



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Practical Security Considerations for Managed
Service Provider On-Premise Equipment

GIAC (GSEC) Gold Certification

Author: William Yeatman, wmyeatman@gmail.com

Advisor: Rob VandenBrink

Accepted: September 2015

Abstract

Many organizations are not adequately staffed to perform 24x7 monitoring of network, systems infrastructure, and security activities such as vulnerability scanning and penetration testing. Use of third party managed service provider to fill this gap is on the rise. It is typical for managed service providers to require the implementation of an on premise device or appliance at the customer location(s). But, who watches the watcher? Service providers must be sure to fully harden any on-premise device placed on a customer network, and they must take steps to protect their own infrastructure against the propagation of an attack or compromise of the customer network and systems. Customers must be informed and work closely with service providers to assure proper placement of the on premise device such that it does not become a vector for compromise against the customer network. Collectively, and in accordance with a set of standards and guidelines, all stakeholders involved in the managed services relationship must be sure to set a sustainable benchmark that sufficiently reduces the chances for 3rd party on premise equipment becoming the root, or a contributing cause of a security compromise.

1. Introduction

Information Technology continues to move ahead, full speed, and no prisoners are being taken. At the same time, IT teams at organizations across the globe continue to feel the stretching effect of having to do more with less. There's nothing new about that - it could be said that IT has always been that way. What might be changing, according to CompTIA's Fourth Annual Trends in Managed Services Study, is an increase in willingness for companies to supplement in-house resources with 3rd party vendors to perform common IT functions such as email hosting, network monitoring, customer relationship management (CRM) applications, storage, backup and recovery (CompTIA, 2015).

The Managed Security Services market, specifically, providing remote management and monitoring of IT security functions, is being fueled by growth overall business growth, the continued proliferation of mobile connectivity, social media, increasingly complex network infrastructure, and lack of capital and skilled IT security professionals (Rajput, 2015). Common security services include security log management and monitoring (a 24x7 function), security scanning (very frequently performed during non-business hours), firewall and other security device management services. The 2015 AMR study indicates the outsourcing of security tasks to managed security service providers has emerged as a lucrative option for many organizations, whether onsite, completely controlled by the provider remotely, or a hybrid of onsite and remote.

Regardless of the specific business model, it is quite common for managed service providers to require the placement of monitoring equipment onsite at a customer location or data center. Take security log management, for example. There is usually a need to place log collection server, appliance, or device onsite at the customer premise to consume and do the heavy lifting to normalize thousands, or multiple millions of log events generated by servers, routers, firewalls, endpoints, and other devices every day. Due to the potentially high volume of raw logs, it makes sense in many cases to send to a local collector, rather than consuming precious bandwidth by sending them across a

remote WAN link. A similar approach is frequently taken by service providers that provide monitoring and management of firewalls, Intrusion Detection Systems (IDS), Unified Threat Management (UTM), as well as more traditional network assets such as routers, switches, telephony/VoIP, and other infrastructure.

Another common scenario involving the placement of vendor provided equipment onsite at the customer location is often encountered with Compliance As A Service (CaaS) vendors. According to Mike Mittel, CEO of RapidFire, changes in regulatory requirements, such as HIPAA or PCI, are complicated for customers to stay on top of and that, combined with an increase in public awareness and regulatory enforcement, have added to the push to meet compliance standards (Kuranda, 2015).

Vendors that require the placement of an appliance or device on their customer network are in a unique and critical position. In a very real sense, the bar is held higher for those service providers that are tasked with watching their customers' networks and infrastructure for availability – and even higher when they are specifically responsible for securing those information assets and assuring an organization's compliance with federal, state, and industry regulations.

Every device placed on a network becomes a potential point of entry, pivot, or exit for an attacker. Know your enemy? It's a philosophy not just used by the good guys. The bad guys want to know their enemy too. Knowledge of network infrastructure, and better yet, the ability to control or manipulate it, is often a goal for the attacker – whether that's a final objective or just a stepping stone in a larger initiative. Repeat: every device placed on a network, whether by the owning organization or a service provider acting on behalf of that organization, becomes a potential point of entry, pivot, or exit for an attacker.

It is imperative for on premise service provider equipment to be initially deployed, and to remain in a hardened and trusted state. While industry standard practices should be applied to all information and computing assets, commensurate with risk, the watcher is always held to a higher standard.

But, even for the managed service provider, how much security is enough for on premise equipment and the ancillary systems and interfaces associated with that equipment? The answer is always “it depends”. After diving deeper into some of the scenarios introduced above, and the high risk areas they present to organizations and managed service providers involved, a set of countermeasures and controls will be offered as a solid starting point for any organization that provides onsite equipment as part of their service offering, or as the basis of due diligence review for any organization considering such services.

2. Example MSP Scenarios

As a point of reference for discussion of the risks and countermeasures associated with MSP on premise equipment, the MSP scenarios are explored further. Note that the terms Managed Service Provider and Managed Security Services Provider, and their acronyms (MSP and MSSP, respectively) are used interchangeably since the concept of securing onsite equipment is the same whether the services provided entail general networking and systems, or are focused on information security relevant services.

2.1. Security Log and Event Management and Monitoring

Monitoring and analysis of security events for evil (also frequently referred to as Security Incident and Event Management, or SIEM services) is a 24x7 job that requires expert security analyst skills. The time commitment, talent requirements, and computing resources necessary sifting and storing large volumes of log data, are compelling reasons that make outsourcing of this function to a capable SIEM service ideal for understaffed IT shops. And while cloud logging services exist, it is very common to find, at the enterprise level, these types of service vendors placing log collectors on premise at the customer location.

It is important to consider and understand the connectivity requirements of an onsite log collector. Usually the requirement is that the collector be able to receive logs from a variety of sources on the customer network. The only time these collectors typically need to initiate communications is to relay normalized log information back to the service provider’s systems and/or Security Operations Center (SOC) for further

analysis, or compressed logs offsite for archival. It is also normal for the service provider to have access to the collector in order to perform routine operations, administration, and maintenance. There is usually no reason for the collector to initiate communications to any other device or host on the customer network as it should only be consuming logs from in scope devices on the customer network.

The following diagram depicts a typical deployment scenario for this type of service:

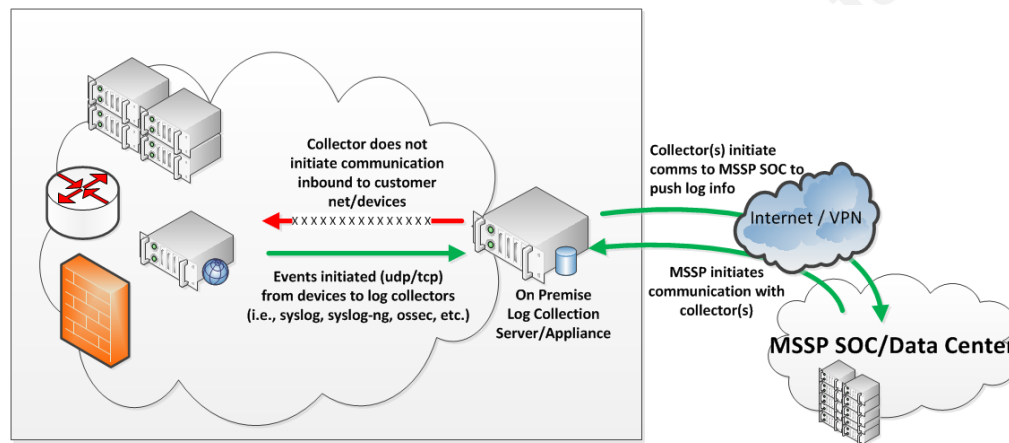


Figure 1. MSSP Log Management/Monitoring Scenario.

As with all things, there are design and architectural exceptions to this model. In some log monitoring deployments, it is possible to have the log collector initiate a connection to the devices being monitored. For example, Splunk can be configured to monitor event log channels and files remotely using Windows Management Interface (WMI), which adds complexity to the scenario we are considering (Splunk Online Documentation, 2015).

2.2. Network and Infrastructure Management and Monitoring

A variety of network and infrastructure monitoring services are available to organizations today to perform a range of services - from simple uptime/availability monitoring and alerting to full-fledged management of network gear, telephony and communications systems, and servers. Similar to the prior log monitoring scenario, MSP's often require the placement of a device on premise at the customer location and

often include log monitoring/management as part of their services, as described in the preceding scenario.

The connectivity and access requirements for these types of services are generally more extensive than what is required in the prior and relatively simple scenario of collecting logs. These types of services often require the on premise device to initiate connections inbound, toward the managed customer network and infrastructure. Additional connectivity must be permitted for protocols such as SNMP (UDP/161) for status polling, SSH (TCP/22) for device management, SMB (TCP/445), RDP (TCP/3389) for management of windows infrastructure, and any number of other ports for performing service discovery and probing for availability that may be part of the managed service offering.

Additional inbound ports to the MSP device must be permitted for service providers that fully manage network devices – for example, SNMP traps (UDP/162) may need to be sent from monitored devices to the on premise collector/device, TFTP (UDP/69) would need to be allowed for the use of TFTP for configuration/firmware image management, and at least one service provider in this space requires TCP/22 and TCP/21 for scp and legacy file transfer capabilities to the on premise management device.

That last point may raise a totally separate issue (encouraging the use of insecure/plaintext protocols, a topic for another paper), but one thing is clear: in this scenario the on premise device requires a considerably higher level of accessibility compared to that of the seemingly simpler log management/SIEM service.

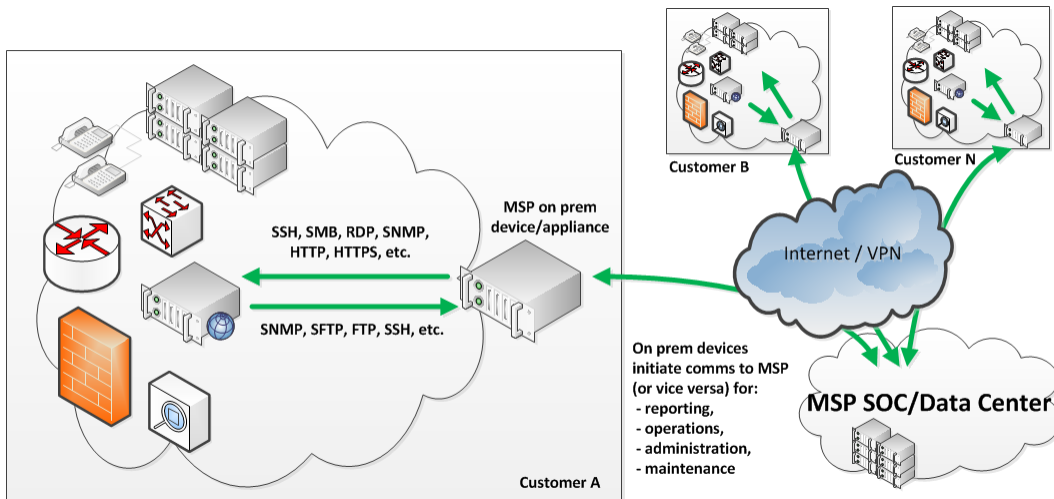


Figure 2. Network/Infrastructure Managed Service Provider

To continue contrasting with the first scenario, these types of services also often require sophisticated toolsets be installed on the on-premise device in order to perform the variety of functions required of such a service, as well as SNMP community strings and login credentials to manage and monitor customer devices. Clearly, the protection of the MSP on-premise devices is incredibly important.

2.3. Compliance As A Service

Another type of service that often requires placement of on-premise equipment is known as Compliance As A Service (CaaS). A popular example is Payment Card Industry (PCI) compliance, where the CaaS vendor places an appliance onsite to manage and perform many, if not most of the operational aspects required by the Data Security Standard (DSS). This includes placing the appliance in or near card data environment in order to perform log collection, vulnerability scanning, penetration testing, filesystem and database scanning for unprotected credit card data, and various other tasks that involve not just network level access, but frequently also requires administrator credentials to execute properly. The following diagram illustrates a potential CaaS arrangement:

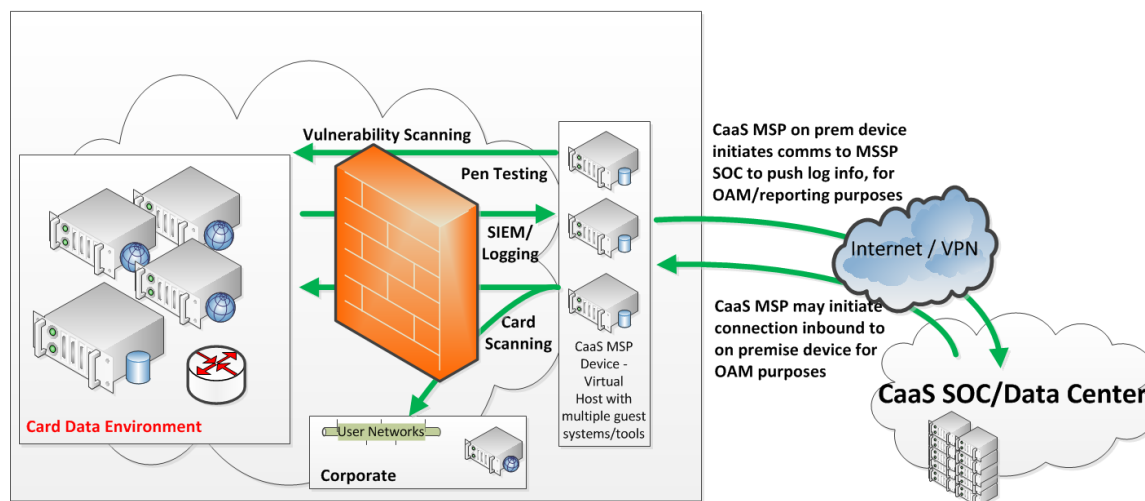


Figure 3. CaaS Scenario.

As we've progressed through these scenarios, additional complexity has been introduced. As the number of functions performed by a Managed Service Provider increases, additional tools and systems are needed to perform the service. What has become increasingly common in the on premise device design is to deploy a virtual host with capacity for multiple virtual guest operating systems and platforms to scale to the growing number of services provided.

Further complexity may be introduced by the fact that the CaaS device will most likely need to initiate communications to systems outside of the card data environment (CDE) in order to assure compliance. For example, card scanning of non-CDE platforms is required in order to assure plaintext card data does not exist outside of the approved CDE. While PCI scoping will not be discussed further, it's an interesting paradox that points to the cross-boundary complexities that must be considered when placing third party owned and operated equipment on customer networks.

It is also a common strategy for MSP's that provide multiple network/server monitoring and management services, as indicated in the second scenario, to employ virtualization based architecture as is describe here for CaaS providers. In any case, the on premise devices and platforms, whether bare metal or virtual, must be properly secured and remain in a trusted state.

3. Risks

Building upon our understanding of how MSP equipment might likely be placed on premise at a customer site, we can further explore the risks and undesirable events that could come about as a result of this relationship.

3.1. Risks to the MSP

In the scenarios presented, the on premise device usually has the ability to connect back to the MSP central management systems. Often this is a persistent connection over a VPN established across the Internet between the device and the MSP operational center. MSP's take this approach since it helps reduce the overhead and administration of the customer needing to open/allow inbound access to the device for operations, administration, and maintenance (OAM) purposes by the MSP. In other cases the device may simply phone home periodically to retrieve job related instructions, or there may be instances where the device is placed in a customer public DMZ and is reachable from the Internet (this isn't an ideal placement in terms of potential for security exposure as we'll discuss in section 4.; however, in such cases, a host based firewall should be implemented on the device to only allow connections from the MSP operations center, with the customer applying similar rules at their border firewall as a secondary layer of defense). Regardless, the device, per se, serves as a channel back into the MSP systems, and one can see how infiltration of an on premise device could serve as a strategic foothold and pivot point for gaining entry to the MSP systems.

Taken a step further, if an adversary were to compromise an on premise device as a stepping stone to infiltrate the MSP central management systems, s/he may likely gain access to not just the MSP, but potentially access to information about other customers -- if not possibly gaining access to the customer networks and systems under the care of the MSP! This is a very plausible tactic that could be used by a nation state or other well-funded adversary - becoming a customer of an MSP in order to seek out new targets, gain information about an existing target, or seeking to infiltrate, destabilize, or render the target inoperable. And while this does bring to light the need to secure and harden the central MSP SOC platform and management structures, properly securing and hardening

each device deployed in the field serves as the front line of defense against pivoting inward toward the MSP and its other customers.

A compromise of an MSP's onsite device could cause, to a large extent, reputational damage. After all, especially for managed security service providers, the inability to defend your own house connotes the inability to defend others' houses. The old saying "the cobbler's children have no shoes" does not apply here.

3.2. Risks to the Customer

Similar to the MSP's need to address the risks of a compromised onsite device becoming a pivot point into its backend management systems, the customer organization should be addressing the risk of the on premise device being used for pivoting further into and potentially exploiting its network and infrastructure. This is complicated by the notion that wide accessibility is an inherent requirement of many service offerings – the device has to communicate and have access in order to do its job. So, the on premise device is frequently receiving information from and about potentially many devices in the customer network and, perhaps unbeknownst to the organization, serving as a treasure map for an adversary seeking to know more about the customer's assets.

By virtue of the requirement to take in logs, data, and establish connections with many sources within the customer network, a vulnerable service running on the collector (i.e., vulnerable syslog daemon), a misconfiguration, or improper access controls, could be exploited by an adversary that has gained access to the customer network either physically or logically. Unsecured consoles that require physical access to the device could be exploited by a malicious insider to gain unauthorized access to it.

Clearly, the customer has a vested interest in assuring the proper hardening of the MSP's onsite device. The earlier analogy of protecting the MSP's central systems by way of securing the devices as first line of defense applies here as well, just in reverse. Harden and make sure the device stays that way in order to protect the customer.

An added concern that customers should have is the trustworthiness of staff working for the MSP. By nature of the information to which they are exposed, and the privileges and access that they may have in order to perform their job duties, MSP staff

are in a position of trust, yet this often carries the potential for misuse. Arguably, abuse by a malicious staffer at an MSP providing read-only log monitoring services may be less of a concern than one that requires staff to have access to manage network devices and servers.

4. Mitigating Controls and Countermeasures

Now that we have discussed reference scenarios and some of the risks germane to the MSP – customer relationship, it's time to let the sun come out and shine and discuss ways that will, if executed and sustained with dedication and focus, significantly reduce the adversary's chances of compromising the on premise device, the MSP central monitoring and operations platform, and the customer's network.

Many of the controls prescribed and described are taken directly from the Council on Cybersecurity's Top 20 Critical Security Controls (CSC Top 20) and are referenced accordingly.

4.1. Contractual Considerations

While some service provider arrangements allow for customers to have access to on premise equipment, this is a practice that should be permitted only under extenuating circumstances and read-only / limited privileges granted. Contractual terms should be included to assure the installation of the on premise equipment in a physically secure and locked rack. Service Level Agreements (SLA's) should be in place governing reconstitution/replacement of an onsite device in case it was to become compromised.

4.2. Physical Device Hardening

In accordance with CSC 3-1, MSP on premise hardware should have standard, secure configurations applied to them. The following controls and countermeasures set the bar for preventing unauthorized access to the device by an individual with physical access. This is critically important given the diversity of untrusted environments to which the MSP may be shipping and deploying their devices.

Console access must be configured to require authentication using either two-factor authentication or a sufficiently complex passphrase (CSC 10-4, 12-4). Because

console access is often limited to a static password, enforce per-customer unique passphrases that are tightly controlled using a password management system (i.e., residing in the MSP SOC/data center), access to which requires two-factor authentication.

All unused media interfaces, including USB, CD-ROM, unused network interface ports (including wireless), must be disabled (CSC 3-1, 7-5).

The MSP must enable and monitor the device for chassis related events (case open), system reboots, network port state changes, local console logins (failed and successful), and reconcile all detected state changes against an approved change/maintenance/support request daily (CSC 14-5). Any irreconcilable event must be escalated for immediate and formal incident response, and possibly forced shutdown (aka “ejection button”) of the device (CSC 18-1, 18-4). While the information processing tasks that many on premise devices perform are varied and diverse, the physical state of the device itself would not typically be subject to frequent changes, so monitoring for, detection of, and response to such events should be considered critical and responded/resolved immediately.

Full disk or volume level encryption using publicly vetted algorithms should be implemented (CSC 17-1, 17-2). In some cases, full disk encryption may be difficult to deploy in an operational environment on these types of platforms, though it does depend on the specific platform. In these cases, the design should consider, at minimum, the use of folder/file level encryption to protect sensitive information collected by the device.

The MSP must document the standard hardening configuration and update it as changes are made (CSC 3-1).

Finally, on devices that utilize a virtualized architecture/design, any local access to the hypervisor would need to be treated and controlled in the same manner as local console/shell access described in this example.

4.3. Logical Network Placement and Access Control

Although specific conditions may warrant otherwise, MSP devices should be placed on a separate internal and isolated segment, controlled by firewall rules that explicitly allow the necessary communications with both the customer’s devices under

their management, and the MSP central operations center. Because MSP devices usually contain sensitive data about the customer network, placement on a DMZ or public network (CSC 19-1, 19-4) should be avoided where possible.

From the perspective of the customer, isolating the MSP device on a separate network and implementing network level firewall rules is necessary. By allow only the communications necessary for providing the service, there is a reduction in opportunity for misuse of the device, and if it were compromised, a reduction in its ability to pivot or for exfiltration of sensitive information. It is worth reiterating that this isn't a stopgap measure given the often wide accessibility required by MSP devices, and it is even more complex if the MSP is the one controlling and managing the customer network firewalls! Additional controls around change/access management help to mitigate these and are discussed later, though it is worth mentioning that not putting all of an organization's eggs in a single one-stop shop MSP basket is a strong consideration to be made.

In the second scenario introduced - network/infrastructure management and monitoring services - the MSP may be required to have ongoing access to the customer environment in order to fulfill the obligations of the service contract. However, contrast this with the CaaS scenario. Activities like vulnerability scanning and filesystem/database scanning for plaintext card data are typically scheduled activities. In this arrangement, the customer can schedule activation of the firewall rules necessary for the duration of the scan activity. When completed, the pertinent firewall rules are deactivated. This lessens the likelihood of a compromised CaaS MSP device being used to further pivot into the customer network by either an unauthorized attacker or authorized, but malicious, insider.

From the perspective of the MSP, customer firewall rules should serve as a second layer of network access control in protecting the MSP device from potentially malicious traffic. As a first line of defense, MSP devices must implement host based firewalls that enforce a strict deny-by-default policy (CSC 11-2). On the interfaces that provide the services to the customer, configure rules that allow only the specific IP versions (4/6), protocols, ports, and source/destinations necessary to provide the services that have been contracted to be performed, dropping all other packets. For management

of the MSP device itself, allow inbound/outbound rules on a physically separate management interface that permits communications with the MSP central management platforms, deny all else, including all other outbound Internet connectivity. All MSP device updates should be controlled and distributed centrally from within the MSP operations center.

4.4. Authentication and Access Control

Standing access to the MSP device and any customer systems by MSP operators/staff should not be permitted (Perry, 2014). Access to the MSP device (and subsequently to any customer systems) is temporarily granted only when tied to a specific and approved service request and/or maintenance/change ticket. By implementing approved and controlled windows for when MSP staff can access the on premise device and customer assets, the chances of misuse of the MSP device and/or customer systems, by a malicious or miscreant MSP staff member, can be significantly reduced. Periodic, sample based reviews of access to the devices by an internal security team, separate from the core MSP operations, as well review of actions performed while logged in, will also serve as a deterrent for misuse and is an additional layer of detective control for identifying other anomalous patterns (CSC 12-1, 16-1, 16-3, 16-5, and 16-11).

Access to the MSP devices and virtual systems/platforms contained therein, should utilize two-factor and/or mutual certificate/key based authentication (CSC 10-4) and certificate/key management practices and procedures must be formalized and implemented.

For environments where static passwords must be used, use a sufficiently complex and long password (20-30 character length) that is unique for each MSP device. These passwords can be centrally managed within the MSP central operations platform/systems, but access to the password repository itself should require two-factor authentication in order to retrieve static passphrases. Rotation/changing of static passwords should be performed periodically, ideally every 90 days.

The value of two-factor authentication cannot be overemphasized. In a formal alert posted by managed security services firm Dell SecureWorks Counter Threat Unit

(CTU), attackers are frequently using compromised credentials coupled with legitimate administrative tools to move laterally and horizontally within infiltrated environments (Dell, 2015). The use of two-factor authentication for remote access and for the management of network devices helps to remove credential (password) theft from the adversary's attack arsenal.

Direct access using the root, administrator, or other highly privileged account on the MSP device, over a network (and locally if possible), must be disabled. Legitimate connections to the MSP device must use a general user login and then elevate (su, runas, enable, etc) to root/admin to perform functions that require privileged access (CSC 3-1, 12-1).

While general MSP operators and staff may be conditionally permitted access to the MSP device to utilize the permitted toolsets and to perform their assigned duties, they should not be permitted access to directly manage and administer the MSP device (i.e., making changes to installed tools, host based firewalls, etc.). This function should be limited to a dedicated engineering group within the MSP and all changes must go through a formal change control process that involves both the MSP and the customer(s) (CSC 3-1).

4.5. Vulnerability and Security Posture Scanning

From the MSP central management platform, perform weekly authenticated port scans, vulnerability scans, and CIS benchmarking against the devices in the field. For virtualized architectures, perform against hypervisor and all guest operating systems in use (CSC 3-1, 4-1, 11-3). All medium or higher findings should be patched/remediated in a test environment and then pushed to the production devices once verified (CSC 3-2).

4.6. Application Whitelisting, Process Monitoring, FIM

In the scenario where an attacker or malicious insider were successful in bypassing the aforementioned controls on the hardened MSP device (e.g., via 0-day exploit, accidental misconfiguration), application whitelisting should be deployed to prevent execution of programs/binaries that would aid the attacker in establishing

backdoors/persistence, pivoting further, and exfiltration of data from the MSP device and any virtual guests that may be hosted within it.

Additionally, implement process and network socket/connection logging, file integrity monitoring, and have the internal MSP security team aggressively monitor to detect unexpected programs/processes that start or die on the MSP device and for any unusual changes to important files. Any unexpected network connections, process creation, or changes to files that cannot be directly tied to a specific support, maintenance, or change request ticket must be immediately investigated and resolved, up to and including disconnection of the device to prevent further propagation of the attacker.

5. Additional Considerations

The items discussed so far have centered on the security and hardening of the MSP device residing onsite at a customer location. However, there are three related and relevant issues that should be covered in any conversation regarding the security of a managed service provider – credential reuse, portals, and costs.

5.1. Credential Reuse

Section 4.4 discussed that, when two-factor authentication isn't doable, passwords of sufficient length and strength, unique to the individual customer device, could be used. This same concept applies to any instance where a static password is used in the course of providing the service. For example, a network infrastructure monitoring and management company may need to poll devices using SNMP. The community string used for polling one customer's network devices should never be the same as the community string for a different customer. Mechanisms and solutions are available for MSP's to securely manage these types of static credentials. It is strongly mentioned that two-factor authentication and auditing be enforced for access to any repository of customer related credentials.

5.2. MSP Portals

It is very common for MSP's to provide customers a web based portal that allows a view into the customer's managed network/devices, statistics, and other account information. These portals are tied very closely to the underlying systems that communicate with and receive information from the devices deployed on the customer network. Similar to the notion of on premise customer devices serving as a treasure map for any single customer's networked information assets, the portal serves many customers and could be viewed as a treasure map of treasure maps! Improper access control and management, failure to identify and remediate portal vulnerabilities, and failure to monitor portal activity could increase the risk of an adversary or unauthorized individual accessing managed services customer information.

While the Top 20 Critical Security Controls should be applied in earnest and full to a customer portal, there are three controls that are an absolute must for protecting these. The first is two-factor authentication. Given the focus on eliciting password credentials from unsuspecting users, or via other means, the use of one-time passwords (OTP), client certificates, or other multi-factor authentication method is an absolute must for gaining access to portal systems. Further, when portals do not require wide accessibility via the Internet, access to them should be whitelisted by source IP address or via VPN connection.

Next, MSP portal applications must be tested initially and before deployment of new releases for web application flaws (CSC 6). Portals should be rigorously tested using static, dynamic, and manual penetration testing methodologies, seeking to find and remediate all types of web application vulnerabilities (see OWASP for additional information). There should be some focus also on prevention of framing/forgery attacks, cookie replay (vis-à-vis XSS attack), and others that involve an element of social engineering. Internal security monitoring and response teams should be prepared to respond to any reports of phishing/waterhole types of attacks that could indicate targeted efforts to compromise credentials for accessing the portal. Web application testing should also be combined with protection by a Web Application Firewall that is actively monitored and tuned to defend against the latest attack methods (CSC 6-2).

Finally, MSP's must actively review portal access. In many cases, management of the portal access is often delegated to a central point of contact within the customer organization. However, the MSP should, as part of the customer engagement and ongoing service delivery process, provide oversight in this area by facilitating quarterly access reviews.

5.3. Security as a Cost Factor and Sales Tool

Talking about the need for security is easy, especially in this day in age with its cyberwar and rumors of cyberwar. Actually providing and sustaining security in the face of getting a product or service to market is another thing. But there is a balance point in between lip service and "gung ho" – a theme that this paper has attempted to convey. Providing reference scenarios, highlighting some of the more likely risks, and then drawing upon a vetted and prioritized set of "Quick Wins" that are considered most likely to prevent, detect, and halt an infiltration, offers MSPs a starting point for assuring their on premise devices remain whole.

MSP's are held to a higher standard, especially managed security services providers. The onus is on them to assure the trustworthiness of devices deployed onsite at customer locations. As such, their business and pricing models must reflect the ongoing requirements to sustain this higher degree of trustworthiness. It's just a cost of doing business in that space. With more buzz in the industry around frameworks like NIST Cybersecurity, the Top 20 Security Controls, ISO compliance, and the like, MSP's can and should leverage the controls identified herein and use those as selling points for their services. They need to understand the costs of doing security well versus the cost of an actual security incident or data breach – whether that directly impacts just the MSP, a single customer, or multiple customers – and parlay that into the service model and sales collateral.

6. Conclusion

Organizations are seeing an increase in willingness to outsource various aspects of their IT and information security functions to managed service providers. While onsite, cloud, and hybrid arrangements exist, it is very common for a service provider to

require placement of a server, appliance, or other device on premise at a customer location and to manage it from a central operations center. There are many types of service provider arrangements - three of which were introduced and described. The most concerning risks are those that incur MSP reputation/brand damage, breach of customer network or other sensitive information, and intermediate concerns such as credential theft, insider misuse, and lateral/horizontal movement by an attacker into either the customer or MSP systems.

The managed service provider device must be properly secured and remain in a trusted state since it will frequently have access to a vast array of customer assets. The Council on Cybersecurity Top 20 Critical Security Controls, particularly the “Quick Win” items, can be used as the basis for securing the on premise device.

Of particular importance in this scenario are physical device hardening, disabling of unused physical interfaces and logical services, device placement (in a locked cage and on a separate/isolated subnet), implementation of network and host based firewall rules, enforcing two-factor authentication, implementing a no-standing access policy/procedures, formally associating all access/changes a valid and approved support/request/maintenance ticket, application whitelisting, and other focused host event monitoring. Additional considerations around virtualization must also be made since each virtual guest in an MSP device represents a platform in need of hardening. The CSC Top 20 framework contains additional controls which should be reviewed and implemented in accordance with the needs of the MSP.

Focused and sustained devotion to maintaining the MSP device security posture, in conjunction with the proper hardening of MSP central monitoring systems, MSP staff workstations/environments, and customer portals can and should be factored in to the cost of the services provided. The cost of providing a verifiably secure service can and should be factored in to the business/service model and contrasted to the cost of a security incident or breach, and the potential resultant harm.

References

- Clarke, P. (2014, May). From Inside the Cloud: Who has access to your data within Office 365? Retrieved from <https://blogs.office.com/2014/06/02/from-inside-the-cloud-who-has-access-to-your-data-within-office-365/>
- CompTIA. (2015, June 1). Customer Awareness and Adoption of Managed Services Trending Upward, New CompTIA Research Reveals. Retrieved from <https://www.comptia.org/about-us/newsroom/press-releases/2015/06/01/customer-awareness-and-adoption-of-managed-services-trending-upward-new-comptia-research-reveals>
- Council on Cybersecurity. (2014, October). Council on CyberSecurity Critical Security Controls for Effective Cyber Defense. Retrieved from <http://www.counciloncybersecurity.org/critical-controls/>
- Dell SecureWorks CTU. (2015, September). Dell SecureWorks Warns Organizations of Hackers Using Little or No Malware to Breach their Targets. Retrieved from <http://www.secureworks.com/assets/pdf-store/other/Media-Alert-09022015.pdf>
- Kuranda, S. (2015, January). 2015's Big Opportunity For MSPs: Compliance-As-A-Service. Retrieved from <http://www.crn.com/news/managed-services/300075263/2015s-big-opportunity-for-msps-compliance-as-a-service.htm>
- Rajput, N. (2015, April). World Managed Security Services Market - Opportunities and Forecasts, 2013 - 2020. Retrieved from www.alliedmarketresearch.com/managed-security-services-market

Splunk. (2015). Monitor Windows event log data. Retrieved from

<http://docs.splunk.com/Documentation/Splunk/6.2.5/Data/Monitorwindowsdata>

© 2015 SANS Institute, Author retains full rights.