Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

Interested in learning more?

Check out the list of upcoming events offering
"Auditing Systems, Applications, and the Cloud (Audit 507)"
at http://www.giac.org/registration/gsna

Securing Single Points of Compromise (SPoC)

GIAC GNSA Gold Certification

Author: David Belangia, dwbelangia@hotmail.com
Advisor: Stephen Northcutt

Abstract

Securing the Single Points of Compromise that provide central services to the institution's environment is paramount to success when trying to protect the business. (Fisk, 2014) Time Based Security mandates protection (erecting and ensuring effective controls) that last longer than the time to detect and react to a compromise. When enterprise protections fail, providing additional layered controls for these central services provides more time to detect and react. While guidance is readily available for securing the individual critical asset, protecting these assets as a group is not often discussed. Using best business practices to protect these resources as individual assets while leveraging holistic defenses for the group increases the opportunity to maximize protection time, allowing detection and reaction time for the SPoCs that is commensurate with the inherent risk of these centralized services.

1. Introduction

Single Points of Compromise (SPoC) are key enterprise central services that could be misused by an intruder or an insider to compromise critical portions of an enterprise's computing environment. When determining what services should be classified as a SPoC, consider services where a compromise would allow login or root login on many assets within the environment ensuring a complete ownership of the institution's environment by an adversary. A compromise that provides available attack surfaces against other computers in the enterprise environment should be considered a SPoC. Finally consider compromises that would yield information of high value to an adversary allowing an understanding of the environment conducive to attacking the enterprise. Protecting these SPoCs will require changing processes and developing methodologies that force additional hurdles for adversaries but will increase complexity for those employee managing the central service. The decision to identify a SPoC and have it comply with a suggested protection is a decision the institution will have to make based on the institution's environment, information, and their adversary's motivation and skillset.

Protecting an individual critical component can require a variety of tools including constant control, auditing, and responding to any change. Some of these assets would include Active Directory, Authentication Services, Configuration Management, Domain Name Services, File Transfers, Firewalls, Infrastructure on Demand, Logging Services, Vulnerability Scanning, and Web Services used to manage a SPoC.

Active Directory is not easy to secure but there are steps that can be taken that include 1) following Administrator best practices, 2) following Domain Controller best practices, 3) following delegation best practices, 4) monitoring and auditing the Active Directory organization, and 5) preparing for the worst. (Allen, 2005)

Authentication services are about keeping unauthorized access from occurring and ensuring authorized access does occur. Authentication verifies the user's identify, while authorization verifies the permissions and rights the authenticated user has. (Shinder, 2001) Various methods can be used for authentication (e.g. password, smart

David Belangia, dwbelangia@hotmail.com

card, and biometrics) that add additional value to ensuring only authorized access is permitted.

Configuring assets in a secure manner will involve literally hundreds of settings. There are many settings for each type of asset that can be configured that provide protection and developing the configurations from scratch is too labor intensive. It is highly recommended that the starting point for an asset's configuration be modelled on a standard like the United States Government Configuration Baseline (USGCB) or DoD Security Technical Implementation Guides (STIG) configurations. From the starting configuration, specific changes can be made to adapt to the unique requirements that are required by the institution.

The accuracy and integrity of the Domain Name System (DNS) records are critical for successful execution of activities within the institution. Misconfigurations can result in blocked or miss-delivered email, hijacking of a web site, and even root compromise of the system. (Griffin, 2015)

File Transfer services are a necessary evil and can provide an easy vector for introducing malware. Managing a secure transfer with auditable information exchange is challenging. Most companies in every industry and government organization are required to exchange information online with their partners, suppliers, customers, and other constituents. (ContentSecurity, 2015) Securing this service can prevent the introduction of malware during the transfer process. By scanning the information prior to introduction, there is a chance of ensuring good condition (absence of malware).

Firewall management is still a primary defense. The requirement to ensure that these assets are configured and managed appropriately will require attention. Dell advises that effective firewall management must include: 1) clearly define change management processes, 2) the test of a change before implementation, 3) backup the configuration prior to making a change, 4) monitor and log all changes, and 5) regularly check the configuration against policy. (SecureWorks, 2015)

Many organizations are now using Infrastructure on Demand to optimize equipment use and to leverage improved security. Paul Zimski, VP of Solution Marketing at Lumension, articulated that managing virtual images has inherent risk. He advised
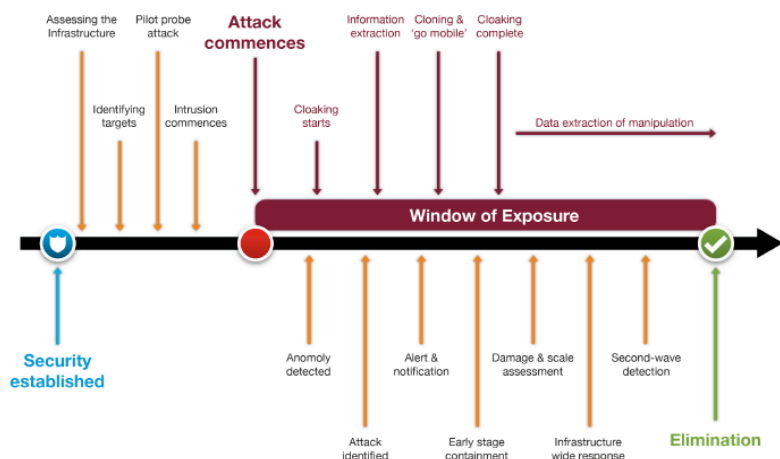
David Belangia, dwbelangia@hotmail.com

patching, involvement of security professionals in deployment, and managing access control are paramount to ensuring a secure environment. (LeClaire, 2013)

Logging Services can provide early alerts to a potential issue and are necessary for the post moratorium analysis after the compromise. Without the correct logs, correlation, and analysis, the victim is left without being able to understand and/or mitigate an issue. NIST 800-92, *Guide to Computer Security Log Management* provides recommendations for log management. With the explosion of network servers, workstations, and other connected assets and the increasingly sophisticated threats, the need for logs has exploded. (NIST, 2006)

Vulnerability scanning is the automated process of proactively identifying security vulnerabilities of computing systems in a network in order to determine if and where a systems can be exploited and/or threatened. (Vulnerability Scanning, 2015). Keeping up with vulnerabilities without a scanning service is impossible. This service helps automate security auditing looking for the thousands of vulnerabilities in the environment and will provide a prioritized list that describes the vulnerability and steps on how to remediate the vulnerability. (Grier, 2014)

Using web services to facilitate management of a SPoC should be undertaken with caution. The SANS Institute provides an excellent checklist outlining best business practices that raise awareness and help implementers to identify and work through issues relating to their use. (SWAT, 2013)

It has been argued that every organization will be compromised or has been compromised. Using a Time Based Security (TBS) Model can help the institution survive this inevitable event. The TBS Model is composed of three concepts; protection, detection, and reaction.

David Belangia, dwbelangia@hotmail.com

Protection must be greater than the combination of detection and reaction. If an attack can be detected and reacted to before the protections fail then the defense in depth has provided sufficient time to detect and react. (Schwartau, 1998) This relationship is depicted in the graphic above. (Gigamon, 2012)

Protecting SPoCs individually is well understood and an important part of operations that the institution must accomplish. Protecting the SPoCs as a group using common protections and the additional layering of controls will provide additional protection time while reducing detection complexity due to centralization of these layered controls. It will make the reaction time less due to solid management of the SPoC increasing the likelihood of success.

2. Protecting the Critical Assets

2.1.   Critical Assets (SPoCs)

Protecting SPoCs within the environment by identifying the critical central service and then applying a level of rigor to those assets commensurate with the asset's importance will assist in multiple areas of the TBS formula. Critical central services that could be misused by an intruder or an insider to compromise large sections of an enterprise's computing environment are highlighted using research of the most common issues affecting the service. Some suggested SPoCs are Active Directory, Authentication Services, Bastion Hosts, Configuration Management Services, Domain Name System, File Transfer Services, Firewalls, Infrastructure on Demand, Logging Services, Scanning Services, and Web Management Service for a SPoC.

2.1.1.   Active Directory

Securing Active Directory requires special attention to include:

1. Following Administrator best practices,

2. Following Domain Controller best practices,

3. Following delegation best practices,

4. Monitoring and auditing the Active Directory organization, and

David Belangia, dwbelangia@hotmail.com

5. Preparing for the worst. (Allen, 2005)

Following Administrator best practices requires the use of special consideration with regard to administrative accounts. Separate administrative accounts must be required. Start by disabling the guest account and renaming the Administrator account including removing the default description of these two accounts. Administrators should always use the non-privileged account for all daily work and only exercise the administrator account when absolutely necessary.

Using "runas" or a similar tool can allow the execution of programs as an administrative user while logged into a non-administrative account. Runas is a command in the Microsoft Windows operating systems that allows a user to run a specific tool or program as a privileged user much like the Unix commands sudo and su. By performing the command on a Windows system and immediately exiting the command prompt or program after completion, the use of privilege is limited to that command. This provides an easy way to allow an unprivileged user elevate privileges for a specific activity requiring a password.

Maintenance of group membership for administrators is extremely important to prevent privilege creep. Monitoring unauthorized for any unauthorized additions is important. Window Group membership must be actively managed and an alert generated whenever someone is added to the Group. As personnel leave or change positions, the appropriate addition and deletion must be made.

Protecting the Administrator account password and using it only as a last resort is important. Having a process to quickly change the Administrator account password is important. This password should be changed frequently even though the distribution of the password is very minimal. Finally the institution must ensure the process for quickly disabling an administrative account is understood and can be executed quickly to avoid a rogue administrator from doing a lot of damage.

The next Administrator best practice is management of the Domain Controller(s) (DC). Maintaining physical control of the DC is imperative. If physical access can be obtained, an adversary can subvert almost any security control. Ensuring DC(s) are protected physically will provide additional security.

David Belangia, dwbelangia@hotmail.com

Using a starting standard configuration as a guide (USGCB or STIG) and building a secure DC, with minimal modifications, there is some assurance that the DC will be secure. DCs replicate quickly and must be consistently applied to other DCs. Important updates for DCs are provided frequently and by the time the update is received by the organization, many of the vulnerabilities that the update would fix are actively being exploited by adversaries. The DC(s) must be updated quickly with subsequent replication.

Versions of Windows 2003 or earlier were vulnerable to a create object vulnerability because of the lack of limits on object creation. If there is Server 2003 or previous version in the environment, upgrade them. If upgrading is not possible, investigate the creation of a reserve file to assist in mitigating the create object vulnerability.

The execution of virus-scanning software is critical in this environment because the DC will replicate file content through File Replication Services (FRS) that are distributed to a large number of servers. The lack of scanning could propagate a virus quickly using the FRS mechanism.

Update the Directory Service Restore Mode (DSRM) password regularly. This password is used to capture a copy of the Active Directory offline. This password can provide the ability for a local operator to copy

David Belangia, dwbelangia@hotmail.com

the NTDS.DIT (AD Database) off the server and reboot the service unnoticed allowing the potential for a compromise.

The third Administrator best practice is to follow delegation best practice. Misconfiguring Access Control Lists (ACL) can provide numerous avenues of attack. Try to keep delegations as simple as possible. Do not assign permission to user accounts but apply permission to a select user group designed to minimize the ability of members to impact other functions. This will help manage attrition or movement in the work force by allowing management by group and not individual. Maintain permission by granting rights to Organizational Units (OU) or a select container. Again this simplifies the organization of the OUs. Ensure the process on granting a right and how that request is made is well documented. A useful tool in this area is Dsrevoke. This tool is described by Microsoft as a command-line tool that when used on a Domain Controller will report all the permissions for a specific user or group on a set of OUs in a domain and can be used to remove a permission for a particular user or group. (Microsoft Download, 2015) In the screen shot provided by Microsoft, Ed Price has two Discretionary Access Control List (ACEs). These ACEs can be removed using the tool.

The fourth best practice is monitoring and auditing of Active Directory (AD) ensuring when a mistake is made or if a change is made, there is a greater possibility of detecting and correcting any issue. Documenting the Active Directory configuration can be tedious but required to understand the configuration. Areas of interest include high-level structures like forest and domain configurations, OUs, top-level directory security, trust relationships, site topography, and any manually connected object. Documenting this information can be easier using a tool like Group Policy Management Console (GPMC). (Dueby, 2006) There are many recommendations on items to monitor such as performance monitoring of the AD (LDAP writes, authentication request, etc…) and security events (collect what is use, understand what is require). Suggested audit items include failed logins, successful and failed account management, object access, and policy change. Alerting on a threshold level in this space is important. The support staff that manage this service is always stressed and overworked. By making the alerts timely and appropriate, the opportunity for success will increase.

David Belangia, dwbelangia@hotmail.com

There will be a successful attack or mistake made. Planning for this failure is paramount to maintaining the environment. Schedule actual tests by performing forest recovery from the backup. Execute these test in a test environment to prevent mistakes and potentially a catastrophe result.

### 2.1.2. Authentication Services

The primary goal of authentication is keeping unauthorized access from occurring and ensuring authorized access occurs. Authentication and authorization are sometimes confused. Authentication verifies the user's identify, while authorization verifies the permissions and rights. (Shinder, 2001)

The organization must develop a process that will facilitate the maintenance of the authentication process. It would be preferable to have dynamic or near real time feeds to an automated tool managing any account change. When an employee with access changes their position or leaves the company, there must be a review and appropriate action taken on all related accounts. Failure to perform this function will result an account lying dormant over time increasing the attack surface of the asset. Maintenance of access groups must be managed. There are tools that allow a dynamic review of groups with associated alerts. These tolls must be leveraged to discover and resolve any access issue.

Multifactor authentication is preferred on a SPoC with potentially multiple logins using a 2-factor one time only token. The storage of a credential is a common flaw in many instances of software. This can allow potential harvesting of the credential and subsequent replay attack.

The separation of normal and privileged account use is a must. The administrator must not be using the privileged account to perform any routine activity not requiring the privilege. The privileged account must be reserved to activities requiring that privilege capability. In addition, the privileged account must be assigned to the user so attribution is possible when malicious activity occurs. Any transmission of a credential should be encrypted.

David Belangia, dwbelangia@hotmail.com

The user authentication market is mature with many vendors offering a range of products. Gartner Group provides a magic quadrant analysis of the marketplace on a frequent basis. This graphic provides an overview of the market as of December 2014. (Gartner, 2014)

2.1.3. Configuration Management Services

The United States Government Configuration Baseline (USGCB) provides guidelines for creating security configurations for federal agencies to apply to a specific computing asset. It evolved from the Federal Desktop Core Configuration (FDCC) mandate. (NIST, 2015) The Department of Defense provides the Security Technical Implementation Guides (STIGs) that are mandated by DODD 8500.1 and DODI 8500.2 for all DoD systems. (Greenwell, 2010) The Center for Internet Security provides additional guidance and tools. While each of these guidelines are slightly different, it is imperative that something be used by the institution as a starting point and that baseline configurations be established and managed. There are valid reason for an organization to modify these recommended settings, but any change should be understood and documented.

Once a secure configuration has been developed, the organization should the capability to audit and maintain the secure configuration. There are a variety of tools available to perform an integrity check for a change. These monitoring tools provide an indication when something has changed and provide the opportunity to know a change has occurred allowing investigation into why the change was made.

Change Control must be implemented to ensure any change is recorded, vetted, and understood before implementation. A formal Change Control process will assist with

David Belangia, dwbelangia@hotmail.com

anticipated change by understanding the potential impact to the institution. Each change must go through this process ensuring a new baseline is recorded that will allow the monitoring tool to audit the new baseline.

2.1.4. Domain Name System

Domain Name System (DNS) are an easy target for an attacker. These records are paramount for successful execution of the network. Unpatched software can provide a large surface area for attack. Ensure the latest version of software is running or updated as available. A badly configured DNS server can display a problem in a variety of ways. The result of misconfiguration can be slow response or even the prevention of a connection. The correct configuration will prevent an unauthorized zone transfer and improve integrity associated with cache pollution or updates to the zone file.

Most internal hosts do not need to resolve Internet host names. In the design of the DNS name resolution infrastructure, configure the infrastructure to minimize Internet name resolution not allowing an internal asset to perform a recursive lookup for a public name. Recursion should be disabled unless required. A separate DNS structure should be deployed in the organization's DMZ to resolve Internet names. Configure forwarding zones externally to support web and any other asset that will require resolution to an internal host name.

This split DNS configuration provides two hardware solutions and zones with very separate functions. Using a split DNS infrastructure will allow the protection of the private zone on the Intranet never exposing the asset to an external user. On the private DNS zone, modify the root hint file to contain only DNS assets on the organization's Intranet. In addition limit zone transfer in any private zone.

If Active Directory is being used, the software can enforce registration of resource records when dynamic name registration is enabled. This will prevent a non-domain member from being able to register their name. Using discretionary access control lists (DACLs), the software can control which asset is able to register a change to address information. Check and apply zone-checking tools to ensure all assets are configured properly. DNSWALK (SourceForge), dlint, and DOC (Domain Obscenity Checker) are some common tools that provide this capability.

David Belangia, dwbelangia@hotmail.com

DNS cache can provide advantage by storing successful queries in cache preventing the requirement to query another DNS server. Queries have a Time to Live (TTL) assigned when cached. The performance advantage makes this configuration pretty common. Enable cache locking on the DNS server to reduce the risk of an attacker overwriting cached records.

A DNS server will often have multiple IP addresses assigned to the NIC or even multiple NICs. A DNS server is frequently configure to listen for DNS queries on all interfaces. Audit and lock down the interface and IP address where possible. Another useful configuration is to increase the number of source ports available and perform randomization. This will increase the socket pool reducing the predictability of the source port for response.

Monitoring DNS Traffic can help to identify potential security threats. Several items to be concerned about are queries from spoofed source address (or unauthorized), malformed DNS queries, queries requesting name resolution from a bad domain, responses from the institution's domains that are resolving to addresses that are not part of the authoritative zone, or suspicious addresses. (Piscitello, 2014) Firewall rules should be defined to prevent IP Spoofing. This will involve defining a rule to prevent DNS queries from outside the organization's IP space. This will prevent the ability to use the organization as a reflector in a DDoS attack. In addition, enabling inspection of traffic for suspicious or anomalous traffic will protect against some software exploit attacks. Use the organization's Intrusion Detection System to alert on short TTL, DNS queries using TCP, DNS queries to non-standard ports, or unexpected large DNS responses. Use traffic analyzers to identify malware traffic. Review logs from the resolver investigating DNS traffic. (Piscitello, 2014) This approach might help prevent the blacklisting of the institution's address. Staying abreast of issues, vulnerabilities and best practices can be challenging. Subscribe to on-line DNS resources to solicit new information on a routine basis.

2.1.5. File Transfer Services

It will always be necessary to move information into the protected zone of the network. Frequently there will be a requirement to move patches, applications,

David Belangia, dwbelangia@hotmail.com

monitoring profiles, signatures and other information that are required to keep the environment current and well protected. Each transfer will increase risk if not done securely and can provide an avenue for the migration of malware, viruses and other potential harmful payloads. The file transfer service must be well designed with the capability to allow movement of information with attribution to the move requester and approval from another individual preventing an unintended transfer. In addition, the information to be moved must be inspected preferably using various techniques such as signature based scanning and behavior base analysis. Information should not be moved unless it can be assured of good condition prior to the move.

There will be requirements to transfer dynamic information to support Active Directory feeds and potentially domain information. This transfer is a little trickier and must be automated and controlled/logged. The compromise of assets using these dynamic feeds could compromise the entire SPoC environment.

2.1.6. Firewalls

A firewall is designed to control the flow of network traffic. Historically these devices have been used for the perimeter protection but are now leveraged to protect some internal critical assets. There are four major recommendations provide by the National Institute of Standards Special Publication 800-41 to consider when implementing these technologies. They are: 1) Create policy that manages inbound and outbound network traffic (IP addresses and ranges, protocols, applications, and content type), 2) Identify requirements for the firewall based on what will be required (packet filtering, stateful inspection deep packet inspection, application-proxy gatewaying), 3) Use rulesets balancing requirement and performance, and 4) Ensure the management of the architecture policy, software and other components throughout the life cycle. (Scarfone, 2009)

2.1.7. Infrastructure on Demand

Virtualization has achieved wide spread acceptance enabling better utilization of hardware resources and reducing power consumption. "Virtualization is the simulation of the software and/or hardware upon which other software runs." (Scarfone, 2011) While it would appear that the enterprise is simply virtualizing a physical server and the security

David Belangia, dwbelangia@hotmail.com

challenge should be similar, this is far from the truth. Security must be analyzed from the hypervisor platform architectural and the hypervisor baseline function perspective.

Juniper has identified five best practices to ensure the institution leverages virtualization without sacrificing security. They are: 1) only use applications and services required, 2) monitor and ensure protection of the hypervisor, 3) apply access control monitoring all traffic and block any not required, 4) use layered defenses (apply policies, monitor, block, review logs, anti-virus protection, and alert), and 5) build the virtual machine with the end purpose in mind and standardize the build. (Juniper, 2012)

Detailed recommendations for security considerations are provided in NIST 800-125-A. The Special Publication provides 22 recommendations to address the most common threats to hypervisor baseline functionality. (Chandramouli, 2014) These recommendations should be considered by the institution while using virtualization.

Maguire provided 13 tips on securing virtual machine environments:

1. Install only what is required and keep them patched.

2. Isolate each virtual machine and only allow required protocols.

3. Install antivirus programs and keep current.

4. Use strong encryption between the host and virtual machines.

5. Do not allow internet surfing from host computer.

6. Secure accounts on the host machine.

7. Only use what you need and shut anything not used down.

8. If connection between computers is not necessary do not allow.

9. Monitor the log and security events on the host and virtual machines.

10. Ensure hardware use is designed for VM use.

11. Strictly manager remote access.

12. Provide replication and continuity for single points of failure.

13. Avoid sharing IP addresses. (Maguire, 2012)

David Belangia, dwbelangia@hotmail.com

2.1.8. Logging Services

Aggregation of logs to a central server is necessary and important to ensure logs are being collected and are not susceptible to modification. The analysis of the logs can be time consuming and difficult. Ensuring all logged assets are time synchronized allowing correlation of an event is extremely important. The information to be collected must be determined by the institution. The period of time to keep logs is another consideration. Logs can grow quickly, but not having enough during a compromise can prevent a meaningful forensic effort. Aggregation and correlation of log information can be challenging and will require special attention. Gartner Group June 2014 Magic Quadrant on Security Information and Event Management provided information on the market leaders at the time

Monitoring log events correctly will provide a robust status of the network in real time. The problem is that there is a lot to monitor. Ensuring that important events are monitored and correlated with an alert will allow quicker detection and reaction. This ability is paramount to surviving an attack and supporting Time Based Security response. The Cyber Kill Chain advocates that an attack is based on a series of events. (Engel, 2014) Mandiant states, "In 2013, the median number of days attackers were present on a victim network before they were discovered was 229 days, down from 243 days in 2012. On the other hand, organizations still have difficulty detecting when they've been breached. In 2013, only 33% of the organizations to which Mandiant responded had

David Belangia, dwbelangia@hotmail.com

discovered the intrusion themselves, versus 37% of the organizations we helped in 2012."
(Mandiant, 2014)

2.1.9.   Scanning Services

Authenticated scanning will provide the ability to obtain vulnerability information on a device by authenticating to the scanned device and querying the operating system and installed software, including configuration and missing patches. The very nature of this scan will require a very careful approach to address the risk of the administrative credentials and the propagation to the scanned device. The correct design of a solution must consider the following practices:

1. Use a special account dedicated to the execution of the scan.

2. Ensure the ability of this special account to authenticate to the targeted asset.

3. Stay away from clear-text authentication protocols.

4. Mitigate the man-in-the middle attack.

5. Only allow the special account to be valid when scanning by disabling when not in use.

6. Automate this disablement.

7. Restrict host/ip address from which the scan runs. (Berkeley Security, 2015)

Choosing a vulnerability scanning tool should be undertaken with some care. Not all scanners are created equal. Ensuring the scanner is accurate, reliable, scalable, and will provide robust reporting will minimize the headache later during implementation and subsequent use of the tool. There are commercial tools, as well, as free scanners. Some available free scanners are OpenVAS, Retina CS Community, Microsoft Baseline Security Analyzer (MBSA) Nexpose, SecurCheq, and Qualys Freescan. (Grier, 2014) The better known and more highly rated commercial products are Nessus (Tenable Network Security), Secunia CSI, and Core Impact (Core Security). (Lindros, 2014)

David Belangia, dwbelangia@hotmail.com

Nessus offers many different policy options to scan for vulnerabilities or misconfigurations as shown below. Performing a scan is straight forward and can offer a lot of information.

David Belangia, dwbelangia@hotmail.com

2.1.10. Web Management Services

Using a web management tool for the management of a SPoC will require special consideration. Web interfaces are inherently poorly designed. Typical best practice include adhering to suggested mitigation strategies for the OWASP Top 10 Most Critical Web Application Security Risk. These risk include SQL injection, broken authentication and session management, cross-site scripting, insecure direct object references, security misconfiguration, sensitive data exposure, missing function level access control, cross-site request forgery, using known vulnerable components, and invalidated redirects and forward. (OWASP, 2013)

Almost all of these risk can be mitigated by ensuring a Web Application Firewall (WAF) is strategically placed to intercept and cleanse input. The use of a proxy server will resolve some of the remaining risk. If the particular SPoC does require a web frontend for management ensure the application of mitigating controls. Scan the management interface with a web scanner (WebInspect, Burp, Netsparker, etc…) A WAF can be expensive and will require time to implement. The WAF configuration will take time and patience and require a level of rigor to ensure the best use of the WAF capability.

2.2. Best Practices for the Defense of SPoCs

The items identified in this section were developed by the Los Alamos National Laboratory (LANL) Information Architecture Working Group as best practices with the identification of special areas for consideration. Alex Malin led this effort for LANL and in Section 5.0 Credits, the IA Working Group Members are identified.

A fundamental question facing many organizations is if the organization should centralize or segment the management of the information structure of the institution. This argument has been raging for decades. There are advantages in either approach.

Centralization will provide the opportunity to manage assets by a core group of experts and the leveraging of centralized defenses. This can be argued as an approach that will allow tighter control and better use of critical resources. One issue with this approach is that if one asset falls, the whole institution could fall. All of the eggs are in one basket.

David Belangia, dwbelangia@hotmail.com

Segmentation of the service can be a defense-in-depth cyber defense approach. Having good segmentation will allow for an additional layering of defense. If one critical asset falls, the whole institution might not fall. Guaranteeing the asset is well managed in a segmented environment can be difficult. There might be more management layers and different people responsible for the segmented environment.

All SPoC assets must be hardened. Services and/or ports should be configured to not respond to queries, configuration options must be selected in a fashion to minimize the attack surface of the asset, patches must be applied and tested quickly, and all default accounts must be disabled or eliminated. There are automated configuration guides that provide a solid starting point to applying the best security posture to a particular asset. The Defense Information Systems Agency provides Security Technical Implementation Guides (STIGs). These guides can be retrieved at http://iase.disa.mil/stigs/Pages/index.aspx. The United States Government Configuration Baseline (USGCB) can be obtained from http://usgcb.nist.gov/. The Center for Internet Security offers benchmarks, tools, and metric definitions at https://benchmarks.cisecurity.org/downloads/. These are a good starting point for hardening a SPoC asset.

Once hardening has been accomplished, the baseline for each asset must be developed and frequently checked to ensure detection of any change. There are tools within the various operating systems that allow automating this activity. In the SANS AUD 507 class, tools and suggestions for developing a baseline and alerting on any change are highlighted with solid examples and suggestions. There are many commercial tools, operating system utilities and free tools that allow the collection of key characteristics of the baseline that allow comparison on a frequent basis using scripts/cron jobs with potentia email alerts to responsible personnel. Any change must be compared to the formal Change Management Process for each SPoC. The change process must be formal and it is a good idea to ensure that the voting members on the Change Control Board are SMEs for the different SPoCs. If one SPoC is allowed to be compromised, it will place the whole enterprise at risk. These SMEs are the most knowledgeable of the environment and have a vested interest in keeping the environment secure.

David Belangia, dwbelangia@hotmail.com

The patching of a SPoCs must be quick and thorough. SPoCs are the prized jewels for an adversary ensuring a compromise of a SPoC will provide the most leverage for their attack. Another issue with the SPoC is that not only are they critical, but the enterprise expects the SPoC to be available 7x24. Scheduled downtime must be planned and patches must be installed.

Frequently access to SPoCs is managed using Firewall Access Control Lists (ACL). As these rules get more complicated, the ACLs can get complex and create a condition that is unintended. Using a bastion host can eliminate a lot of this confusion. Limiting access to private address space will provide additional assurance that access is managed to an internal specific address.

A bastion host is defined as any computer that is setup specifically for securing a private network. (Joyent, 2015) This device can be referred to as jump host, golden host, jump box, or a bastion host. (CyberArk, 2014) Most commonly, this device is used for network separation. Based on the nature of the configuration, there is an expectation that this host will be more exposed than the SPoC they protect. Various measures are required to protect this asset to include 1) limiting available services and daemons, 2) disabling or limiting available user accounts, 3) limiting the number of network protocols, and 4) closing all ports that are not needed or used.

Limiting services, daemons, and other packages will minimize the attack surface where a software bug or configuration error could lead to a security concern. The host should be as small a profile as necessary to serve the function.

All access should be 2 factor-derived, one-time-password (OTP) and in many cases should require multiple logins. By leveraging a bastion host and using 2 factor / OTP authentication, the institution has provided a layered defense that would require multiple compromises to be successful. Accessibility to the bastion host and any asset behind this host must be tightly controlled including restricted services and source address. The SPoC should be on a physically or virtually separated network space to ensure accessibility is restricted. VLANS are acceptable but vulnerabilities from this implementation approach must be addressed.

David Belangia, dwbelangia@hotmail.com

Using a VPN connection to the bastion host can add another level of security. Authentication to the VPN service should be two-factor / OTP authenticated and limited to the selected asset. After accessing the VPN service, then the bastion host should force additional authentication improving security. Time outs must be enforced on the connection to ensure after inactivity the connection is terminated. Additional restrictions can be enforced on the bastion host allowing granular restriction of users and what can constitute a connection.

The disablement or removal of unneeded user accounts on all SPoCS will ensure only those accounts necessary exist. Audit these accounts carefully and make sure all accounts are still required.

The same logic applies to unneeded protocols. Start with a "deny all" approach and then add only those protocols required. Monitor the protocols regularly ensuring they are still disabled or if any change occur.

All unnecessary services must be disabled. Within this environment only those services explicitly required should be enabled and all other services should be default deny. Minimizing services and available ports will simplify the network monitoring while ensuring only necessary items are available reducing the attack surface.

Logging should be designed with alerting tightly controlled. All activity should be monitored with a deliberate decision on what is important and how an event is handled. The collection of logs should be done centrally. This is a good mitigation avoiding modification of logs by a malicious insider or someone who has gained access to the SPoC.

All services (LDAP, Authentication, DNS, Time, etc…) provided to the SPoC must be protected solidly. These services should be controlled with firewall policy. There may be a good argument to include some of these services within the out of band network inside with the SPoC.

If a web interface is required for management of a SPoC appropriate controls must be implemented. These include a Web Application Firewall and potentially a Proxy Server. Secure coding should be required but even with solid secure coding processes,

David Belangia, dwbelangia@hotmail.com

commercial software is likely to have vulnerabilities that will need to be mitigated by the firewall or proxy server.

3. Credits

Mike Fisk, Los Alamos National Laboratory Chief Information Officer chartered this working group November 2, 2014.

Special expertise was identified and Subject Matter Experts (SME) were added to the working group. Alex Malin was tasked with leading the technical discussion with the SMEs. Under Alex's leadership, the working group explained each individual SPoC and the approach used to secure the SPoC. The contributors identified here are only a few of the SMEs whose expertise was tapped. The SMEs who supported this effort include: Mike Fisk, Steve McLenithan, Daniel Vollans, Jeff Johnson, Aaron Morrison, Steve Howard, Karl Pommer, Dale Leschnitzer, Georgia Pedicini, Edward Brown, Gregory Lee, Jason Holladay, Paul Brown, Heath Davis, Scott Miller, Bob Knight, David Kennel, Bill Ebanks, Mark Martinez, Jeffrey Click, Kelly Koch, Susan Coulter, Brian Sedlacek, Mark Lorenc, Timothy Hemphill, Ekkehard Koch, Sarah Hooks, Chris Olsen, George Brehm, Randy Cardon, John Parrack, Mathew Wheeler, Dan Walters, and others. The author of this paper wishes to thank this group for their hard work, caring nature, and expertise.

4. Conclusion

Not all assets are created equal. Adversaries have desired targets and the compromise of one of these targets (a SPoC) could mean complete control of the enterprise's environment. Once an adversary owns a SPoC, the enterprise is lost.

The individual protection of each SPoC is important. Ensuring that these assets are patched, configured correctly, and using authentication properly are a good starting point. Understanding normal behavior is important but can seem overwhelming. Using some scripting tools to collect baseline information and scheduling routine comparison of those baselines to operating status for change can provide the ability to understand normal behavior and provide an alert when something is abnormal.

David Belangia, dwbelangia@hotmail.com

Going beyond these expected approach for individual SPoCs and implementing additional controls for the SPoCs as a core group of assets can provide additional value. Move attractive targets into network segregation that will provide the ability to implement additional controls. Minimizing traffic to a subnet or VLAN can simplify auditing of traffic and make alerting easier. Aggregating logs and focusing more attention on what is normal behavior for the SPoC is simplified by reducing the noise.

Requiring alternative and multiple authentication into these network segmentations will provide additional layering of control. The application of a bastion host and the restriction of access to the bastion host will increase the complexity for an adversary and provide assurances that protections will last longer.

If the enterprise can protect the SPoC, the ability to increase the protection time in the TBS Model will provide some hope to mitigate an attack protecting the crown jewels (SPoC) valued by the enterprise, stockholders, and potentially the nation. Plan defense of the SPoC and implement those defenses in a manner that is commensurate with the SPoC's value in protecting the enterprise. Providing additional layered controls by managing the SPoCs as whole will provide increased protection time and support better detection and reaction time.

5. References

Allen, R. (2005). 5 Steps to a Secured Active Directory. *Windows ITPro*. Retrieved from
    http://windowsitpro.com/windows-client/5-steps-secured-active-directory.

Allan, A., Sigh, A., Ahlm, E. (2014, December). Magic Quadrant for User
    Authentication. Retrieved from http://www.safenet-inc.com/authentication-magic-
    quadrant/.

Berkeley Security. (2015). Authenticated Scans Guideline. Retrieved from
    https://security.berkeley.edu/content/authenticated-scans-guideline.

Center for Internet Security. (2015). Security Benchmarks. Retrieved from
    https://benchmarks.cisecurity.org/downloads/.

David Belangia, dwbelangia@hotmail.com

Chandramouli, R. (2014, October). Security Recommendations for Hypervisor

Deployment. *NIST Draft SP 800-125-A*. Retrieved from

http://csrc.nist.gov/publications/drafts/800-125a/sp800-125a_draft.pdf.

ContentSecurity. (2015). Managed File Transfer. Retrieved from

http://www.contentsecurity.com.au/security-solutions/data-protection-

management/managed-file-transfer/.

CyberArk. (2014). Next Generation Jump Servers for Industrial Control Systems.

Retrieved from http://lp.cyberark.com/rs/cyberarksoftware/images/wp-CyberArk-

NextGenJumpServer-5-28-2014-en.pdf.

Dueby, S. (2006). 19 Smart Tips for Securing Active Directory. *TechNet Magazine*.

Retrieved from https://technet.microsoft.com/en-

us/magazine/2006.05.smarttips.aspx.

Engel, G. (2014, November). Deconstructing the Cyber Kill Chain. *Dark Reading*.

Retrieved from http://www.darkreading.com/attacks-breaches/deconstructing-the-

cyber-kill-chain/a/d-id/1317542.

Fisk, M. (2014). Single Point of Compromise WG. Los Alamos National Laboratory

Information Architecture Working Group Formulation Email.

Geier, E. (2014, April). 6 Free Network Vulnerability Scanners. *Network World*.

Retrieved from http://www.networkworld.com/article/2176429/security/6-free-

network-vulnerability-scanners.html.

Gigamon. (2012). Out-of-Band Security Solution. Retrieved from

https://www.gigamon.com/PDF/SolutionOverviews/Security_Out_of_Band.pdf.

Greenwell, R, (2010, July). STIGs, SCAP and Data Metrics. Retrieved from

http://www.disa.mil/news/conferences-and-

events/~/media/Files/DISA/News/Conference/CIF/Briefing/IA_STIG_SCAP_and

_Data_Metrics.pdf.

Griffin, A. (2015). How to Secure a Domain Name Server (DNS). Retrieved from

http://csrc.nist.gov/groups/SMA/fasp/documents/network_security/NISTSecuring

DNS/NISTSecuringDNS.htm.

David Belangia, dwbelangia@hotmail.com

Joyent. (2015, January). Setting up a Bastion Host. Retrieved from
    http://www.darkreading.com/analytics/threat-intelligence/5-ways-to-monitor-dns-
    traffic-for-security-threats/a/d-id/1315868.

Juniper. (2012). Five Best Practices to Protect Your Virtual Environment. *Juniper*
    Networks. Retrieved from http://www.advantel.com/wp-
    content/uploads/2013/11/Five-Best-Practices-to-Protect-your-virtual-
    environment.pdf.

Kavanagh, K., Nicolett, M., Rochford, O. (2014, June). Magic Quadrant for Security
    Information and Event Management. Retrieved from
    https://scadahacker.com/library/Documents/White_Papers/Gartner%20-
    %20Magic%20Quadrant%20for%20SIEM.pdf.

LeClaire, J. (2013, February). Managing Security Risks in a Virtual Environment.
    Lumension A Heat Software Company. Retrieved from
    http://blog.lumension.com/6413/managing-security-risks-in-a-virtual-
    environment/.

Lindros, K, Tittel, E. (2014, September). How to Choose the Best Vulnerability Scanning
    Tool for Your Business. *CIO.* Retrieved from
    http://www.cio.com/article/2683235/security0/how-to-choose-the-best-
    vulnerability-scanning-tool-for-your-business.html.

Mandiant. (2014). 2014 Threat Report. *Trends, Beyond the Breach.* Retrieved from
    https://dl.mandiant.com/EE/library/WP_M-Trends2014_140409.pdf.

Microsoft Download. (2015). How to View or Delete Active Directory Delegated
    Permissions. Retrieved from
    http://social.technet.microsoft.com/wiki/contents/articles/6477.how-to-view-or-
    delete-active-directory-delegated-permissions.aspx.

Maguire, M. (2012, May). 13 Tips to Secure Your Virtual Machine Environment. *State of*
    Security. Retrieved from http://stateofsecurity.com/?p=2471.

NIST. (2006). Guide to Computer Security Log Management Special Publication 800-92.
    Retrieved from http://csrc.nist.gov/publications/nistpubs/800-92/SP800-92.pdf.

NIST. (2015). The United States Government Configuration Baseline (USGCB).
    Retrieved from http://usgcb.nist.gov/usgcb_faq.html.

David Belangia, dwbelangia@hotmail.com

OWASP. (2013). OWASP Top 10 – 2013. *Open Web Application Security Project.*
Retrieved from
https://www.owasp.org/index.php/Top10#OWASP_Top_10_for_2013.

Piscitello, D. (2014, June). Monitor DNS Traffic & You Just Might Catch A RAT. *Dark Reading.* Retrieved from http://www.darkreading.com/attacks-breaches/monitor-dns-traffic-and-you-just-might-catch-a-rat/a/d-id/1269593.

Piscitello, D. (2014, September). 5 Ways to Monitor DNS Traffic for Security Threats.
Dark Reading. Retrieved from http://www.darkreading.com/analytics/threat-intelligence/5-ways-to-monitor-dns-traffic-for-security-threats/a/d-id/1315868.

Schwartau, W. (1998). Time-Based Security Explained: Provable Security Models and
Formulas for the Practitioner and Vendor. *Computers & Security.* Retrieved from
http://ac.els-cdn.com/S0167404898801004/1-s2.0-S0167404898801004-
main.pdf?_tid=fd140b82-f011-11e4-ba7c-
00000aab0f27&acdnat=1430492257_929843032312cb45d55bb7a754f2c708.

Scarfone,K, Hoffman, P. (2009, September). Guidelines on Firewalls and Firewall
Policy. *NIST SP 800-41.* Retrieved from
http://csrc.nist.gov/publications/nistpubs/800-41-Rev1/sp800-41-rev1.pdf.

Scarfone, K., Souppaya, M., Hoffman, P. (2011, January). Guide to Security for Full
Virtualization Technologies. *NIST 800-125.* Retrieved from
http://csrc.nist.gov/publications/nistpubs/800-125/SP800-125-final.pdf.

SecureWorks, (2015). Five Critical Rules for Firewall Management: Lessons from the
Field. Retrieved from http://www.secureworks.com/assets/pdf-
store/articles/five_critical_rules_for_firewall_management.pdf.

Shinder, D. (2001, August). Understanding and selecting authentication methods.
TechRepublic. Retrieved from
http://www.techrepublic.com/article/understanding-and-selecting-authentication-
methods/.

Shinder, D. (2011, April). DNS Security: DNS Security Steps Prior to Deploying
DNSSEC. Retrieved from http://www.windowsecurity.com/articles-
tutorials/windows_server_2008_security/DNS-Security-Part2.html.

David Belangia, dwbelangia@hotmail.com

SWAT. (2013). Securing Web Application Technologies. Retrieved from

http://www.securingthehuman.org/media/resources/planning/STH-poster-winter-2013.pdf.

Vulnerability Scanning. (2015). Webopedia. Retrieved from

http://www.webopedia.com/TERM/V/vulnerability_scanning.html.

David Belangia, dwbelangia@hotmail.com