Ten Recommendations for Improving Government (or anyone's) IT Security:

July 5, 2015
By Ken Foster

In part one "Why Federal Government IT is both a Target and Hard to Fix" we identified six areas that were core to the problem [1]. In this segment, we will explore possible solutions to address these shortfalls. Before we begin, it's critical to recognize most government IT professionals go to work every day and put forth a significant effort to keep our citizens information safe. The public never hears about these efforts, only the media whaling how could this have happened when something goes wrong. Often as we will discuss further, these professionals are frustrated by budgetary and management roadblocks. I hope this article will both provide an avenue for renewed support of these efforts or at least recognition for those cyber warriors that their efforts are heading in the right direction.

1. Identifying Magic Bullet Syndrome:

This is the fallacy that single vendor solution will magically solve all your cyber security problems. I refer to this as the "Security as a SKU" mentality. Cyber is a team sport with tools and solutions that bring a variety of unique capabilities to the network. In the historical sense, we call this Defense-in-Depth. By layering different mitigating measures, we increase the complexity to an attacker and lower the likelihood of an undetected successful penetration.

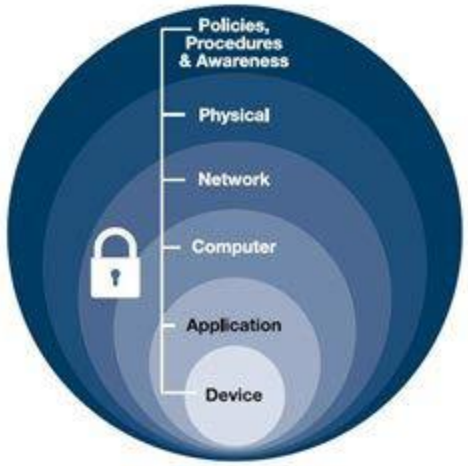

However, regardless of how good the tools within the layers, a simple truth remains. Who's watching the castle? Deploying tools that are passively monitored, where alerts are not reviewed, cleared or escalated to an Incident Response team is honestly a "Management Feel-good" solution, not a cyber security program. Often this practice is can correlated to either management pushing new solutions based on slick vendor presentations that no one wanted or asked for, or the lack of understanding that new processes must be supported by adequate manpower. For example, say your shop recently deployed Host-based Intrusion Prevention (HIPS) into an environment that did not previously have this technology. Whose responsibility is it to tune, monitor, and mitigate detected issues? How does this impact their existing workload and do they have existing capacity to adequately take on this task? Have Standard Operating Procedures (SOPs) been updated to reflect the tools capabilities and integrations? Likely the answer to all these questions in a resounding No! Even with tools such as Einstein monitoring your network, you have to be engaged enough to follow up and research anomalies [2]. A cyber security tool unmonitored is a one that is underutilized.

2. Applying the 80/20 Manpower Rule to Cyber Security:

This directly correlates to the problem of the Magic Bullet syndrome. Too many tools, not adequately monitored. The Single Paine of Glass management tool is typically a single vendor implementation which raises issues regarding exploits that impact the entire vendor framework. In cases where a true multi-vendor dashboard exists, it's often very expensive and not agency extensible when new technologies are added or upgrades cause incompatibility. What is really needed is a back to basics approach to tool usage. This could be achieved if more organizations structured their cyber defense teams using the 80% manpower model. In this model, team members are never tasked above 80% of their total available time. This allows them to spend the remaining 20% reviewing information, troubleshooting problems, and identifying process improvements. Using this model, alerts are effectively analyzed, escalations made, and the meantime to detection is lowered dramatically. This push begins in the trenches with effective process documentation and SOP development. Once the 'What', 'Where', and 'How' are documented, then measure the required 'Time' to establish the baseline standard. My multiplying the 'What' and 'Time' you can develop the manpower requirements. Next you must develop a dispassionate presentation (PowerPoint or Whitepaper depending on organizational culture) that documents the manpower requirements. Once the requirement is documented, correlate the workload capacity at current staffing and highlight the delta of unaddressed events based on current manpower. By graphing the risk delta, it helps management visualize the need for additional staffing in terms of passive unrealized risk acceptance. It's not a guarantee of increased manpower, but it does signal to management there is a problem that you have documented and brought to their attention. As a manager, the first step to problem resolution is identifying the issue. It's also critical to acknowledge that no manager wants to be 'that guy' who was told about the problem and chose on their own to not take action. Never underestimate the power of self-preservation within your management team.

3. Compliance and Security are not the Same Thing:

In government, we are required to adhere to various regulatory best practices and laws. These are often generalizations designed by legislative bodies with the best of intentions and designed to encourage higher levels of baseline security because to address shortcomings found in the practices of those whom came before you. Unless you are in the Department of Defense and familiar with their mandatory Security Technical Implementation Guide (STIGs) most of the compliance guidelines tend to be rudimentary at best. Take for example NIST 800-53 revision 4, control IA-5, Authentication Management [3]. This control states both minimum complexity and length of passwords are organizationally defined. So if an organization decides to set a minimum password length of six characters and no complexity, so long as they document this in their agency policy it would meet the compliance requirement. Whereas a STIG might state the minimum password length is ten characters (establishing a floor) and must be complex including upper, lower, numbers, and special characters (or to the maximum a device supports) as the requirement. Regardless of the level of regulatory guidelines and agency enhanced requirements, these are still intended to set the baseline standard. Because each agency is unique with regards to requirements, technology, and supported constituency, cyber security must identify granular risks associated by logical system boundaries. For example, workstations that access the tax information of citizens such as their social security, associated banking information, and other identifiable client data represent different risks and detection / mitigation controls than the postage machine in the mail room. By identifying the types of data processed, internet access required, and associated applications required by role, agencies can build risk profiles. Once risk profiles are developed (see FIP 199), regulatory controls and additional mitigation measures can be layered [4]. This customizes the defense-in-depth measures applied based on risk and sensitivity of the information processed. By doing so, it both allows cyber security professionals to maximize their tools and time while reducing costs. This is a very time consuming process and no automated tools exist to accomplish this task. This is an area of opportunity for private industry to develop tools and services to meet these requirements. Absent of vendor support in this area, the executive level management team should empower a workgroup to collect, identify, and provide recommendations for the agency. Those assets will likely need to be tasked to this effort fulltime for an extended period due to the projects complexity. Even in moderately complex agencies, this should be achievable within 90 days. The key is to treat this like any other project with deliverables and timelines. Don't forget to share the results with the Disaster Recovery team. You may help identify critical dependencies not previous known.

4. Data Encryption as a Cyber Security Enhancement:

Thinking about the recent OPM data breach as an example, encryption could have inhibited the data exfiltration, maybe? That doesn't sound as definitive as many journalists in the early days of the story have written. Why, because this is truly a 'depends' answer. Encryption provides data protections in two different situations. First, when the data is at rest (not in use) it can be stored in an encrypted format. Simple SQL bulk extraction (data dumps) from the tables using tools such as SQLMAP would have returned the encrypted data, which would require offline cracking of the key. However, if the attackers had administrative access to the server, they could extract the key using various techniques (not covered in this article) to recover and decrypt the data. The other scenario where encryption enhances security is for data in transit. If the connections between the data source and the consumer are encrypted, it inhibits successful interception, thus achieving confidentiality. However, this can be circumvented though man-in-the-middle attacks. If the bad actors had control of the web server hosting the lookup page, there would have been numerous methods that could have been used to capture the unencrypted data at time of display for exfiltration. What is a best case recommendation for encryption usage? Unfortunately this requires a lot of work to recode applications to accept encrypted data. As a general best business practice organizations should:

- Normalize data storage
- Encrypt at a minimum personally identifiable data columns within each table (e.g. SSN, Name, Address, Account Number, Banking / Financial Data, etc…)
- Create an internal unencrypted user ID and associate that in place of the SSN for lookup fields, reporting, etc…
- Resist exposing or using the citizens SSN in all but the most extreme or regulated circumstances
- Establish role-based access to unencrypted PII whereas only the minimum elements are available; use partial making wherever possible (e.g. only display the numeric value of the street address if used in identity validation process)

- Enable full logging, set alerting on failed/denied access attempts
- Establish an SOP that requires monitoring and analysis of all alerts within 60 minutes (You can steal a bunch of data in < 60 minutes)

By embracing these concepts you alter your data exfiltration likelihood from low hanging fruit to a significantly harder task. In the context of risk, consider that the majority of embarrassing data dumps on Pastebin are by moderate-level attackers. Encrypting the data at rest could prevent sensitive information exposure, changing the breach notice from one requiring data monitoring to one of simple notice.

5. Patching, I think Bob does that?

I can't tell you how many times I have been involved in a cyber security assessments where the organization really didn't have a good handle on their security patch state and didn't know it. In IT organizations with limited resources, patch management is often handed off to a mid-level administrator as an additional duty. Thinking back to the 80/20 manpower rule, failing to provide the appropriate time and resources to accomplish this will result in unknown risk 100% of the time. In many Windows shops, patch management is relegated to a Windows Server Update Service (WSUS) server [5]. In this method our administrator Bob selects the patches available for the known Operating Systems (OS) and Applications (Apps) in the environment and releases them for installation from the WSUS server to the client systems. There are two important caveats in the prior statement; known OS and Apps. If the organization does not have a good handle on legacy standalone Applications such as Visio, they may fail to continue to provide updates. Another scenario that occurs is patch application anxiety. This is a state where prior incomplete testing has resulted in prior outages, causing organizations to not apply patches in a timely fashion or ever. This is the low hanging fruit that Penetration testers and bad actors seek. According to the Verizon 2015 Data Breach Investigations Report (DBIR), patches are being converted and added to exploit tool kits in ever increasing speed; sometimes within a month of their release [6]. It only takes one client-side exploit hosted on a malicious website to obtain that critical first foothold. Another related issues is the lack of 3rd party patch management and reporting tools, unfortunately a far too common situation. While there are numerous options in this space to manage both Microsoft and 3rd party patches (e.g. Adobe everything, Java, Apple, ManageEngine, etc…) in a single solution, most organizations either do not do this or fail to validate their successful deployments. This brings up another point worth making. To quote a co-worker, "You must inspect what you expect". Just because Microsoft System Center Configuration Manager (SCCM) successfully deployed a patch, doesn't mean it mitigated the risk. The textbook example was the GDI+ patch released in 2011. Microsoft diligently released a patch to address this privilege escalation issue, however that is not the only place where GDI+ could have been introduced into your systems. If you use HP printers, they graciously provided the exploitable dll for months after the patch was released within their driver and software installer packages. How often do enterprises update printer drivers and the associated applications they deploy? Since Microsoft only addresses the OS deployed vulnerability, the HP risk still existed and as was exploitable on many machines for years [7]. Agencies need to procure, develop SLA's, deploy, and use vulnerability scanners within the enterprise. Recognizing that some patches can introduce higher risk to operations than others, creating a

VM test network that mimics production is critical. It allows patches to be tested on representative systems without causing critical service outages. These SLA's should require a full testing of patches of vendor developed and home growth application (to include testing documentation) within 30 days of release. In our organization, the administrator responsible for our enterprise patch management solution also performs the pre-deployment testing. It's a natural marriage since he would also likely be the one to remediate a patch that causes failures. This methodology can only be successful in an 80/20 environment where there is adequate time for testing, deployment, and monitoring.

6. Remediate or Quarantine the Long and Short-term Solution:
In IT security we have to expect that a certain percentage of systems will encounter issues. Those may be related to corrupted patch catalogs that will no longer accurately accept patching or systems that have corrupted registries. Regardless of the situation, if you can't successfully remediate the problem, then you need to either quarantine or reimage. There are situations where a system runs a legacy critical application that cannot be replicated. In those cases, your only choice is the quarantine the system (e.g. complete logical isolation and prohibited internet access) until a long-term solution can be identified. This compromise solution allows for the application to continue to function temporarily at the cost of user convenience. However, sometimes production workstations and servers simply stop taking security patches. When this occurs, wasting hundreds of hours troubleshooting an errant registry key value provides far less value than exercising the Disaster Recovery Plan (DRP), spinning a new host, applying the required applications, and swapping the new server into production. It you are not embracing virtualization for critical servers and using server templates as part of your DRP, you need to rethink your plan. Don't forget to update the patch levels on your templates no less than quarterly. It shortens the time to deploy a fully patched server.

7. Know Who's on the Wire:
If you are a medium to large organization with multiple remote facilities, knowing exactly what systems are riding your network can be challenging. It's not uncommon for well-meaning employees to graciously extend your network via home quality wireless access points without your knowledge. They do this because they believe they need the convenience these tools offer do perform their jobs more effectively. Since they're not cyber security professionals they often do not understand the additional risk they are introducing. Another scenario that can occur are devices that bridge external and internal network (e.g. Cellular hotspots hosted on laptops) providing gateways for uninvited guests on your network. Typically agencies would address this issue using some sort of Network Access Control (NAC) solution. These solutions are both costly to install and require constant care and dedicated monitoring to catch rouge devices once they connect. So why go through all the grief if they are so hard to manage? An employee who plugs in their personal laptop to the agency network to listen to music may also be accidently introducing an infected machine with ransomware, destructive malware, worms, or other uninvited guests. Since it's unreasonable to expect a home user to provide the same level of cyber security as the enterprise, why would you allowing uncontrolled port-level access to the agency networks to those same devices? In doing so you are passively allowing users to bypass potentially hundreds of thousands of dollars in invested perimeter and network defenses. There

may also be potential risks associated with insider threats and untrusted devices. Consider this, an uncontrolled device with unknown software and an unknown infection state can do pretty much access, scan, document, attack, and exfiltrate data from any system the devices subnet is allows it to access. At a minimum embrace a stand naming convention and routinely scan DHCP reservations for unusual devices; track them and take appropriate action. Consider randomly dumping DHCP scopes and seeking outlier MAC addresses which could be a sign of a device intentionally attempting to spoof your logical controls and ride the wire undetected. For example, if you don't buy Panasonic Toughbook's, why is one on your network in accounting? If you decide to deploy a NAC solution, seek vendors with proven track records in NAC deployment. Ensure you Integrate into your existing policies agency guidance prohibiting connection of personal IT hardware to the network. This should include mobile devices. Agencies without approved Bring your Own Device (BYOD) policies should not allow employees to connect personal devices to agency systems for charging via USB cables. Doing so allows the device to open a storage-level connection to the agencies systems that can be exploited (unintentionally or maliciously) in numerous ways. It you have an existing BYOD policy, ensure you have developed both an isolation strategy that limits the devices access to sensitive areas of the network and the authority to capture and possibly wipe a device that is storing non-public agency information if the employee is terminated . If the management and legal teams are not willing to provide that capability, then your organization is not ready to introduce that risk to the enterprise, regardless of the convenience and capabilities it provides. If management will not back a whipping policy it's a management culture problem whereas they do not understand the risk and what they are accepting. Just make sure you document your recommendation and their refusal for when the music stops; this will help ensure the person short a chair is not you!

8. Training Your Users to be Successful:

We hire employees and contractors to do work that makes our agencies more successful at providing their core services. Since the majority of those employees are not hired to be cyber security professionals, we shouldn't expect them to conduct themselves in that manner. If we want to use our employees to help us detect malicious activities (User's as a Sensor or UaaS), we need to provide them meaningful, relevant, and engaging training appropriate their level. The challenge is to identify training these is generally role appropriate, current, and filled with relevant examples that emphasizes the point(s) you are trying to make. I find using a mixture of modalities tends to reach the most users. Modalities I would recommend include audio (story telling), Visual (screen shots), and Kinesthetic (activity-based) learning methods. In this example one might breakdown the topic of phishing as:

- What Phishing is (Audio)
- How it's used (Audio / Visual via PowerPoint)
- Success rate associated with the tactic (Audio / Visual via PowerPoint)
- Indicators of an active attempt (Audio)
- What to do if employees encounter a suspected phishing attempt (Audio / Visual via PowerPoint)
- Screen shots of various Phishing Attempts (Visual)
- Jeopardy Game to Detect Signs of Phishing (Kinesthetic)

Consider the process of Gamification of your training. Divide your students into mixed groups’ representative of several departments each. Pre-arrange with management to award the group that is most successful a Pizza Lunch immediately following the training, 2 hours off with pay, or some other incentive that enhances their focus and teamwork towards the reward. It will be the cheapest and most effective risk mitigation measure you undertake in cyber security.

9. Know what your Network is Saying About You:
A lot can be learned from a simple open source scanning of your network and general public web sites. Does our DNS server share internal network systems names across the internet? Are there more ports open to DMZ devices that necessary? What about from your DMZ to the internal network? Do you block ICMP from the internet to your internal network? When was the last time you validated your external controls? What about Requests for Proposals / Quotes; sharing the hardware and security measures you use? Consider performing external port and connect scans no less than quarterly to ensure you’re only exposing what you intend to share. An example of such a scan might look similar to:

```
nmap -v -T2 -O --top-ports 500 --min-parallelism 10 -oN
  external_scan_result.txt 192.168.10.1-24
```

Obviously you would need to replace the IP address range with your own. If you find excess open ports, block them before someone else uses them in ways you did not intend. Another technique you can use is called Google Dorking [8]. This leverages the power of Google indexing to find sensitive information posted by or about your company that you may not know is being exposed. An example of types of searches you could conduct might include:

```
Inurl acme.com intitle:"index of" mysql.conf OR mysql_config
```

Assuming that was your domain and you use MySQL. Clearly there are many different dorks you can perform. Many of these have been automated into Pentesting (or attacker) scripts, simplifying the detection overhead required. If you do find something concerning, mitigate the risk where possible (e.g. change passwords, apply Access Control Lists, submit takedown requests, notify affected parties, etc…).

10. Hold Contractors Accountable:
Finally the hardest and least likely recommendation to be followed is holding contractors responsible for the configuration and viability of the code they deploy. Before you buy an application, have the vendor demonstrate they patch their code, how often it’s patched, and how much longer the application will remain under support. If an application will reach the end of the support cycle next October, it doesn’t make sense to buy it this May. When buying an application insist the vendor provide you support for a minimum of three years from date of delivery or provide a no-cost upgrade to a supported version. That support should specify functionality with all host OS security updates. That will be a hard sell, but at the end of the day they want your continued business more than you realize. If they don’t, then find one that does. If you are buying custom code, insist on the same level of support and place a timeline of no

more than 60 days from patch release to tested solution. Collaborate with other agencies to include this language in their contracts. Share between contracting agencies those vendors that meet those standards in order to drive more vendors into compliance. Remember, at the end of the day it's the agency IT team that will be impacted most when a vendor doesn't support their products. Don't get stuck holding their bag and stop doing business with those who put you in that position.

Just remember, if IT Security was easy, everyone would do it… Until the next article.

References:

[1] Part One Article: Why Federal Government IT is both a Target and Hard to Fix, https://www.linkedin.com/pulse/why-federal-government-both-target-hard-fix-ken-foster

[2] DHS Einstein 2, Network Flow Analysis Engine, http://www.dhs.gov/publication/einstein-2

[3] National Institutes of Standards and Technology (NIST), Special Publication 800-53, revision 4, http://dx.doi.org/10.6028/NIST.SP.800-53r4

[4] Federal Information Processing Standards (FIPS) Publication 199, Standards for Categorization of Federal Information and Federal Information Systems, February 2004, http://csrc.nist.gov/publications/fips/fips199/FIPS-PUB-199-final.pdf

[5] Microsoft Windows Server Update Services, Technet, https://technet.microsoft.com/en-us/windowsserver/bb332157.aspx

[6] Verizon 2015 Data Breach Investigations Report (DBIR), Verizon Inc., http://www.verizonenterprise.com/DBIR/2015/

[7] HP Compaq Notebooks ActiveX Remote Code Execution Exploit, Author Porkythepig, 2007, https://www.exploit-db.com/exploits/4720/ Note: this is an example of one such method HP exposed customers, not the only one.

[8] The Google Hacking Guide, KMBL Security,http://www.uscyberwarrior.com/Articles/Stories/Google%20Hacking%20Guide.pdf