Critical Analysis of an Organization's Network

Andrew Bruce, CISSP, PMP, FITSP-D

CTO, RiVidium Corporation (http://www.rividium.com/)

andy.bruce@rividium.com

20 November 2010

Topic Summary:

- Describe and analyze the current state of our network
- Describe and analyze how this network reached its current state
- Analyze the anticipated future business needs, and the anticipated demands on this network
- Propose recommendations

Executive Summary

Department of Defense (DoD) information systems contractors face considerable problems today. Despite the efforts starting with then-Secretary Rumsfeld's Defense Transformation initiative in 2001, the nation has seen years of increased spending due to the conflicts in Iraq, Afghanistan, and the overall Global War on Terror (GWOT). The present gives us quite a different landscape: President Obama, Defense Secretary Gates, and Chief Information Officer Kundra have imposed real spending restrictions throughout the Federal Government. In this paper, we examine how this heightened cost awareness has impacted our organization's security posture.

Our Approach

Government is demanding that private industry improve delivery efficiency while simultaneously decreasing costs – all the while meeting ever-more-stringent information assurance requirements. To achieve these goals, organizations of all sizes must ensure that their networks are run with a strict eye toward cost containment. The key to cost containment lies in increased productivity (powered by scalability) and automation (powered by a centralized network infrastructure). Our organization's network is the starting point for this analysis. The goal is to analyze how increased productivity and automation have been deployed using our industry's best security practices to provide a high level of confidentiality, availability, and integrity. The results of our security decisions are reviewed and the lessons learned are presented. We include a series of recommendations to enable decision makers to select new technology solutions to ensure a continued strong security posture while minimizing procurement and maintenance costs.

Our Audience

This paper is addressed to fellow corporate officers, corporate information managers interested in reviewing and evaluating possible security solutions to common infrastructure issues, general security practitioners curious to see how these common issues have been addressed within a real-world environment, and students interested in learning about the fascinating and complex world of security management. As succinctly noted by Krehnke and Krehnke, "each organization must establish and maintain an effective enterprise information security architecture that contributes to its own security, its employees, customers, business partners — and that of the nation." [**]

Our Goal

By providing a working case study on how network upgrade decisions can help to decrease costs while simultaneously increasing security, other organizations may use this information to help in providing the best possible value to Government customers. Of particular interest to those of us in the DoD space, improvement in the way that Government functions overall ultimately translates to better support for our primary customer – the Warfighter.

[**] Peter R. Stephenson, ed. *Information Security Essentials: Section 2* (Auerbach Publishing, ISBN 978-1-4398-0030-0, 2009), pg. 3.

Table of Contents

Illustration Index

About the Author

Andrew Bruce is Chief Technical Officer for RiVidium Incorporated. RiVidium provides professional services to the Federal Government and Department of Defense. Its core competencies include customizing and developing architecture and governance models leveraging our proprietary technologies as well as designing / building engineering solutions. Mr. Bruce's job responsibilities include: working directly with customers and partners for new business development, managing RiVidium's network infrastructure, working with project managers to ensure project completion, and managing software development efforts throughout the entire system life-cycle. Mr. Bruce is focusing on Information Assurance to achieve his goal of building and managing large data centers providing cloud computing utility services for commercial and Government customers.

1.0 Introduction

In the world of Government and DoD contracting, the requirements and expectations on private industry have shifted dramatically towards cost control and improved time-to-completion. After years of increased spending with the conflicts in Iraq, Afghanistan, and the overall Global War on Terror (GWOT,[1] now known as the Overseas Contingency Operation),[2] we see a different mindset today. With the economic downturn of 2007 and the ensuing sluggish recovery, increasingly loud calls have been made to lower overall Government spending. As far back as 2007, the American Association for the Advancement of Science warned that "R&D in the Defense Agencies would fall […]3.4 percent" in fiscal year 2008.[3] The change in administration has made this imperative more pronounced, with Defense Secretary Gates announcing that he "wants to cut $100 billion over the next five years, and reduce defense contracts by ten per cent."[4] More to the point, Chief Information Officer Kundra has issued an Office of Management and Budget Memorandum effectively halting work on a number of high-profile modernization efforts, saying that we need to "end a culture in which we continue to throw good money after bad."[5]

However, there is no corresponding relaxation of the need to maintain compliance to all existing laws, regulations, and policies. Delivered systems must continue to maintain or improve the existing security standards as well as adapt to changes in these standards.[6] In that light, this paper covers the following:

1. Analyze the state of our current network, especially in light of improving our security posture to meet ever-more-stringent Federal and DoD requirements.
2. Review how our network decisions over the past year were shaped by changes in Information Assurance thinking and practice within the DoD, how these decisions supported our overarching business drivers of improved scalability and cost containment, and whether we were able to realize the anticipated benefits from those changes.
3. Plan for our future business needs, especially as they pertain to our network infrastructure and Information Assurance.
4. Make specific recommendations on how we can improve our network initiatives, accounting for the best possible use of scarce resources.

One key to our continued success lies in improving our overall productivity in managing our network. To meet these goals while still improving our network productivity, we concentrate our focus upon *scalability* and *automation* in the management of our network's infrastructure and the overall delivery of confidentiality, availability, and integrity.

2.0 The Current State of our Corporate Network

Business networks exist to enable core business objectives. In this section, we cover our own business objectives, the Information Technology (IT) processes that support those objectives, and how Information Assurance (IA) has shaped those IT processes.

2.1 Our Business Goals and Objectives

Our primary business goal is to provide the most efficient and productive solutions for our customers, always keeping in mind that while the customer may be the DoD or some affiliated agency, the ultimate *beneficiary* of our hard work is a twenty-two year old Soldier serving in harm's way.[7] Three specific strategies help to accomplish this goal:

1. *Building trust* with our existing customers by meeting or exceeding contractual requirements.
2. *Pursuing additional work* by teaming with other large and small vendors.
3. *Expanding our model* to include commercial opportunities as well as state and local government contracts.

2.2 IT Processes Supporting Goals and Objectives

We use our network to help deliver our message to our target audience, educate and share knowledge with our colleagues and end-users, enable collaboration between our team members, and to provide a solid hosting environment for proofs-of-concept.

2.2.1 Corporate Presence

To be successful, we must deliver our corporate message to our target audience: enterprise architecture analysts, Defense and Government agency decision makers, and partners interested in teaming with us for current and emerging opportunities. This translates to a set of IT requirements including:

- Web server – Must be highly available;
- Really Simple Syndication (RSS) feeds;
- Social Web (Web 2.0) support, including Twitter broadcasts, Facebook integration; and,
- Technology Demonstrations.

2.2.2 Search Engine Presence

When our target customers seek information from the Internet, we want our content to be available easily and at any time. Our corporate Web site provides and publishes resources related to our areas of expertise (properly tagged to be found by Web crawlers):

- Free white papers on industry topics;
- Sponsorship of open source initiatives; and,
- Interactive user experience via blogs and Wikis.

2.2.3 Collaboration

We use IT to support:

- Phone and email support (our phone system uses Voice over IP, or VOIP);
- Internal corporate portals supporting business functions including Human Resource and Payroll functions;
- Secure instant messaging;[*] and,
- Controlled remote access from our team members, partners, and customers to corporate assets.

2.2.4 Hosting Environment

A significant part of our revenue comes from building and hosting proofs-of-concept for our customers. Multiple hosted customer systems must be capable of running simultaneously without having knowledge of or impact upon each other. Additionally, our future plans include moving towards compliance to the National Industrial Security Program Operating Manual Supplement (DoD 5220.22-M-Sup 1).[8]

2.3 Information Assurance and our Organization

IA has a significant impact on the way that we do business. While all organizations realize the dangers inherent in providing a public-facing network, we have additional concerns.

- *Protecting corporate assets* – The sensitive nature of our work requires us to exercise due care in how project information is stored, transported, and accessed.
- *Security compliance* – Our development projects generally are built for use within a secure computing facility, thus we must be aware of relevant Army and DoD security requirements.
- *Contractual obligations* – Teaming agreements (TAs) with our partner organizations require us to maintain documentation on our internal network processes.
- *Organizational certifications* – We are currently pursuing Capability Maturity Model – Integrated Level 2 (“Managed”) as a strategic move to enhance our ability to qualify for new work.

[*] One security problem with standard instant messaging solutions like AIM is that all messages are sent through and stored on a central server maintained by the provider in cleartext. In our organization's case, we use an open source solution called OpenFire (http://www.igniterealtime.org/projects/openfire/) that allows us to host the instant messaging central server in our own secure network, minimizing the potential for our team members’ instant messages to be tracked or hacked. We ensure confidentiality and integrity by requiring all communications over our OpenFire server to be encrypted using SSL (this is a configuration option available for the OpenFire server).

2.4 Current Network Configuration

The following high-level diagram provides an overview of our corporate network and its functions. (We do not show the physical devices such as firewalls and routers.) We use a combination of physical machines for critical processes that require high performance and virtual machines (running under a hypervisor) for the vast majority of our systems.

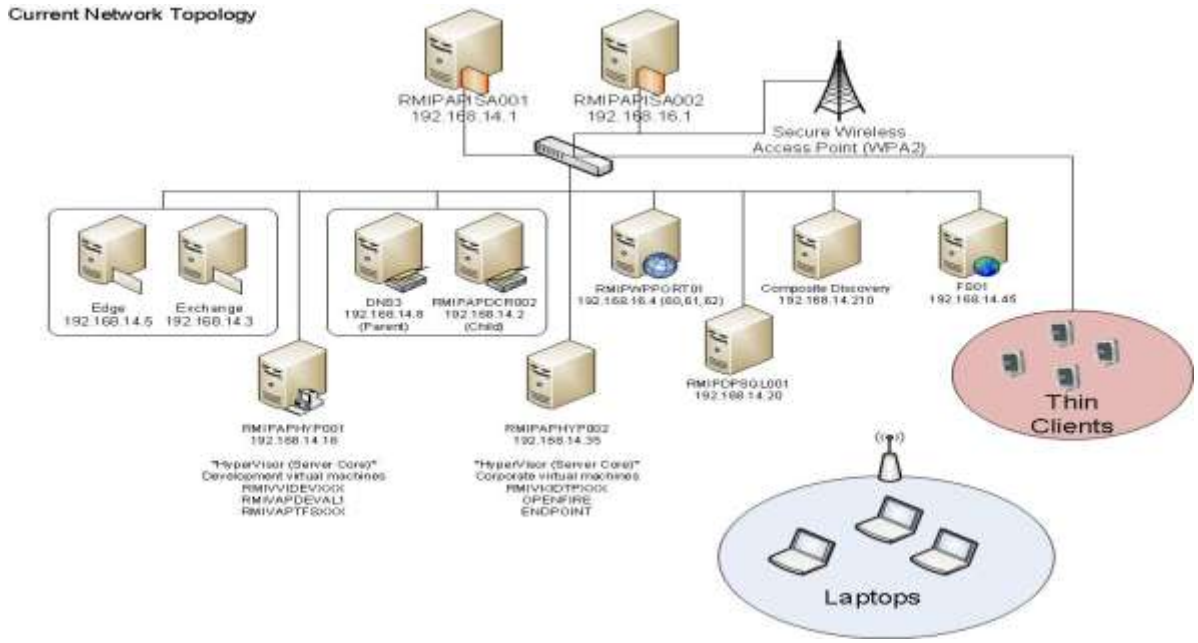

Figure 1 - Current Network Topology

2.4.1 Business Objective Support

Our existing network configuration supports our business objectives by providing a specific computing resource solution for each of our needs, as we map below:

- *Corporate Presence:* Our portal server RMIPWPPORT01 is a powerful physical server that serves our corporate Web site and RSS feeds. Actual Web site content is maintained in our corporate database server RMIPDPSQL001, and access is controlled by our domain controllers (Active Directory).
- *Education and Knowledge Sharing:* We use our portal server to provide white papers and host our "contact us" capabilities.
- *Collaboration:* We use our eponymous email servers EXCHANGE and EDGE to provide email support, and we have a separate Voice-over-IP (VoIP) network to provide phone support. We use an open source instant messaging package called OpenFire (secured using SSL) on our aptly-named OPENFIRE server. Finally, our corporate gateways RMIPAPISA001 and RMIPAPISA002 provide Virtual Private Network (VPN) access for staff on travel.
- *Hosting Environment:* As needed, we install, provision, and deploy physical servers within our network environment for customers. These servers are always installed in their own security domain with

limited access from our corporate network (typically one-way communication from trusted hosts only). We can also provide virtual machines that are logically segmented into their own security domain by using one of our hypervisor servers; we can protect access to these virtual machines using the same techniques as if they were physical machines.

- *Development Staff:* On our hypervisor servers, we provide a set of "desktop" (RMIVKIDTPXXX) and "developer" (RMIVVIDEVXXX) virtual machines (not physical desktops). We have a number of standard configurations that we use based on the specific end-user's needs; most administrative and marketing users get a standardized Windows XP setup while most developers get a standardized Windows Server 2008 Standard Edition setup. Typically, end-users are not allowed to store documents on their assigned virtual machines; instead, we use our corporate file server FS01 to store all user documents (allowing us to manage and govern corporate data files from a central location).

2.4.2 IT Processes Support

Our IT processes support our business functions by ensuring that our network remains secure and available. These processes include supporting new hires, providing help desk support, automatic corporate-wide data backup, centralized system access control, vulnerability scanning, and network threat detection. Our network supports these processes by emphasizing centralization of management functions, especially in relation to managing our user base. We avoid issuing physical machines to new hires, instead assigning them a "thin client" that can only connect to a specific assigned virtual machine. This allows us to ensure standardization, apply system updates, and troubleshoot problems without ever needing to be at the user's station.

Corporate data backup support is made simpler by our document storage policy. By requiring most users to store all documents on the corporate file server we ensure that we can control software installations as well as providing a centralized location to find backup data. When personnel are assigned laptops (such as during travel) they are provided with automatic synchronization software to manage locally stored files – primarily these laptops are running Windows XP.

We control system access by using a centralized user directory within our network domains (Active Directory). We leverage Microsoft's Group Policy Objects (GPOs) to ensure that all machines connecting to our network meet minimum security standards; for guest access we provide a separate "quarantine" area that allows Internet access but no access to any system on our corporate network. We also use GPOs to install and configure vulnerability scans via the DoD-approved "Retina" vulnerability scanner as per the Joint System Administrator Checklist,[9] and to install anti-virus from our ENDPOINT central anti-virus server.

2.4.3 Responding to Security Needs

In our organization, security needs can be expressed as mitigation or remediation efforts. Mitigation efforts attempt to prevent security events by using defense-in-depth and include:

- Educating our team members to be responsible end-users based on our corporate Security Policy.
- Controlling ingress to and egress from our network via our firewalls and edge authentication servers.
- Detecting malicious network activity via our Network Intrusion Detection Systems (NIDS).
- Detecting malicious host processes via our Host Intrusion Detection Systems (HIDS).
- Providing redundancy to account for failures (for example RAID 5 disk storage).

We address these needs in a number of ways. New hires must all sign an Acceptable Use document for the corporate network (based on best practices as outlined by the Department of Veterans Affairs).[10] As a practical example of our network support for this process, we use our GPO capabilities to support this requirement by providing sign-on banners as shown below. Our configuration addresses ingress / egress by using commercial firewalls to prevent network attacks and VPN servers to control remote access. We address intrusion detection by automatically deploying anti-virus software and updates to machines as they connect to the network. Redundancy support is made simpler by our usage of a virtualized environment where we have the minimum number of physical devices; while this approach does increase the risk of a single error affecting multiple servers, for our small organization the advantages of centralized management outweigh this risk.

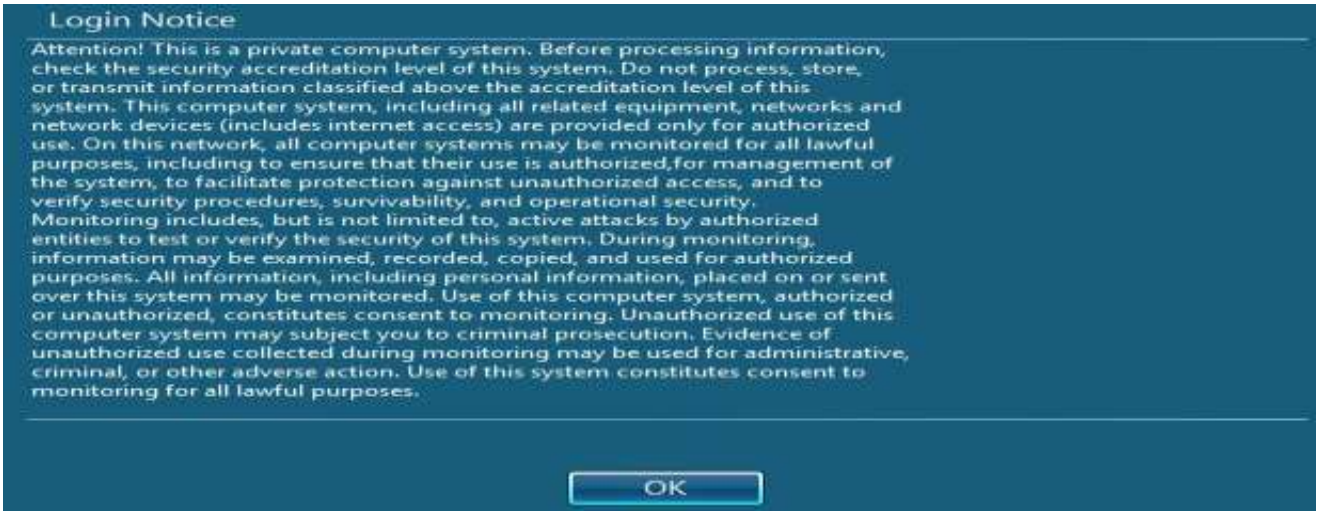

Figure 2- An example of a login banner

Remediation efforts recover from security events based on risk identification and a tested recovery plan. These risks include:

- Loss of Internet service
- Hardware failure
- Data destruction
- Security breaches (such as a successful attack into the network)

Our network configuration addresses loss of Internet service by providing multiple Internet service providers. We address hardware failure by building in redundancy at key points; for example, we have two domain controllers and either is capable of handling our network load. Data destruction risk is mitigated by our corporate backup policy and our high-capacity, redundant data storage devices. Finally, we can respond to network security breaches by isolating the impacted machine(s) in a quarantine area where we can examine them with little fear that the infection can spread.

3.0 Network Decisions over the Past Year

Our organization was born from a group of individuals working together informally as analysts over the past five years. The author's own position with the company started shortly after its current incarnation in 2009, the network at that time reflected the company's from-the-ground-up start-up environment. The corporate infrastructure consisted of a number of physical servers, including older "beater boxes" built with after-market (consumer) components. The servers and all network equipment were on long tables within an open and carpeted area of the "development bullpen." Each company resource had her own laptop connected via wireless to the company backbone (developers, analysts, executives, even the receptionist). In some cases the laptop was employee-owned.

Over the past twelve months we have made a number of significant changes to the network infrastructure, driven by a number of factors:

- *satisfying our current customers* and gaining trust to help us grow our business within the DoD;
- *providing for reliable and scalable infrastructure* in the face of rapid growth while minimizing administration overhead costs; and,
- *protecting company assets* and minimizing physical device usage.

We examine each of these areas in more detail below, closing with our lessons learned.

3.1 Satisfying our Customers

Our DoD customer projects include custom software packages built for deployment within a secure environment as well as proofs-of-concept to be hosted within our environment. To keep our customers happy, we work closely with them to ensure that we account for changes in Information Assurance thinking and practice within the DoD. The "most cogent description of the issues, challenges and potential solutions"[11] were provided by Deputy Defense Secretary Lynn in an essay written for the journal *Foreign Affairs* in January 2010.[12] Within the essay, Secretary Lynn specifically posits that national cybersecurity defense "will only succeed if it is coordinated across the government, with allies, and with partners in the commercial sector." This particularly affected a proof-of-concept we built for a DoD agency where the goal was to demonstrate how Biometrics data could be passed securely and facilitate information sharing / rapid identification of suspect individuals. Additionally, this proof-of-concept demonstrated compliance to a highly-standardized development model based on the Army Data Services Layer (ADSL).[13] The anticipated benefit from this standardized approach was to "provide enterprise-accepted and trusted information for reuse, analysis, and aggregation,"[14] although it is too early to say whether the DoD will see the expected payoff.

Over the past year, we have built several projects that required server "stacks" (sets of servers that together provide complete application functionality). Each server within a stack must be capable of running in a secure environment and must be compliant to evolving Information Assurance practices within the DoD. For example, we use the Army Golden Master (AGM) as a starting point for each server that we build. The AGM changes over time; as Army security analysts identify additional hardening measures and patches, they must be applied to our systems via Information Assurance Vulnerability Alerts (IAVAs). Also, overall policies change over time: as an example, Army Regulation 25-2 published a "Rapid Action Revision" on March 23, 2009 that required us to modify the operating system sign-on banner to meet new specifications.[15] The anticipated benefit from this approach was to ensure that secure systems stay secure over time; however, as the USB Drive Security Breach of 2008[16] pointed out, a reactive and perimeter-heavy response to overall system security can fail.

3.2 Accounting for Growth

Within our industry, the biggest single change impacting our organization has been "cloud" computing. In this model, the DoD delivers application services not from local computer networks but from centrally-hosted data facilities making up a set of "globally interconnected, end-to-end set of information capabilities" – the Global Information Grid (GIG).[17] The effort is expected to result in massive savings both in time and money. For example, historically it has taken months to provision a secure application server; the Defense Information Systems Agency (DISA) can now use the Rapid Access Computing Environment (RACE) to create the same server in fifteen minutes.[18]

Our organization has been personally affected by this cloud initiative. From a revenue standpoint, over the past year we were able to provide support for the Army's Area Processing Center initiative championed by Army CIO LTG Sorenson to consolidate the current Network Enterprise Centers (formerly Directorates of Information Management, or DOIMs).[19] At the network infrastructure level, the drive to use the cloud has impacted us as well. One important step in moving to the cloud is to ensure that servers can run in a virtualized environment; Ang Li et. al. note in a 2010 paper that virtualization is a key requisite in "elastic computing" (the dynamic and efficient scaling of application resources based on user demand).[20] Over the past year we have purchased two powerful hypervisor servers to run virtual machines (VMs). Virtualization was one solution that allowed us to satisfy both our company growth while reducing our administration overhead costs. In a virtualized environment we can create new systems quickly based on templates; we have templates defined for our developers, our analysts, and even multiple types of servers (such as Web or database servers). Because these virtual machines are all controllable via a single management interface, we are able to create new VMs, start and stop existing VMs, and perform backup / restore functions easily and quickly. In our case, the anticipated benefits of centralized management and convenient server provisioning have been fully realized.

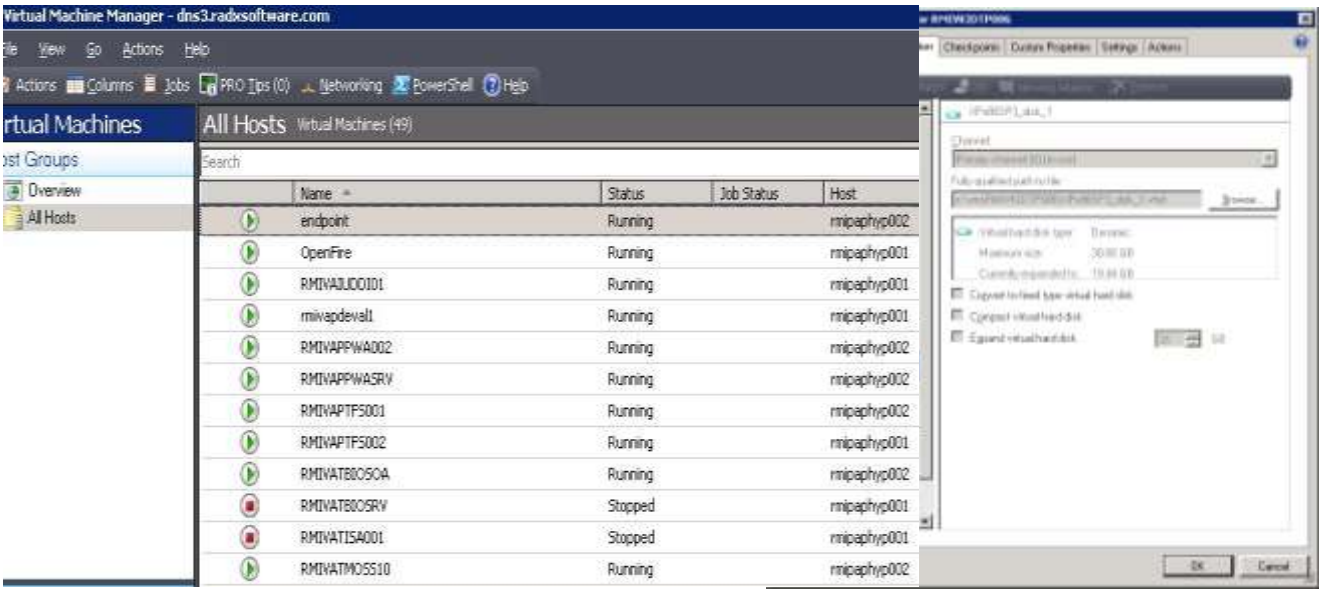

Figure 3 - A subset of installed virtual machines, and a specific example

3.3 Protecting Company Assets

Within the Federal Government, the past year has seen a significant focus on cybersecurity. President Obama's 2009 address on the issue stated that the "status quo is no longer acceptable" and that protecting network infrastructure "will be a national security priority."[21] This new security impetus has led to much new legislation, the latest being "The Protecting Cyberspace as a National Asset Act of 2010" by Senators Lieberman, Collins, and Carper.[22]

With the increased emphasis on security, we have had to implement a number of improvements based on requirements from our own customers. We have concentrated on applying these security improvements and building a solid network infrastructure that allows us to grow and ensures that we meet customer expectations. As one practical example, we have worked hard to address the issue of company-issued and employee-owned laptops all connected to a wireless access point:

1. *Security Policy* - We started by defining our overall corporate security policy consisting of our corporate positions on Acceptable Use, Antivirus, Backup, Password, Email, Encryption, Ethics, Privacy, Software Installation, Local / Remote Access, and Education / Training.
2. *Physical Inventory Management* – Our next step was to identify all company-owned equipment and ensure accountability. As part of this effort, we identified and removed all employee-owned equipment from our network.
3. *Logical Inventory Management* – On each of our systems, we looked at installed software packages and ensured that each package existed to support a business requirement. One useful byproduct of this effort was a set of templates for different types of machines (developers, analysts, executives, and so on).
4. *Security Domain* – We implemented a company-wide policy for all devices regarding software installation, anti-virus, sign-on banners, and network access. We used group membership to assign roles to each of our resources and servers and applied rules to each role. As a result, a software developer has required programming tools installed while an analyst might have adaptors for accessing DoD Web sites containing architectural documents.
5. *Server Area* – Our server infrastructure started life on long tables in a large "bullpen" area. We defined an area for these servers, replaced the carpeting with tile to reduce the risk of static electricity, purchased racks for holding our Uninterruptible Power Supplies (UPS devices) and our rack-mountable servers, paid electricians to rewire our area, setup a cabling patch panel that we could use to connect our phones and network cables to a single location, and then proceeded to build and document our improved server infrastructure. We found a number of items that we could correct, for example, we located and moved several network cables that crossed fluorescent lights within our dropped ceilings – by removing the radio frequency interference (RFI) we found that we had a measurable improvement in our network connectivity.

6. *Thin Client* – One area where we were not happy was in the need for resources to have a company-provided laptop simply to connect to the network. To eliminate that need, we used our hypervisors to create a pool of standardized virtual machines for our team members to use. Initially, we had each team member use their laptop to connect to and work on her assigned virtual machine. After this shakedown process completed, we provided a "thin client" device to the team member. This device contains neither hard disk nor CD-ROM bay and it allows connection to a keyboard, mouse, and monitor only. (USB disk storage devices can be disabled via administrative configuration.) The thin client has one function: to allow the team member to connect to their assigned virtual machine. As part of this effort, we also provided physical wiring to each work station to eliminate the need for open wireless access to our network. This major project helped ensure that our corporate assets (both physical and logical) were more secure.

3.4 Lessons Learned

Our primarily lesson learned can be summed up simply: security is not static. New vulnerabilities are being uncovered all the time. As a case in point, an August 2010 audit of the National Cyber Security Division (NCSD) reports that "a significant effort is needed to address existing security issues" and that "NCSD needs to focus on deploying timely system security patches."[23]

In our own organization, we have applied that lesson to ourselves. Over the past year we have implemented automatic detection software to alert system administrators when security patches are available for our critical infrastructure servers. For our end-user virtual systems we require that operating system patches are applied automatically. Such an automatic approach can be dangerous to use on a production system due to unanticipated side effects.[24] However, for us the advantages of an automatically-updated system far outweigh the risk of a particular operating system patch negatively impacting our users.

Our choice of thin client hardware contributed to our lessons learned. We started with a highly reputable vendor and one of its most popular models. Additionally, we brought in a sample unit for a full month to test it on our network in a controlled environment. All functions worked quite well for us and we gave the go-ahead to purchase it as a company standard, and that's when the trouble started. We found that some of our higher-end monitors did not work well with the thin client – the display could be fuzzy or simply wrong. Also, as part of our network requirements we wanted to move toward a smartcard environment where our users could login without requiring a user name and password. Our test thin client worked quite well in allowing network login using both a locally-created smartcard (one we provisioned ourselves) as well as with a DoD-issued Common Access Card (CAC).[25] However, some of our analysts had to use a CAC to logon to various secure DoD Web sites. Our thin clients did not perform well for us in that scenario; we had problems that required a good deal of time and effort to work around!

This thin client experience taught us to think beyond current needs and to anticipate the future. We knew that we needed to support DoD CACs, but we assumed that if our thin client solution supported this with a local network login then everything was "A-OK." That was a valuable lesson – *follow testing through to its logical conclusion*.

4.0 Future Business Needs of Our Network

We have analyzed our current network state and milestone events from the past twelve months. Now we analyze how our anticipated future business needs will affect demands upon our network:[26]

- *Growth* – We are actively pursuing additional work and submitting proposals. Our network must be capable of scaling out to handle this growth very quickly and reliably.
- *Physical Compliance* – Our customers require us to maintain a highly secure server environment.

In accounting for these future needs, we see three primary areas where we can plan for growth: our ability to handle blended threats;[*] our ability to prepare our network for new users and systems; and, our ability to maintain a secure facility for our network servers. Additionally we must be cognizant of cost by continuously focusing upon scalability and automation of our network infrastructure. Fortunately for our organization, our future business issues do not basically conflict with any information assurance issues.

4.1 Addressing Blended Threats

Blended threats are especially worrisome to us. As we host more users on our network and provide more Web sites for our customers, our network becomes a larger target for ever-mutating malware. We believe that the key to handling blended threats economically is to centralize our network security management.

4.1.1 The Integrated Threat Management (ITM) System

Our current network architecture consists of the following high-level layered approach:

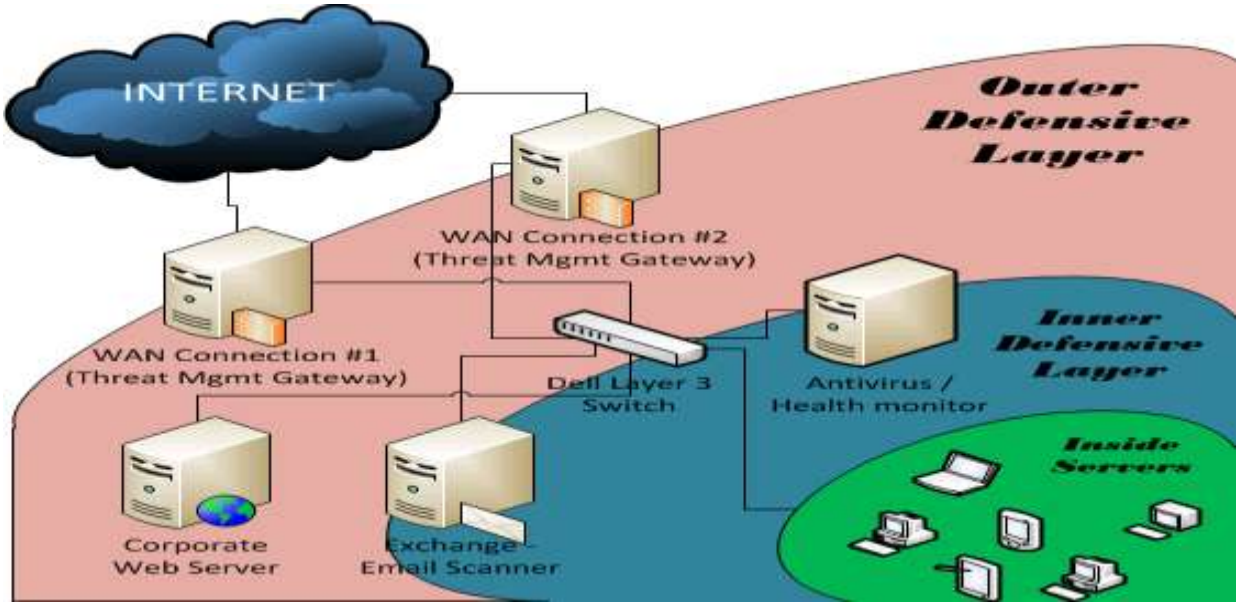

Figure 4 - Our current network layers

[*] Threats consisting of more than one classic attack vector, such as a piece of malware combining the characteristics of a worm with malicious code. Blended threats are especially harmful on the Internet.

According to McBride, ITM requires an **integrated** "platform to monitor and manage the ITM" as well as "components [that] have their life-cycle activities consolidated."[27] In other words: a centralized dashboard display with views to show the operator the current state of the environment along with alerts and required remedial action.

4.1.2 Edge Firewall and Blended Threat Protection

In our organization's case, we have dual firewalls – one for each of our external network connections. The firewall itself has a number of features enabled to provide true application-level (OSI Layer 7) protection, as shown below:

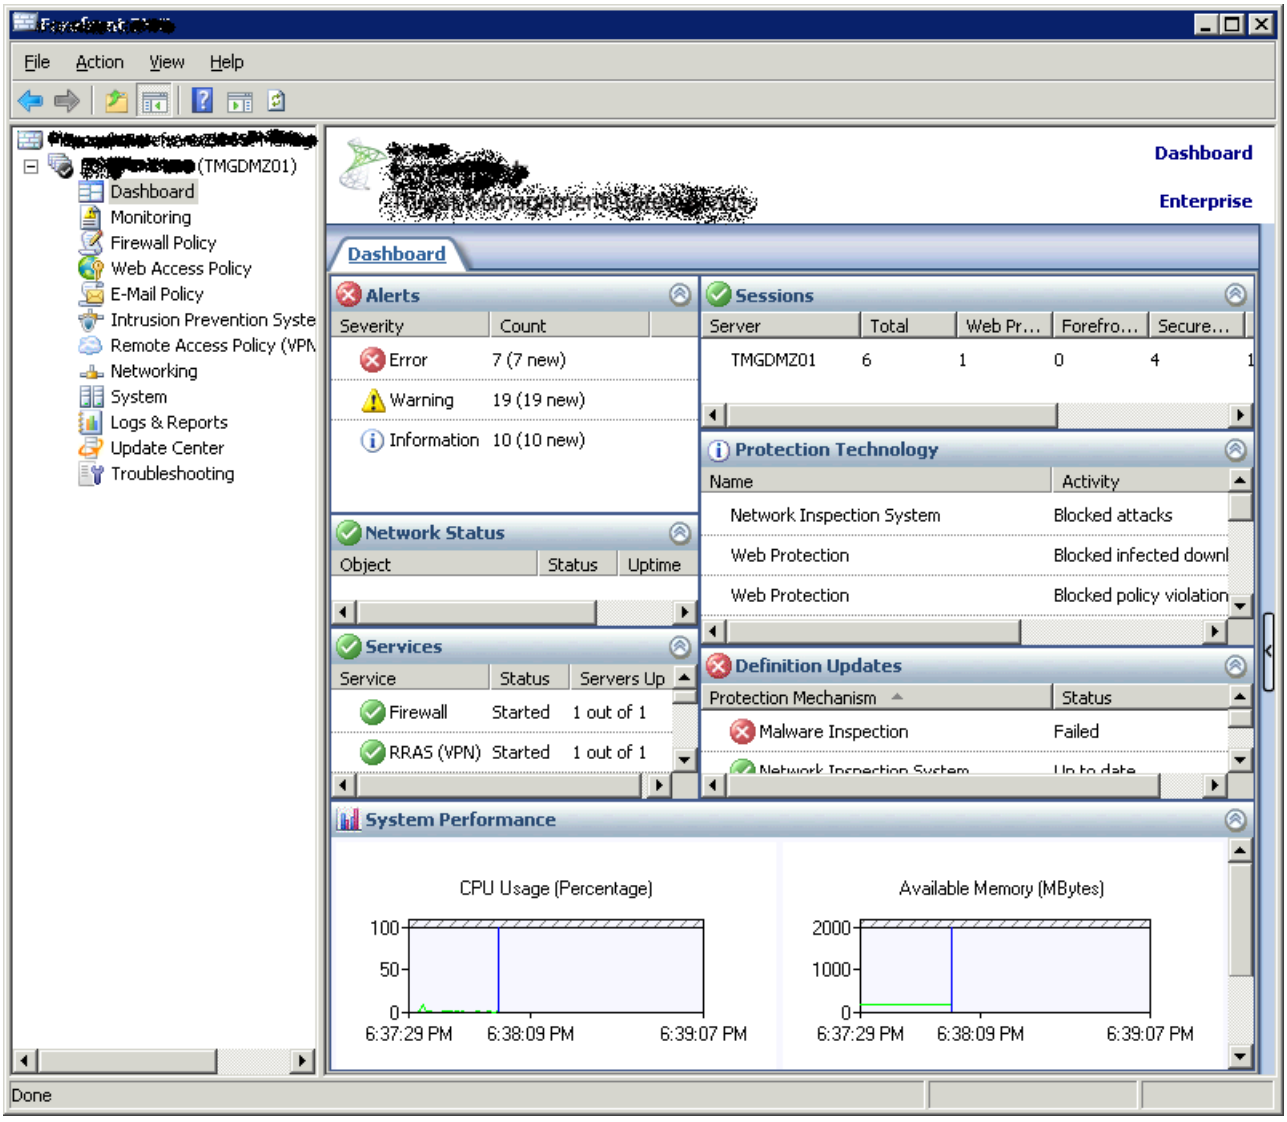

Figure 5 - An Integrated Threat Management Gateway

Our current tool satisfies key characteristics of ITM:

- integration of multiple types of tools (not just firewall);
- ability to scan traffic (including email traffic) and to perform network traffic analysis (intrusion prevention);
- support for Virtual Private Network connections, including access policy; and,
- automatic notification upon failures or errors.

Our solution's built-in scanner fully analyzes each network packet, tracking file downloads for each client and submitting these downloads to a full scan based on signatures updated automatically every fifteen minutes as shown below:

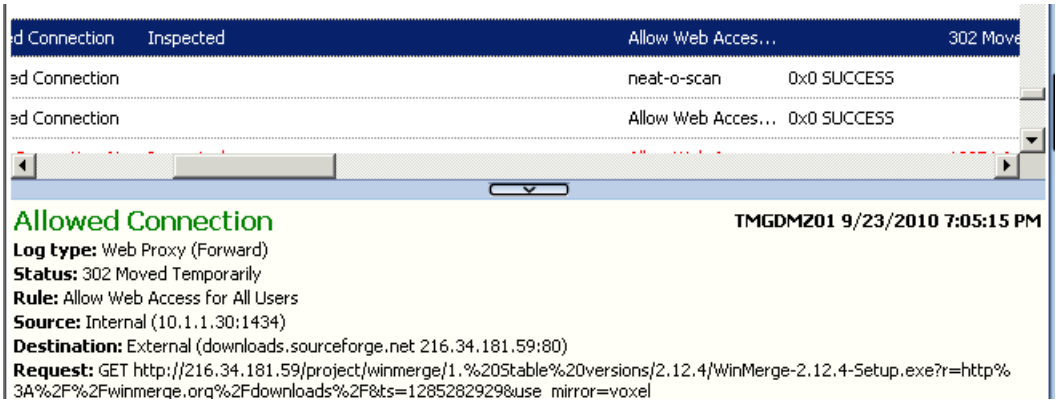

Figure 6 - ITM scanning Web content

4.1.3 A Missing Link – NIPS

While our current blended threat protection is good, we suffer from a lack of a Network Intrusion *Prevention* System. Our existing protection systems analyze data and events as they occur, and our firewall applies detection rules to prevent denial of service attacks, but we have no centralized method on our network that will proactively identify traffic and / or activity that indicates an attack may be imminent or that other suspicious activities are occurring. Even when a system under analysis is simply "misbehaving"(such as a compromised end-user workstation where a zero-day worm has disabled the local antivirus client and is busily engaged in fingerprinting all adjacent workstations in preparation for replication), consider the possibilities: vast amounts of data can be dumped to the network (consuming bandwidth and requiring immediate remediation from IT staff), or disk files could fill up, or user keystrokes could be lost or corrupted. While a NIPS would not address *all* of these issues, it would certainly help to identify those issues that involve the network. Not having a NIPS is a real problem for us that we address in our recommendations.

4.2 Grooming for Growth via Virtualization

Growth is a significant factor affecting our network over the next eighteen months. The proposals we have already submitted, combined with the proposals that are in our pipeline, make it critical that we can provision and support desktop environments as well as development and testing servers in a timely fashion. (Most task orders require the vendor to be fully staffed and running on the awarded work within a thirty day window.) In this environment we can expect that the increased network traffic can test the limits of the physical equipment that we have installed, which can result in a negative impact on our network's overall availability – a key element of our Information Assurance posture.

4.3.1 Virtualization and…Network Cards??

One area where we may have a problem is in the way we have implemented network communications support within our physical hypervisor. Each hypervisor communicates with the physical network by using Network Interface Cards (NICs). In a highly virtualized environment such as ours we have a problem specific to our NICs – namely, that of overall network throughput. As can be seen in the drawing, the hypervisor contains a number of virtual machines (VMs); each one of these VMs is a complete and running system and therefore has at least one network connection ("virtual NIC"). The hypervisor itself manages all of these virtual NICs and routes all network communications (inbound and outbound) through its physical NIC.

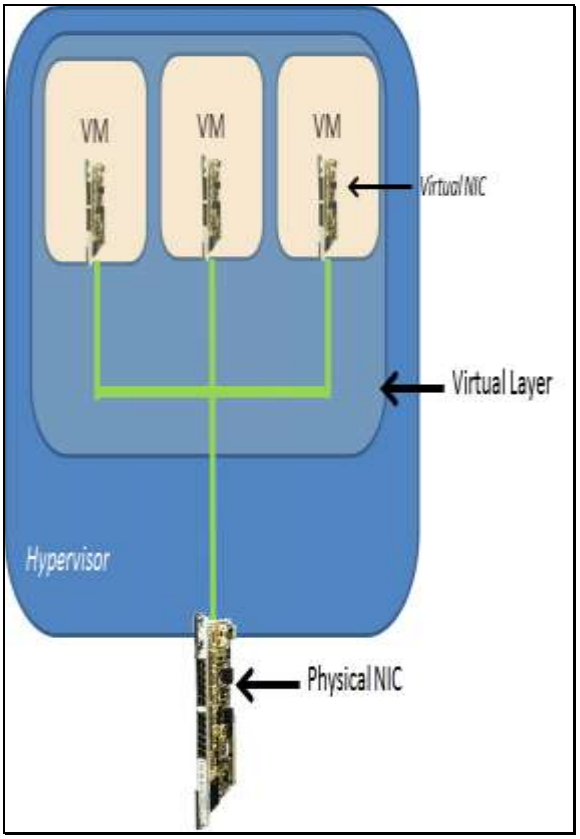

Illustration 1: Single NIC for a hypervisor

Problems with this approach include:

- *All network traffic for all virtual machines hosted on the hypervisor is serialized through the single physical network card.* Consider a common case where the physical network interface card uses "Fast" Ethernet (100Mbps). In this situation, having multiple VMs all performing large numbers of network operations (such as a Web server or a file sharing computer), the sum total of network traffic from all VMs can overwhelm the capabilities of the single card. (This can also occur if the hypervisor has a Gigabit Ethernet 1000Mbps NIC, although is more rare.) We can expect this to become a problem for us over the next eighteen months as we increase staff.
- *Failures to the single network card effectively bring down the entire set of virtual machines hosted on the hypervisor.* While NICs are extremely reliable, failures certainly do occur. In a situation where the hypervisor is hosting a number of dynamically allocated VMs that are providing functions for a paying customer, this type of failure risk is unacceptable.

4.3.2 Virtualization and Disk Storage

In order for virtualization to work, our hypervisors must be able to store and load virtual machines (VMs) to perform required functions. (As an example, each software developer is assigned a specific VM for her daily work.) Currently, we have sufficient high-quality and fault-tolerant storage not only for our current VMs, but for many more. However, we have two problems with this approach that we need to address:

- What happens when we run out of space?
- What happens when we need to assign a VM to a different hypervisor (for example, adding a new hypervisor)?

The two questions are closely related, although perhaps not at first glance. The problem lies in how we have our hypervisors physically configured: they each have their own disk storage. In our case this means that when we have a specific VM that we want to run, we must place that VM on a specific hypervisor. To add a new VM to an existing hypervisor, it must have sufficient disk space available. Adding a new hypervisor into the mix (for example, to reallocate disk resource usage from an existing hypervisors) requires us to physically transfer any existing VMs to that new hypervisor.

The real problem is that we are missing *common storage* for each of our hypervisors to use. In our recommendations below, we address this problem by means of a Storage Access Network device.

4.3 Physical Server Security

In our organization's network, our server setup reflects our start-up past. We have a single dedicated area in our development bullpen where we maintain our locked server racks as shown in the figure.

The problems we face with this setup are many. First, it eliminates any option of being qualified as a secure computing facility (a critical component of our future business model). Also, as we hire more staff we must bear in mind that we have a significant risk in having the servers physically near the development team. Finally, keeping our servers located in our facility does not provide that level of assurance and trust which our partners will want to see.

We are currently defining a long-term home not just for our current servers but for additional equipment we have already planned on ordering for new work. As part of this decision, we are accounting for Army and DoD requirements and ensure that these new servers (and their long-term home) meet or exceed information assurance needs as well as network management issues.

5.0 Recommendations

We aim our recommendations at improving our current level of information assurance in our network's organization and preparing for changes currently in the planning stages. Each recommendation clearly supports our overarching goals of increasing our network's scalability while containing costs through the effective use of automation techniques.

5.1 Redundant NICs in Hypervisors

In our future business needs, we discovered that our expected company growth over the next year will add significantly to our network demands. One significant risk is the loss of availability due to a failed NIC within one of our hypervisors or the simple overload of the NIC card's capacity. **We recommend using NIC Teaming** to address both scalability (increased throughput) and availability (redundancy). This low-risk and inexpensive solution should be implemented as soon as possible. (For details on NIC Teaming and the network card problem with our hypervisor, see 4.3.1 above.)

5.1.1 Benefits and Downsides

Implementing NIC teaming ensures that if a network card fails, the hypervisor does not lose network connectivity. NIC teaming works only with Microsoft Hyper-V on Server 2008 R2 – but this is our standard hypervisor installation. As a downside, only a few NICs have shown themselves to be reliable, limiting our choices.[*]

5.1.2 Additional Resources

We require our Chief Operations Officer to approve this purchase.

5.1.3 Approximate Costs

A server-class Intel PRO/1000 MT Dual Port Server NIC retails on Google Shopping for $177.00,[**] making the total expected cost for five hypervisors less than $1,000.00. (Five hypervisors are sufficient for our growth over the next 12 months.) Network installation is somewhat tricky, as it requires the hypervisor to be configured with static MAC addresses (similar to static IP addresses but at OSI Layer 2). Total time involved should be a single weekend as this installation effort is well-documented on the Web. As an example, Joachim Nässlander has a complete writeup on implementing NIC teaming for our environment (Microsoft Hyper-V in Server Core mode) on his technically-acclaimed Web site[***] (although numerous other excellent write-ups exist).

---

[*] "Broadcom NIC Teaming and Hyper-V on Server 2008 R2," *confused?!amused*, http://tinyurl.com/323oub8 (accessed: November 15, 2010).

[**] Verified on November 1, 2010 using Google Shopping search and the first listed offer from CompUPlus.com (http://www.compuplus.com/Network/Intel-PRO1000-MT-Network-Adapter-1019830.html).

[***] Joachim Nässlander, "NIC Teaming in Server Core or Hyper-V Server," *nullsession.com*, http://www.nullsession.com/2010/08/15/nic-teaming-in-server-core-or-hyper-v-server/ (accessed: November 12, 2010).

5.2 Implementing Hardware NIPS

In our Blended Threats section, we identified a missing component to our defense-in-depth: Network Intrusion *Prevention*. **We recommend the purchase of a dedicated NIPS appliance** to prepare for the increased team members and the added attack surface these team members represent.

5.2.1 Benefits and Downsides

While Intrusion *Detection* Systems work by means of known attack signatures, NIPS works by having a highly secure device listening in promiscuous mode on the LAN and analyzing traffic patterns for anomalies that may indicate a security event is taking place. A NIPS appliance is largely self-maintaining and requires little configuration; it can be integrated with our existing email and phone services to send automated alerts upon detecting a condition. Although previously maligned for issuing excessive numbers of false positives, latest-generation NIPS devices are quite reliable in this respect. Additionally, the NIPS device provides true defense-in-depth to our network's security posture as it is completely independent of and works in conjunction with our existing solutions.

There are downsides to this purchase. The value of NIPS lies in its ability to detect unusual network activity, which means that we must allocate a "learning" period for the NIPS device to detect normal network traffic. Additionally, standard maintenance windows where we either take servers offline or perform load testing on the network will generate false alarms unless the NIPS is configured to ignore these situations. Finally, as threats and attack patterns evolve, the NIPS software itself must be updated. This adds a yearly maintenance expense for a support contract (which can be expensive). For this recommendation, we are monitoring VM communications using the standard appliance and not via special software running within the hypervisor stack.

5.2.2 Additional Resources

Our Chief Operating Officer must approve this purchase as well as the annual maintenance cost. Having a support contract may mean that we are locked in to the vendor we select, so there could be a need to run the support contract past our retained legal counsel. Also, we need an opinion on our legal right to sniff, decode, and analyze all traffic on our network (including guests using this network). While our existing corporate security policy and login banners may suffice for this level of control, we may need to modify our verbiage or exclude one or more subnets of traffic from analysis.

5.2.3 Approximate Costs

We analyzed the Juniper IDP75 appliance, the Cisco ASA 5520, and the 3COM TippingPoint 50. All three products have excellent ratings and similar throughput; all are designed for small businesses and meet our needs for the near future. The difference becomes cost.

Product	Cost*	Support / Subscription Annual Maintenance	Total (Year 1)	Total (Year 5)
Juniper IDP75	$5,440.00	$1,171.20	$6,611.20	$11,296.00
Cisco ASA 5520	$5,868.00	$3,461.00	$9,329.00	$23,173.00
3COM TippingPoint 50	$4,381.00	$2,000.00**	$6,381.00	$14,381.00

Based on the above, we **recommend the Juniper IDP75** for our current needs.

Each of the systems we investigated provided free online training courses along with paid classes; the average length of time recommended for completing the courses is at least three days. The selected Juniper IDP75 model also offers a paid three-day training course starting from $2,500.00 (making it still the least expensive solution).***

5.3 Implementing Storage Access Network (SAN)

Each of our current hypervisors has its own high-performance and fault-tolerant disk storage for managing the virtual machines (VMs) stored on it. However, our need to account for massive growth in conjunction with our requirement for centralized management makes our current solution untenable. **We recommend the purchase of at least a 12 terabyte (TB) SAN** so that we may store all of our virtual machines on a high-speed and shared drive.

5.3.1 Benefits and Downsides

Implementing a SAN has very few downsides apart from cost. The SAN itself can be partitioned in numerous ways such that we can use one data partition for our end users, another partition for the shared storage between our hypervisors, and yet another partition for highly confidential information. The SANs we have examined all feature highly secure data partitioning schemes that can be configured to disallow access from entire subnets of data consumers – thus, an end-user with access to the SAN for storing documents simply has no way to access any other protected partition. Additionally, SANs feature pluggable and fault-tolerant disk arrays to which we can add new disks dynamically in order to extend a data partition – without taking either the SAN or any data consumer offline.

---

* Verified from Google Products search as of October 27, 2010.

** 2007 numbers from PC World based on subscriptions to Digital Vaccine and Web content filtering (http://www.pcworld.com/businesscenter/article/132071/3com_network_security_gear_for_small_business.html).

*** Verified from Juniper Web site as of October 28, 2010 (http://www.juniper.net/us/en/training/technical_education/courses/EDU-JUN-IIDP.htm).

5.3.2 Additional Resources

Our Chief Operating Officer must approve this expense due to the significant cost of the SAN.

5.3.3 Approximate Costs

In determining a specific SAN solution, the first question to answer is whether to use *iSCSI* or *Fibre*. The iSCSI solution uses a standard network cable (RJ-45) to transfer data between the SAN and the attached servers, while the Fibre solution uses fiber-optic cable. Fibre is significantly faster than iSCSI (at the low end, compare 4Gbps for Fibre to 1Gbps for iSCSI), but it is possible to use NIC teaming to increase the overall iSCSI output (four teamed NICs provide close to the same throughput).[*] Additionally, using Fibre requires the purchase of a much more expensive network switch to connect the SAN and the consuming servers and significant staff training. For these reasons, we recommend an iSCSI solution.

We analyzed the Dell EqualLogic PS4000E, the IBM DS3400 SAS Dual-Controller, and the HP StorageWorks P4300 G2 SAS. All are rated as basic SANs having consistently high marks. Our configuration assumes highly-available RAID 1 drive arrays running with SATA drives at 7200RPM; however, the SANs listed below allow SAS drives at 15000RPM to be substituted. The SAS drives are significantly more expensive (almost twice the cost) and hold much less space per disk (max of 600GB for SAS compared to 2TB for SATA).

Product	**Support Package (3-year)**	**Cost (4TB)**	**Cost (8TB)**	**Cost (16TB)**
Dell EqualLogic PS4000E[**]	$2,000.00	$20,293.00	$28,293.00	$42,293.00
IBM DS3400 SAS Dual-Controller[***]	$1,926.00	$15,340.00	$20,456.00	$25,572.00 (12TB max)
HP StorageWorks P4300 G2 SAS[****]	[included]	[n/a]	$26,940.00	[n/a]

Based on the above, we **recommend the IBM DS3400** with 8TB of high-availability disk space. We can expand this space to the 12TB max simply by adding 2TB SATA disks.

Personnel support costs are limited to a required training program for at least two network engineers and a single weekend to install and test the SAN itself. The training program costs differ between vendors; for the IBM DS3400 there is no charge for the online lessons but they do require several days to complete.[*****]

[*] Marc Stanner, "Choosing a SAN technology: Fibre Channel vs. iSCSI," *SMB Storage Tips*, http://searchsmbstorage.techtarget.com/tip/0,289483,sid188_gci1516944,00.html (accessed: November 3, 2010).

[**] Source is Dell Web site, November 3, 2010.

[***] Source is IBM Web site, November 4, 2010.

[****] Source is Google products, November 18, 2010.

[*****] Source is IBM Web site: http://ibmdsseriestraining.com/ds3000/installation-a-configuration/ds3400-ds3300-a-ds3200 (last accessed November 4, 2010).

5.4 Securing our Physical Servers

Our current environment reflects our company's start-up -- we have concentrated on building software solutions using a small team of trusted individuals. Beyond hardening the network to implement least privilege, we have not had to worry about physical access to our physical server stack itself. Upcoming proposals require us to address this deficiency to qualify for desired task orders.

Arguably, properly securing the physical servers can be considered the most critical action to be performed. However, the significant expense of implementing this approach makes it a decision to be made carefully. Nevertheless, **we recommend storing our physical server stack in a secure location.** *Secure* in this sense means that we are *on the way* to being fully compliant to DoD standards, not a true Secure Computing Facility (SCF) as currently defined by DoD standards.

5.4.1 Benefits and Downsides

The first question to decide is whether a secure facility makes sense at all – given the current business climate, would it not make more sense to embrace utility computing and to use a service provider by "colocating" our existing servers and equipment in an already-built, hardened, and available location? The advantages to using a colocation facility are many – we have no permits to apply for, we don't need to engage contractors to build our server room, and we don't need to worry about running out of space (we just rent more space from the facility). However, our choices for colocation do not stop there; as we see from the *Colocation Hosting* Web portal[*] there are a number of questions to review when selecting a colocation provider.

5.4.1.1 Traditional or Managed?

The first question to research is just what type of colocation makes most sense for our organization. Colocation has traditionally been a service where an organization ships existing server hardware to a third-party location. The third-party provides the network connections, power supply, guaranteed uptime, and remote access necessary to manage the shipped servers (*colocation service provider*). In this model, our organization is still fully responsible for managing the servers; the colocation service provider normally simply provisions new servers, attaches them to the network, and arrange for remote access. Any additional services (troubleshooting, hardware replacements, etc.) are managed based on the type of service provider and the colocation plan our organization has in effect.

A managed colocation choice, on the other hand, offers a utility-based model rather than the traditional choice.[**] In the managed model, we don't send physical servers to the service provider. Instead, we are charged for the amount of resources we use – disk, central processor units or CPUs, memory, and network bandwidth. Our negotiated subscription plan generally establishes a minimum fee for these resources, and any usage beyond that is charged at a predetermined rate. This allows us to increase our number of Web servers, for example, simply by requesting and provisioning additional virtual servers from the managed colocation service

---

[*] ColocationHosting.org, "Your #1 Source For Colocation Hosting Information," *ColocationHosting*, http://www.colocation-hosting.org/ (accessed: October 18, 2010). This not-for-profit Web site provides basic information on the key decisions to make when selecting a colocation hosting service plan or provider.

[**] Rackspace, "Managed Colocation Versus Traditional Colocation," *rackspaceHosting*, http://www.rackspace.com/managed_hosting/managed_colocation/comparison.php (accessed: October 18, 2010). This resource, while commercial in nature, highlights the major differentiators between managed colocation and traditional colocation options.

provider. Examples of companies providing this type of service include Amazon's S3 service and Microsoft Azure.

5.4.1.2 Reliable Colocation Service Provider

Regardless of our choice between a traditional or a managed colocation service provider, we have a number of common questions that must be answered, including:

1. Is the service provider reputable? Are there reference customers?
2. Does the service provider provide a backup policy as well as a disaster recovery plan?
3. What type of power and cooling capabilities does the service provider have? Is power limited to the city grid? Are multiple substations involved in providing power to the facility?
4. For traditional colocation providers in the same city, what are the terms for physical access to the facility? Is there an additional charge for after-hours access? What controls are in place to protect our organization's hosted servers and equipment from other subscribers?
5. For managed colocation providers, does the service provider have alternate sites that can be leveraged and provide assurance of both Recovery Time Objective as well as Recovery Point Objective?

5.4.1.3 Physical Server Room: Our Recommendation Now

While managed colocation certainly has numerous attractive features, **for the present we recommend using a physical server room**. Our future business model assumes our ability to provide a true stand-alone secure computing facility, properly hardened and certified to DoD standards. Moving forward, we may certainly outsource our corporate computing needs by using a managed colocation provider. However, we will also certainly still have a need for our own secure facility; over the next two years it will be less expensive to leverage the physical servers we know we require rather than to have both physical servers and the managed colocation facility. In a future paper, we will delve into the colocation choice more deeply.

For our secure physical location, we still have one more choice to make: renting space in an existing hosting facility or building our own space using an appropriate room in our current location. (We do not include Web hosting options here, as our organization's network environment includes substantially more than a simple Web server.) While space considerations preclude us from analyzing both options in this paper and recommending one option over the other, we propose addressing this issue in a future paper.

5.4.2 Additional Resources

Regardless of whether we use a hosting facility or build a secure server room, we will need to include our Contracts individual as we will either be entering into a new long-term agreement either with our landlord or with the hosting facility. Our Chief Operations Officer must approve the final agreement and the plan specifications, and our CEO has the final say on our company’s approach.

To build a server room we must request the appropriate county permits and ensure that we meet minimum requirements (fire-ratings for walls, ceiling and floor clearances, fire suppression systems, HVAC capabilities, and so on). To make the room secure we must arrange for physical security (such as a card reader) and to ensure that our solution provides an accounting capability to control entrance. We'll also need to make sure that we select an interior room (no windows) and that the room has full walls (true floor to true ceiling) and not just floor to dropped ceiling tiles as is normally the case for interior commercial office spaces.

5.4.3 Approximate Costs

In this paper, we look only at estimated costs for building a secure server room. A future paper will look at hosting costs.

Item	Estimated Cost	Time	Notes
Site Location (assumes using space in an existing local facility)	[n/a]	5 days	Phone calls, trips to local facility
Physical Office Space	$10.00 / SF* * (25x20) = $5,000 / year	30 days	Subject to receiving valid permitting, landlord negotiations, and contractual requirements.
Construction	$25,000.00**	45 days	Based on: walls rated two-hour fire resistance, wiring for server racks, HVAC, and required permits
Verification	$2,500***	5 days	Our network engineer ensures that we have valid connections and wiring to the facility, and that we pre-install necessary switches and routers between our home off. Pricing is for two Cisco 24-port SGE2000 switches ($700) along with two Cisco RVS4000 routers ($130) and miscellaneous network costs (cabling, etc.)
Moving	~ $1,000.00	5 days	One weekend to perform the physical move, several days to verify correct operations.

* Cushman & Wakefield, "Northern Virginia Industrial Report," *Market Beat 1Q10* (McLean, VA: 2010). Available online at http://www.cushwake.com/cwmbs1q10/PDF/ind_northernva_1q10.pdf (accessed: October 20, 2010). The square footage price quoted is an approximate average of commercial real-estate prices for the Manassas, VA area.

** Response from Chief Operations Officer based on quotes for a project similar to this effort.

*** Source is *newegg.com*, validated November 1, 2010.

6.0 Conclusion

"If you think technology can solve your security problems, then you don't understand the problems and you don't understand the technology." -- Bruce Schneier

6.1 Final Thoughts

Our organization's network reflects the tremendous forward thinking of our executive staff and the willingness of our Team to put our "money where our mouth is." In a difficult economic climate and while building a solid business practice including new business development, we have implemented numerous enhancements to our network to improve our overall Information Assurance and to provide a high-quality work environment to our team members. We have responded to Government's and DoD's call for improved operating efficiency *and* improved overall security with a series of cost-effective and centralizing implementation decisions that allow us to meet our increased customer demands and obligations without breaking the bank or imposing onerous ongoing maintenance costs on our organization.

However, we realize that we can and must not rest content on our laurels. The one truism in Information Assurance is the assurance that the level playing field is *always changing* – whether from newly attack vectors (vulnerabilities), new attack patterns (such as emerging zero-day blended threats targeting "soft" areas like *external* employee usage of social networking), or new regulations (such as the The Protecting Cyberspace as a National Asset Act of 2010). We as security practitioners must always maintain our vigilance and our commitment to providing a secure environment that provides above all else a cost-effective balance between requirements and implementation. The job of security is not to cost money, but to ensure that the maximum amount of money can be made in a reliable, transparent, and auditable fashion.

6.2 The Bottom Line

In this paper, we have analyzed how our particular organization's network came into being. We have analyzed how Government and technology changes have personally impacted us. We have seen how important it is not only to test technology – but to test it beyond your initial requirements and to think outside of the box. We have looked at specific new technologies that can aid us as we move forward as an organization and we have taken a SWAG at how these technologies can be implemented and maintained in a cost-efficient manner.

These are all good things. But they miss the real point of security…security is not about what your machines can prevent and mediate, it's about what you as a *whole organization* can (to paraphrase Bruce Schneier once more) "detect and respond." Education is more important than technology just as people are more important than machines: a single poorly inserted USB device can bypass any network firewall in existence.

We close this paper by reminding our readers that we need to bear in mind the importance of providing a functioning and functional environment for the end-users, customers, and auditors. As we move to the future, and learn from the past, we need to remember that – given enough time -- attackers can always overcome any defenses. Our job as security practitioners is to ensure that we see this risk clearly, accept the fact that it is not a question of *if* but of *when* breaches occur, and that we have a documented and proven plan in place to handle the expected impact of a data breach on our system. It is only thus that we truly provide the assurance that our principals and users need in order to function effectively in our always-on and ever-more-wired world.

Appendix: Recommendations by Priority

In this section, we provide a brief overview of our recommendations. For details on our recommendations, please refer to Section 5 above.

Item	Overview	Reasoning	Costs
#1: Securing our Physical Servers	Our servers should be moved to a specially-built location within our organization. This location should have strong physical security controls and should allow easy remote administration capabilities. This is most important due to the importance of our physical servers to our ability to perform our business functions.	Implementing this recommendation supports both our short-term concern of ensuring that our servers are physically protected, and our long-term goal to provide a secure computing facility meeting DoD standards.	Personnel: At least ten days of network engineer time. Also requires approval from our CFO, time from our COO, and may require advice from our Legal counsel. Monetary: $33,000.00 (first year), $5,000.00 (additional leasing cost per year)
#2: Redundant Network Interface Cards (NICs)	Our hypervisors are at risk of failing to scale to expanded network load due to the single physical network connection shared between all hosted virtual machines (VMs).	The low cost and complexity of this action make this an easy one to apply.	Personnel: At least three days of research from our network engineer. Monetary: Less than $800.00.
#3: Implement a Storage Access Network (SAN)	Currently we cannot automate dynamic VM hosting within our hypervisors; a given virtual machine is tightly bound to a specific hypervisor. Having all hypervisors work from a SAN allows us to load-balance our VMs based on demand.	As we grow over the next eighteen months, we will be adding numerous VMs and hypervisors. We need to plan ahead to ensure that we can manage our network infrastructure with minimal downtime.	Personnel: At least five days of network engineer time (training / setup). Requires authorization from our CTO and acceptance by our COO. Monetary: $25,572.00
#4: Implement Hardware Network Intrusion Prevention System (NIPS)	Our network environment does not have a centralized NIPS and therefore cannot detect unusual network traffic patterns. These network traffic patterns might indicate that an unidentified (zero-day) attack is in place.	We are well-protected on our network perimeter with strong defense-in-depth (network access rules, anti-virus, anti-spam, and more). However, an infected internally connected device can bypass the edge protections we have built.	Personnel: At least three days of network engineer training time, but minimal maintenance cost after that (the system is advertised as “self-maintaining”). Cost: $11,296.00 without paid training (add $2,500.00 for the paid training – we recommend this highly).

Reference List and Endnotes

End-notes immediately follow this section.

Bosworth, Seymour, M.E. Kabay, Eric Whyne, eds., *Computer Security Handbook: Volume 1, 4th ed.* Hoboken, NJ: John Wiley & Sons, Inc., 2009.

Bush, George W. *The Global War on Terrorism: The First 100 Days*. Washington: The Coalition Information Center, December 20, 2001. Available online at http://ics.leeds.ac.uk/papers/pmt/exhibits/330/GWOT100.pdf (accessed: November 10, 2010).

Centers for Medicare & Medicaid Services. *CMS Policy for the Acceptable Use of CMS Desktop/Laptop and Other IT Resources*. December 8, 2008. Available online at http://csrc.nist.gov/groups/SMA/fasp/documents/policy_procedure/Policy_Desktop_Laptop_Resources.doc (accessed: November 1, 2010).

Charette, Robert. "DoD Confirms Flash Drive Breached its IT Security in 2008," *ieeespectrum riskfactor blog*. August 25, 2010. http://spectrum.ieee.org/riskfactor/computing/it/dod-confirms-flash-drive-breached-its-it-security-in-2008 (accessed: October 28, 2010).

Commonwealth of Massachusetts. "Standards for the Protection of Personal Information of Residents of the Commonwealth." *201 CMR 17.00*. Springfield: Commonwealth of Massachusetts, 2010.

Deffer, Frank W. *DHS Needs to Improve the Security Posture of Its Cybersecurity Program Systems [OIG-10-111]*. Washington: Department of Homeland Security, August, 2010. Available online at http://www.dhs.gov/xoig/assets/mgmtrpts/OIG_10-111_Aug10.pdf (accessed: October 31, 2010).

Department of Defense. "Joint System Administrator Checklist." *Information Assurance Support Environment*. Defense Information Systems Agency, December 22, 2005. Available online at http://iase.disa.mil/stigs/checklist/joint-system-administrator-checklist-dec05.doc (accessed: November 14, 2010).

Department of the Army. "Authoritative Data Source (ADS)." *Army Net-Centric Data Strategy*. http://data.army.mil (accessed: November 7, 2010).

Department of the Army. "Frequently Asked Questions about Recruiting." *Support Army Recruiting.* http://www.2k.army.mil/faqs.htm#age (accessed: November 6, 2010).

Department of the Army. "Information Assurance: Rapid Action Revision (RAR) Issue Date: 23 March 2009." *Army Regulation 25-2*. Washington: October 24, 2007. Available online at http://www.army.mil/usapa/epubs/pdf/r25_2.pdf (accessed: November 8, 2010).

Deputy Under Secretary of Defense for Counterintelligence and Security. "Operating Manual Supplement." *National Industrial Security Program [DoD 5220.22-M-Sup 1]*. Available online at http://www.fas.org/sgp/library/nispom_dod_overprint_rev1.pdf (accessed: November 5, 2010).

Gallagher, Sean. "Army confronts battle to globalize its network resources." *defensesystems*. April 7, 2010. Available online at http://defensesystems.com/articles/2010/04/06/defense-it-1-gnecs-uphill-climb.aspx (accessed: October 29, 2010).

Howard, Rick. "The US DOD Proposes their Cyber Security Plan." *Notes from the Cyber Trenches*. September 9, 2010. Available online at http://blogs.verisign.com/idefense/2010/09/the-us-dod-proposes-their-cyber-security-plan.html (accessed: November 1, 2010).

Intersociety Working Group. "Research and Development FY 2008." *AAAS Report XXXII*. Washington: The American Association for the Advancement of Science, 2007.

Johnson, Nicole Blake. "OMB identifies 26 IT projects for closer scrutiny," *FederalTimes.com*, August 23, 2010. http://www.federaltimes.com/article/20100823/IT04/8230301/ (accessed: November 8, 2010).

Klindt, Todd. "KB973917 breaks SharePoint to pieces." *Todd Klindt's SharePoint Admin Blog*. January 4, 2010. http://www.toddklindt.com/blog/Lists/Posts/Post.aspx?ID=178 (accessed: November 1, 2010).

Li, Ang, Xiaowei Yang, Srikanth Kandula, Ming Zhang. "CloudCmp: Shopping for a Cloud Made Easy." *USENIX Workshop on Hot Topics in Cloud Computing (HotCloud) 2010*. Available online at http://research.microsoft.com/en-us/people/mzh/hotcloud10-cloudcmp.pdf (accessed: November 17, 2010).

Lynn, William J., III. "Defending a New Domain: The Pentagon's Cyberstrategy," *Foreign Affairs*. September / October 2010. Available online at http://www.foreignaffairs.com/articles/66552/william-j-lynn-iii/defending-a-new-domain?page=show (free registration required, accessed November 3, 2010).

National Security Agency. "Global Information Grid." *NSA Web site*. http://www.nsa.gov/ia/programs/global_industry_grid/index.shtml (accessed: October 29, 2010).

Obama, Barack. "Remarks by the President on Securing Our Nation's Cyber Infrastructure." Speech. East Room of the White House. Washington, DC. May 29, 2009.

Phillips, Leslie. "Lieberman, Collins, Carper Unveil Major Cybersecurity Bill to Modernize, Strengthen, and Coordinate Cyber Defenses." *Senate Committee on Homeland Security & Governmental Affairs*. June 10, 2010. available online at http://hsgac.senate.gov/public/index.cfm?FuseAction=Press.MajorityNews&ContentRecord_id=227d9e1e-5056-8059-765f-2239d301fb7f (accessed: November 7, 2010).

St. John, Alison. "Cuts In Department of Defense Spending Could Hit San Diego's Economy." *KPBS*, August 10, 2010. http://www.kpbs.org/news/2010/aug/10/cuts-department-defense-spending-will-hit-san-dieg/ (accessed: October 31, 2010).

Stephenson, Peter R., ed. *Information Security Essentials* (Section 1 and Section 2). Auerbach Publishing, ISBN 978-1-4398-0030-0, 2009.

Suss, Warren. "5 lessons from DOD's cloud computing efforts." *Government Computer News*. September 23, 2010. Available online at http://gcn.com/Articles/2009/09/28/Warren-Suss-5-lessons-of-DOD-cloud-computing.aspx (accessed: October 29, 2010).

Wilson, Scott and Al Kamen. "'Global War On Terror' Is Given New Name." *The Washington Post*, March 25, 2009. Published online at http://www.washingtonpost.com/wp-dyn/content/article/2009/03/24/AR2009032402818.html (accessed: November 18, 2010).

---

[1] George W. Bush, *The Global War on Terrorism: The First 100 Days* (Washington: The Coalition Information Center, December 20, 2001). Available online at http://ics.leeds.ac.uk/papers/pmt/exhibits/330/GWOT100.pdf (accessed: November 10, 2010). This document officially used the term "Global War on Terrorism" to describe the attack on the World Trade Center made by Al-Quaeda on September 11, 2001.

[2] Scott Wilson and Al Kamen, "'Global War On Terror' Is Given New Name," *The Washington Post*, March 25, 2009. Published online at http://www.washingtonpost.com/wp-dyn/content/article/2009/03/24/AR2009032402818.html

(accessed: November 18, 2010). The Obama administration agreed with analysis that the GWOT phrase "mischaracterized the nature of the enemy and it abilities."

[3] Intersociety Working Group, "Research and Development FY 2008," *AAAS Report XXXII* (Washington: The American Association for the Advancement of Science, 2007). The quoted text is in Chapter 5, "R&D in the FY 2008 Department of Defense Budget." A preview version is available online at http://www.aaas.org/spp/rd/dod08p.pdf (last accessed: November 2, 2010).

[4] Alison St. John, "Cuts In Department of Defense Spending Could Hit San Diego's Economy," *KPBS*, August 10, 2010, http://www.kpbs.org/news/2010/aug/10/cuts-department-defense-spending-will-hit-san-dieg/ (accessed: October 31, 2010).

[5] Nicole Blake Johnson, "OMB identifies 26 IT projects for closer scrutiny," *FederalTimes.com*, August 23, 2010, http://www.federaltimes.com/article/20100823/IT04/8230301/ (accessed: November 8, 2010).

[6] Commonwealth of Massachusetts, "Standards for the Protection of Personal Information of Residents of the Commonwealth," *201 CMR 17.00* (Springfield: Commonwealth of Massachusetts, 2010). See especially section 17.04 that directly correlates to computer security for firms storing data of residents of the state *regardless of the data holder's physical location*. Compliance is required as of March 1, 2010.

[7] Department of the Army, "Frequently Asked Questions about Recruiting," *Support Army Recruiting*, http://www.2k.army.mil/faqs.htm#age (accessed: November 6, 2010).

[8] Deputy Under Secretary of Defense for Counterintelligence and Security, "Operating Manual Supplement," *National Industrial Security Program [DoD 5220.22-M-Sup 1]*, available online at http://www.fas.org/sgp/library/nispom_dod_overprint_rev1.pdf (accessed: November 5, 2010). This exhaustive supplement defines the requirements for commercial organizations providing a Special Access Program (SAP), especially if the organization must store and process data classified as SECRET or TOP SECRET. While in our case we are just beginning this process, many of our initiatives over the past year have been driven directly from this document.

[9] Department of Defense, "Joint System Administrator Checklist," *Information Assurance Support Environment*, Defense Information Systems Agency, December 22, 2005, pg. 3. Available online at http://iase.disa.mil/stigs/checklist/joint-system-administrator-checklist-dec05.doc (accessed: November 14, 2010). This checklist contains a series of steps to be run daily, weekly, monthly, quarterly, and annually. In our case we typically concentrate on the initial hardening procedures for delivery to the customer, although for ongoing projects we must apply the process-oriented functions to satisfy Operations and Maintenance requirements.

[10] Centers for Medicare & Medicaid Services, *CMS Policy for the Acceptable Use of CMS Desktop/Laptop and Other IT Resources*, December 8, 2008. Available online at http://csrc.nist.gov/groups/SMA/fasp/documents/policy_procedure/Policy_Desktop_Laptop_Resources.doc (accessed: November 1, 2010). Based on our organization's desire to submit proposals to the Department of Veteran's Affairs, we used this Acceptable Use document as a template for our own.

[11] Rick Howard, "The US DOD Proposes their Cyber Security Plan," *Notes from the Cyber Trenches*, September 9, 2010. Available online at http://blogs.verisign.com/idefense/2010/09/the-us-dod-proposes-their-cyber-security-plan.html (accessed: November 1, 2010).

[12] William J. Lynn, III, "Defending a New Domain: The Pentagon's Cyberstrategy," *Foreign Affairs*, September / October 2010. Available online at http://www.foreignaffairs.com/articles/66552/william-j-lynn-iii/defending-a-new-domain?page=show (free registration required, accessed November 3, 2010).

[13] Department of the Army, "ADSL Service Interface Specifications," *Army Net-Centric Data Strategy*, http://data.army.mil/ADSL_SIS.html (accessed: November 7, 2010).

[14] Department of the Army, "Authoritative Data Source (ADS)," *Army Net-Centric Data Strategy*,

http://data.army.mil/datastrategy_authoritative.html(accessed: November 7, 2010).

[15] Department of the Army, "Information Assurance: Rapid Action Revision (RAR) Issue Date: 23 March 2009," *Army Regulation 25-2* (Washington: October 24, 2007). Available online at http://www.army.mil/usapa/epubs/pdf/r25_2.pdf (accessed: November 8, 2010).

[16] Robert Charette, "DoD Confirms Flash Drive Breached its IT Security in 2008," *ieeespectrum riskfactor blog*, August 25, 2010, http://spectrum.ieee.org/riskfactor/computing/it/dod-confirms-flash-drive-breached-its-it-security-in-2008 (accessed: October 28, 2010).

[17] National Security Agency, "Global Information Grid," *NSA Web site*, http://www.nsa.gov/ia/programs/global_industry_grid/index.shtml (accessed: October 29, 2010).

[18] Warren Suss, "5 lessons from DOD's cloud computing efforts," *Government Computer News*, September 23, 2010. Available online at http://gcn.com/Articles/2009/09/28/Warren-Suss-5-lessons-of-DOD-cloud-computing.aspx (accessed: October 29, 2010).

[19] Sean Gallagher, "Army confronts battle to globalize its network resources," *defensesystems*, April 7, 2010. Available online at http://defensesystems.com/articles/2010/04/06/defense-it-1-gnecs-uphill-climb.aspx (accessed: October 29, 2010).

[20] Ang Li, Xiaowei Yang, Srikanth Kandula, Ming Zhang, "CloudCmp: Shopping for a Cloud Made Easy," *USENIX Workshop on Hot Topics in Cloud Computing (HotCloud) 2010*. Available online at http://research.microsoft.com/en-us/people/mzh/hotcloud10-cloudcmp.pdf (accessed: November 17, 2010).

[21] Barack Obama, "Remarks by the President on Securing Our Nation's Cyber Infrastructure" (speech, East Room of the White House, Washington, DC, May 29, 2009).

[22] Leslie Phillips, "Lieberman, Collins, Carper Unveil Major Cybersecurity Bill to Modernize, Strengthen, and Coordinate Cyber Defenses," Senate Committee on Homeland Security & Governmental Affairs, June 10, 2010, available online at http://hsgac.senate.gov/public/index.cfm?FuseAction=Press.MajorityNews&ContentRecord_id=227d9e1e-5056-8059-765f-2239d301fb7f (accessed: November 7, 2010).

[23] Frank W. Deffer, *DHS Needs to Improve the Security Posture of Its Cybersecurity Program Systems [OIG-10-111]* (Washington: Department of Homeland Security, August, 2010). Available online at http://www.dhs.gov/xoig/assets/mgmtrpts/OIG_10-111_Aug10.pdf (accessed: October 31, 2010).

[24] Todd Klindt, "KB973917 breaks SharePoint to pieces," *Todd Klindt's SharePoint Admin Blog*, January 4, 2010, http://www.toddklindt.com/blog/Lists/Posts/Post.aspx?ID=178 (accessed: November 1, 2010). In this incident, a Windows security patch brought down a Microsoft Office SharePoint Services (MOSS) server.

[25] The DoD provides a Common Access Card, or *CAC*, that contains an encrypted and highly-protected private key for each active-duty or retired military person and his or her dependents (see http://www.cac.mil/ for information). DoD also provides CACs to civilian workers and DoD contractors (contractors require sponsorship). Interestingly enough, a DoD CAC can be used to login to any type of Windows operating system – all that is required is to export the *public* key from the card and associate that key with the user's Active Directory account. This is the approach we used when validating our thin client solution prior to integrating it company-wide (to our chagrin, this was not quite enough testing).

[26] Interview with Chief Executive Officer, October 19, 2010. Our future business needs list came directly from that interview.

[27] Peter R. Stephenson, ed. *Information Security Essentials: Section 2* (Auerbach Publishing, ISBN 978-1-4398-0030-0, 2009), pp. 358 and 359.