which would result in changing the cloud environment from a "public" cloud to a "hybrid" cloud. Another alternative is to require the vendor to provide specific technical configurations designed to prevent data leakage issues. The government's IT support team would need to validate that the configuration is sufficient to protect the government's data, and the procurement team would need to ensure the contract stipulates vendor obligations for the specific configuration.

In addition to contract terms focused on security processes, guarantees, and physical changes to the virtual environment, the government can also require data escrow. This practice, similar to the

concept of source code escrow, can require the vendor to deposit a copy of critical SaaS application data with a third party on a regular basis, to be released to the government upon the occurrence of specific trigger events—including vendor bankruptcy, data loss, or data corruption.

The desire for data retention is not an uncommon occurrence, inside or outside of government applications. According to the "Software as a service" entry on www.itechsolutions.in, "research conducted by Clearpace Software Ltd. into the growth of SaaS showed that 85 percent of the participants wanted to take a copy of their SaaS data. A third of these participants wanted a copy on a daily basis." [14]

Conclusion

Virtualization is a proven, valuable technology that has a long history dating back to the 1970s. Client server architectures of the 80s and 90s initially rendered virtualization almost obsolete, but with the proliferation of distributed systems, ubiquitous use of application software, and massive amounts of data, the efficiencies and scalability of virtualization have helped it regain its relevance.

The simultaneous increase of new software licensing models, growing interest in outsourcing both infrastructure and application management, and the proliferation of the Internet, have converged with virtualization technology to create an explosion of cloud computing in hosted environments.

Not surprisingly, the implications of virtualization for government software procurement are very similar to the considerations for cloud computing. These include the pricing, type and scope of license grants, maintenance and support services, SLAs, data ownership, and data security.

Government procurement organizations should become familiar with the important elements of virtualization and cloud computing to make good decisions about their utility for users, their security, and their costs as compared to other alternatives.

If the government decides to seek SaaS licenses in a cloud with virtualization, the procurement teams need to ensure adequate contract terms and conditions are in place to address the concerns outlined in this paper and to spell out appropriate vendor duties along with government rights and remedies.

About the Author

Tom Crawford is a seasoned executive and consultant who has been leading, growing and advising technology businesses for the past 20 years. After successful senior management roles with leading software companies—including SAP, PeopleSoft, Oracle, and BMC—Tom started his own technology consulting business, lending his expertise to a variety of software-related accounts. His work with DoD ESI draws upon his diverse experience helping clients save significant dollars in software and services procurement, as well as experience as a leader or participant on teams that have closed and negotiated scores of software and services contracts ranging up to $65 million. Tom is a graduate of the U.S. Naval Academy, a former U.S. Navy officer, and holds an MBA from Wharton and a Juris Doctor from the University of Pittsburgh.

[1] Definition of virtualization. Found at http://searchservervirtualization.techtarget.com/definition/virtualization (Last accessed July 31, 2012)

[2] White Paper, "Best Practices for Negotiating Cloud-based Software Contracts", by Gretchen Kwashnik, published by DoD ESI. Found at http://www.esi.mil/LandingZone.aspx?id=273&zid=3 (Last accessed July 31, 2012)

[3] Web site, http://www.vmware.com/virtualization/history.html (Last accessed July 31, 2012)

[4] Article, "Virtualization's Next Steps", by Rutrell Yasin, Government Computing News, February 9, 2009. Found at http://gcn.com/articles/2009/02/09/virtualization-extends-reach-in-agencies.aspx (Last accessed July 31, 2012)

[5] Article, "The Cloud Imperative", by Simson L. Garfinkel, October 3, 2011. Found at http://www.technologyreview.com/business/38710/ (Last accessed July 31, 2012)

[6] Blog Entry, "Understanding Oracle Database Licensing Policies" by Neeraj Bhatia, January 17, 2011. Found at http://neerajbhatia.wordpress.com/2011/01/17/understanding-oracle-database-licensing-policies/ January 17, 2011 (Last accessed July 31, 2012)

[7] White Paper, "Best Practices for Negotiating Cloud-based Software Contracts", by Gretchen Kwashnik, published by DoD ESI. Found at http://www.esi.mil/LandingZone.aspx?id=273&zid=3 (Last accessed July 31, 2012)

[8] White Paper, "Best Practices for Negotiating Cloud-based Software Contracts", by Gretchen Kwashnik, published by DoD ESI. Found at http://www.esi.mil/LandingZone.aspx?id=273&zid=3 (Last accessed July 31, 2012)

[9] White Paper, "When Virtual is Harder than Real: Security Challenges in Virtual Machine Based Computing Environments", by Tal Garfinkel & Mendel Rosenblum, Stanford University Department of Computer Science. Found at http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.103.2536 (Last accessed July 31, 2012)

[10] "Server Virtualization: Top Five Security Concerns" by Kevin Fogarty, May 13, 2009. Found at http://www.cio.com/article/492605/Server_Virtualization_Top_Five_Security_Concerns (Last accessed July 31, 2012)

[11] National Institute of Standardds and Technology (NIST) Special Publication 800-125, "Guide to Security for Full Virtualization Technologies", by Karen Scarfone, Murughia Souppaya and Paul Hoffman. Found Linked fromat http://www.nist.gov/itl/csd/virtual-020111.cfm (Last accessed July 31, 2012)

[12] Article, "PCI Security Standards Council Releases Guidelines for Virtual Environments", by Nicolai Schurko, Esq., June 26, 2011. Found at http://www.databreachlegalwatch.com/2011/06/pci-security-standards-council-releases-guidelines-for-virtual-environments/ (Last accessed July 31, 2012)

[13] Article, "Virtual IT Environments Requiring Tougher Data Security Measures", by Simply Security, November 29, 2011. Found at http://www.simplysecurity.com/2011/11/29/virtual-it-environments-requiring-tougher-data-security-measures/ (Last accessed July 31, 2012)

[14] Article, "Software as a service", http://www.itechsolutions.in/issue_04.html#article1 (Last accessed July 31, 2012)

DoD ESI is an official
Department of Defense initiative
sponsored by the Department of Defense
Chief Information Officer (DoD CIO).

Your Preferred Source for
IT Acquisition Across the DoD

BEST VALUE

EFFICIENT

LOW RISK

VOLUME DISCOUNTS

UNIFIED VOICE

Visit DoD ESI online at **www.esi.mil**

Department of Defense Chief Information Officer
6000 Pentagon
Washington, DC 20350-6000