Simple Approach to Access Control: Port Control and MAC Filtering

GIAC (GCCC) Gold Certification

Author: Bill Knaffl, bill.knaffl@gmail.com
Advisor: Adam Kliarsky

Accepted: August 21, 2016

Template Version September 2014

Abstract

Many times businesses will spend time and money on "Magic Bullet" security and focus on a single technology or threat. This focus can lend itself more towards placing a "check in the box" for compliance rather than on actual security and facing today's threats. Frequently, missing controls can have a cascading effect where because one control was missing or inadequate, other failures occur turning a minor problem into a breach. This paper approaches one such incident, calls out which control was identified as the primary failure and offers an evaluation of a specific tool that could have helped prevent this attack. It covers not only the cost of the tool and the time to implement but discusses other costs such as training, monitoring, maintenance, user impact and offers a guide for a successful implementation.

1. Introduction

While it is true that there are hundreds of Security Frameworks offering thousands of controls designed to help ensure that any particular network is compliant, most of these focus on compliance rather than security for known attack vectors. For instance, the National Institute of Standards and Technology (NIST) Security and Privacy Controls for Federal Information Systems and Organizations (NIST Special Publication 800-53 rev 4) offers 170 controls covering various control families. Many of these high-level controls are then decomposed into even more granular protections based upon the sensitivity of the protected data (National Institute of Standards and Technology, n.d.). Another example, ISO 27002 has “35 control objectives (one per ’security control category’)” with “114 defined controls” ("ISO/IEC 27002 code of practice," n.d.). Regardless of the framework chosen, there are controls designed to bolster the overall security of a system. Many times this focus on meeting the control objective leads to compliance rather than a focus on protecting current attack information.

The approach taken by the Center for Internet Security (CIS) is that the controls from these common frameworks are:

> “… part of a comprehensive risk management framework for USG Agencies, which specifies a full life cycle of security categorization, design and implementation, assessment, authorization, and monitoring. NIST 800-53 is then the starting point for an Agency to select the CIS Controls needed to manage the assessed risk to their information systems” ("Center for Internet Security," n.d.).

By comparing the controls to known attack risk, one is able to develop a prioritized approach to the implementation of controls. Not only does this focus on risk (rather than compliance alone) but it also allows an organization to approach these risks in an iterative process.

1.1. Brief History

In 2008, the United States Government started working on “prioritizing the myriad security controls that were available for cybersecurity” ("SANS Institute - CIS Critical Security Controls: A Brief History," n.d.). From the same article, the goal of this work was to understand known attacks, assign prioritization and then create or document controls to protect against these known attacks. While the initial work was limited to selective participants, the output included agencies worldwide. In late 2008, “the Center for Strategic and International Studies (CSIS) published the CIS Controls for the first time” ("Center for Internet Security," n.d.). Since that time, the Critical Security Controls (CSC) have undergone updates and modifications, with each version published for comment and update. The most recent version (version 6) of these controls is based on protecting against known attacks and is vendor/platform agnostic.

1.2. CIS Breakdown of Controls

The 20 Critical Controls are grouped into three categories: System (controls 1-10), Network (controls 11-15), and Application (controls 16-20) (Taralla, n.d., p. 1-20). The prioritization of these controls allows a great deal of latitude for customization. It is important to note that there is no requirement to implement all the controls nor any specific control at a given time. The goal is to make an informed decision for prioritization based upon the specific risk as perceived by that environment. Since the controls are based upon known attacks and attack vectors, even addressing one control can influence the security of the overall system. Often the determining factor is a cost versus risk comparison. For instance, if a particular threat is of little consequence the control may be delayed or not used altogether. This customization is one of the benefits of using the controls as many standard frameworks. By organizing the controls into subsections, and then selecting the sections and controls pertinent to the business, the security professional can take a layered approach to the increase in security posture.

2. Case Background / Root Cause

The purpose of this paper is to review the ACME case discussed “Case Study: How CIS Controls can limit the cascading failures during an attack” (Knaffl, 2016). That paper discussed how the failures within multiple controls contributed to the overall compromise. While there were numerous issues found during the root cause analysis, the primary failure occurred within the first CIS CSC: “Control 1 - Inventory of Authorized and Unauthorized Devices.” The first failure was that there was no access control for the devices on the network, and the unauthorized access point (AP) went undetected. Later, the infected laptop went undetected as it connected to the AP. When the laptop scanned and collected files from a forgotten FTP server, both the FTP server and the data connections were also undetected. It was only during the exfiltration of data were the above items identified. Ultimately, the failure to detect the rogue system led to a substantial compromise of the network.

The SANS course material for the SEC566: Implementing and Auditing the Critical Security Controls - In-Depth states:

> “Any time a new device is installed on a network, the risks of exposing the network to unknown vulnerabilities or hampering its operation are present. Malicious code can take advantage of new hardware that is not configured and patched with appropriate security updates at the time of installation” ("SANS Critical security controls training course I 20 critical controls I SEC566," n.d.).

This held true for this particular case in that when a user connected an infected laptop to the unauthorized wireless access point (AP). The sub controls that failed included 1.1, 1.2, 1.3, 1.4 and 1.5. While there were multiple controls that failed, the overall issue was that the network protection (as described within CSC #1) failed to detect or react to an unauthorized network device.

Sub control 1.1 calls for an automated tool to scan the network and collect a preliminary list of assets. From that, an automated process (1.4) will continuously monitor the network for new devices not already recorded. Neither the scan from control 1.1 nor the automated processes defined in control 1.4 were present. Using Dynamic Host Control Protocol (DHCP) as described in sub control 1.2 makes it easy to deploy machines but leaves the network vulnerable should controls 1.1/1.4 fail. Sub control 1.3 defines a change in the acquisition process allowing the device to connect to the network. Since none of the previous controls were in place, this was not possible. Finally, sub control 1.5 was not present as there was no 802.1x deployment. Any one of the above sub controls could have been a substantial roadblock to the attack.

3. Architectural Solution

The ACME case leadership determined that there were gaps in their security posture that needed remediation. Management authorized the removal of the unauthorized Wireless AP, as well as the suspension of DHCP services for that network segment (Knaffl, 2016). Given that the lab network in which this event took place has machines that are in constant movement, static IP address was only an immediate solution. While the corrective controls did address the issues, the impact was significant. The future solution needed to address several key requirements:

- Offers DHCP without losing control of the devices making the connection.
- Allows automated processes to shut off a port in case of an unauthorized connection.
- Provides notifications when an unauthorized device is connected and port deactivated.
- Does not require a complete replacement of current networking solutions.
- To accommodate the open source directive from management, the product needs to be open source in nature.

After meeting with Lab Operations and the Information Security team, the consensus was that the best solution would combine port security with standard Media Access Control (MAC) filtering. As there are 16 switch ports and multiple possible devices, the configuration will be dynamic, based upon the operation. In this base configuration, port one is the uplink. Six of the devices are constantly in use and therefore have assigned ports via static MAC filter; these will normally remain unchanged. There are times in which data processing gear from various external areas is on long-term loan thus there will be three ports allotted for this purpose. These will have port security turned on and configured as needed. Port numbers eleven and twelve shall be configured to have port security and ports thirteen through sixteen with MAC filtering. Apparently, this configuration was the best fit for the operations team and could meet a minimal control set for the Information Security team.

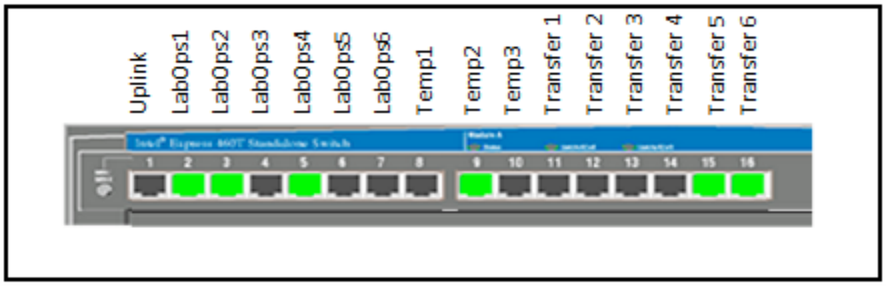

Figure 1 - Planned Port Assignment

3.1. Solution Information

3.1.1. Prerequisites

For the most part, most modern switches support some form of Port Security. For this configuration to meet the requirements, a managed switch allowing port security and mac filtering at the same time is required. In addition to the switch, other requirements include; a full list of the computers and the associated MAC addresses, and the personnel to apply the rules, monitors the devices and report on port changes.

3.1.2. Financial Impact

As there is usually part of any modern network device, the cost for the tool itself is minimal. Many times, the cost changes due to the labor to implement the settings, maintaining the authorized MAC lists, and research when a switch port is set to “disabled” due to an unauthorized connection. In a large network area, this could be potentially a very high cost. When ports are set to “disabled”, the time lost waiting for the port to be reset could also be a factor as well.

3.1.3. Risk Evaluation

The goal of locking down access as defined by the operations team greatly reduced the overall risk this segment added to the network. It also allowed the operations to continue without a complete redesign of the architecture and rewrite of the processes and procedures. Another risk reduction was that a complete operations review was required. This analysis identified inconsistencies in process and procedure allowing for a quick resolution. One weakness of any solution involving MAC addresses is MAC Spoofing. This solution does not address the willful misconduct and attempts to spoof mac addresses of known approved devices. IT Management also informed the Lab

Managers that such tampering is a security violation with consequences that can include actions up to and including termination of employment.

3.1.4. Metric opportunities

Since this is a manual process, there are not many automatic metrics. The only automated process is the shutdown of a port when an unauthorized device connects. One could argue that the time to deactivate is a metric, but that might be the only automated process, and even that does not have automatic tracking. Trouble tickets are one way for a metric collection. Using trend analysis to measure the number of tickets (opened for a deactivated port change) over time could indicate a change in behavior. Another method is to implement port security and run daily reports for SNMP traffic. This option would be able to generate daily reports indicating that a there was a port deactivation. Finally, ticket calls will be required to engage the network engineering team. These calls are required in order to configure new or unknown equipment without putting the rest of the network at risk.

3.2. Port Security

This particular area has many devices deployed to the equipment on the tarmac; this equipment may not always be on the network but can move back and forth between storage, platform and lab network. Essentially these network devices are sometimes portable computers, other times they could be specific hardware for mass storage used for the transfer of data to and from test and production systems. After the initial transfer of data completes, the device is recycled and put back to storage or transferred to another technician and for use in a different subsystem. The decision to use port security reduces the impact of DHCP removal and the meets the overall Concept of Operation (CONOPS) for the lab network operations. Not only does static IP Addressing require administrative rights on the target system, but also a system change every time the device relocates. That will, in turn, drive labor costs. While this is not a perfect solution, as some of the requirements remain unmet, this did answer the immediate operational needs of the lab and positioned to address the unauthorized access with little interruption in service. This solution also allows DHCP services and allows time for the implementation of a more robust solution. In addition to standard security was to be a process to scan the switch hourly and store these reports for the use in metric generation and research upon violation conditions.

3.2.1. Port Security Defined

The goal of port security is to limit the hosts that can connect to the device. Port security “provides the ability to limit what addresses will be allowed to send traffic on individual switch ports within the switched network” ("Switchport Security Configuration," n.d.). While the concept of Port Security is simple, many vendors have created their own implementation, processes, and procedures. When using port security, as a device physically connects to the network port, the MAC address is compared to the list of allowed MAC Addresses. If the device is allowed, the connection is permitted. If the device is not allowed, then usually the port is set to “disabled” and the connection terminated. Again, various manufacturers have different terminology defining the operations that the equipment can perform, but most equipment can support standard MAC Address filtering.

3.2.2. Enable and Configure Port Security

In general, one can configure this type of switch from either a command line environment or a web interface. The instructions will show the web interface and for the sake of brevity, will only address turning on port security for a single port. For the purposes of the desired configuration, Port security is to be enabled to ports eight through twelve.

Start by connecting to the switch using a browser. In this case, the switch is located at 192.168.1.222 and enter the credentials for the device.

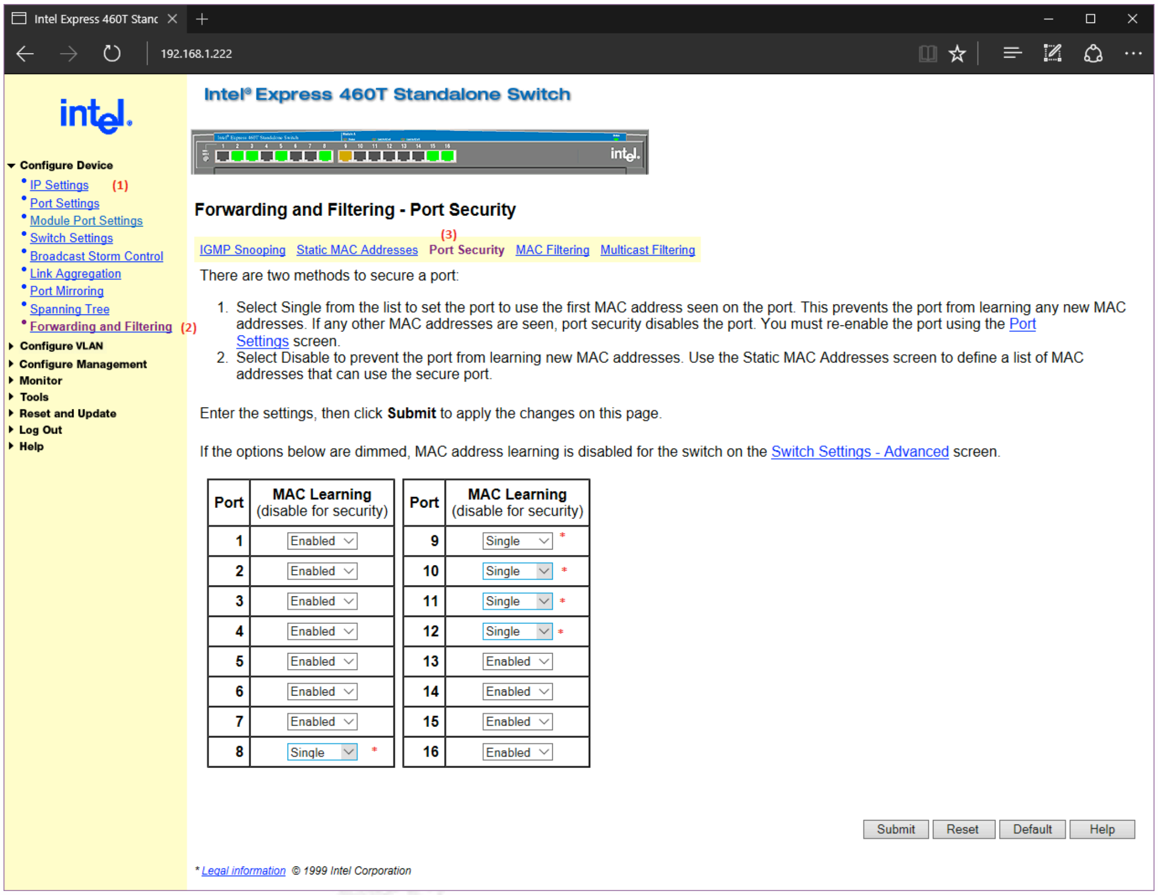

Intel® Express 460T Standalone Switch

- Configure Device
  - IP Settings (1)
  - Port Settings
  - Module Port Settings
  - Switch Settings
  - Broadcast Storm Control
  - Link Aggregation
  - Port Mirroring
  - Spanning Tree
  - Forwarding and Filtering (2)
- Configure VLAN
- Configure Management
- Monitor
- Tools
- Reset and Update
- Log Out
- Help

Forwarding and Filtering - Port Security

IGMP Snooping | Static MAC Addresses | (3) Port Security | MAC Filtering | Multicast Filtering

There are two methods to secure a port:

1. Select Single from the list to set the port to use the first MAC address seen on the port. This prevents the port from learning any new MAC addresses. If any other MAC addresses are seen, port security disables the port. You must re-enable the port using the Port Settings screen.
2. Select Disable to prevent the port from learning new MAC addresses. Use the Static MAC Addresses screen to define a list of MAC addresses that can use the secure port.

Enter the settings, then click **Submit** to apply the changes on this page.

If the options below are dimmed, MAC address learning is disabled for the switch on the Switch Settings - Advanced screen.

Port	MAC Learning (disable for security)	Port	MAC Learning (disable for security)
1	Enabled	9	Single *
2	Enabled	10	Single *
3	Enabled	11	Single *
4	Enabled	12	Single *
5	Enabled	13	Enabled
6	Enabled	14	Enabled
7	Enabled	15	Enabled
8	Single *	16	Enabled

* *Legal information* © *1999 Intel Corporation*

Figure 2 - Switch Configuration - Port Security

To configure Port Security within this device, click on **Configure Device** (1) , then **Forwarding and Filtering**(2), then on **Port Security** (3). From the screen capture above it can be seen that ports 8-12 are configured to allow the first device connected (assigned with assistance from the Network Engineer). Any device that is plugged in will cause the port to error and disable the port until turned back on manually.

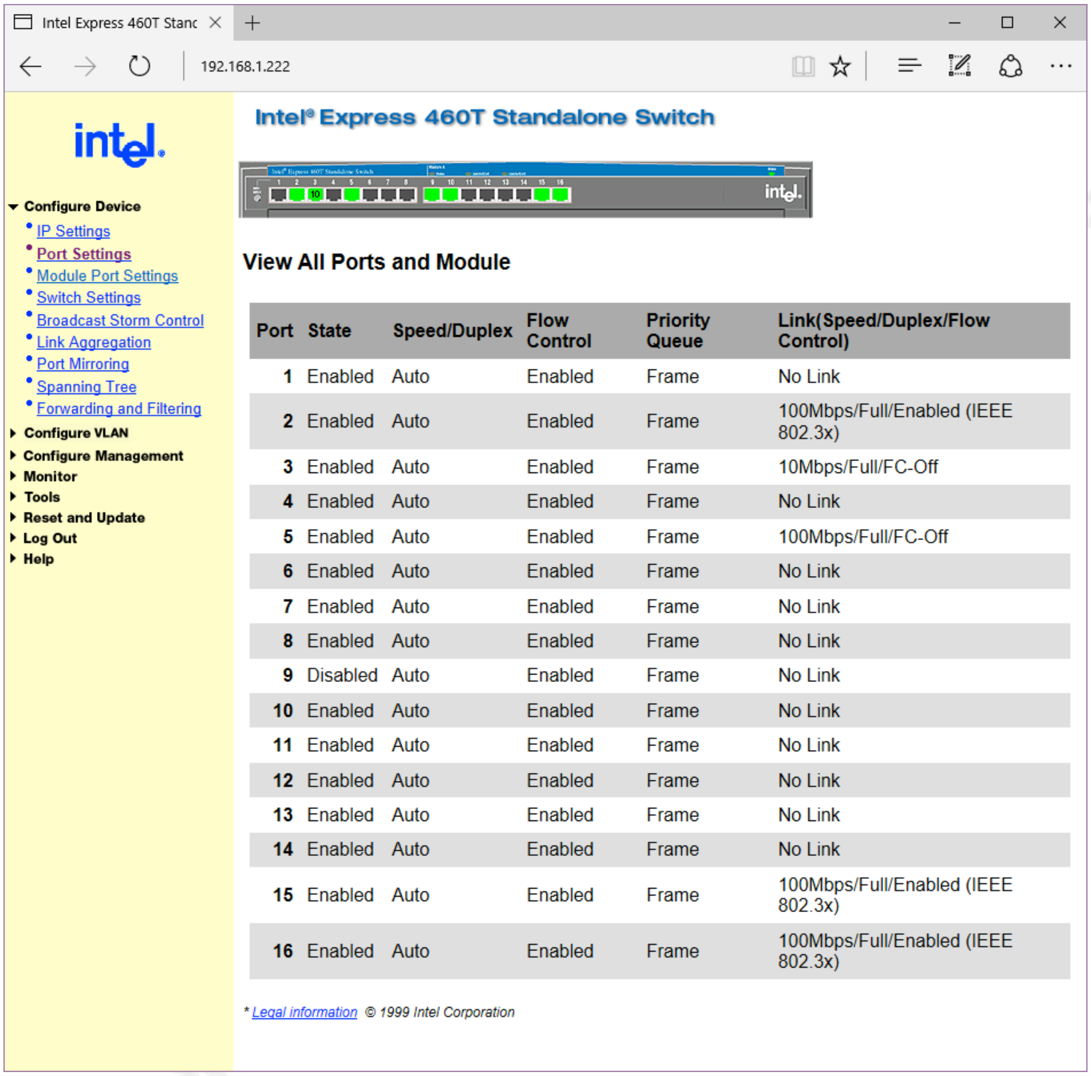

Intel® Express 460T Standalone Switch

View All Ports and Module

Port	State	Speed/Duplex	Flow Control	Priority Queue	Link(Speed/Duplex/Flow Control)
1	Enabled	Auto	Enabled	Frame	No Link
2	Enabled	Auto	Enabled	Frame	100Mbps/Full/Enabled (IEEE 802.3x)
3	Enabled	Auto	Enabled	Frame	10Mbps/Full/FC-Off
4	Enabled	Auto	Enabled	Frame	No Link
5	Enabled	Auto	Enabled	Frame	100Mbps/Full/FC-Off
6	Enabled	Auto	Enabled	Frame	No Link
7	Enabled	Auto	Enabled	Frame	No Link
8	Enabled	Auto	Enabled	Frame	No Link
9	Disabled	Auto	Enabled	Frame	No Link
10	Enabled	Auto	Enabled	Frame	No Link
11	Enabled	Auto	Enabled	Frame	No Link
12	Enabled	Auto	Enabled	Frame	No Link
13	Enabled	Auto	Enabled	Frame	No Link
14	Enabled	Auto	Enabled	Frame	No Link
15	Enabled	Auto	Enabled	Frame	100Mbps/Full/Enabled (IEEE 802.3x)
16	Enabled	Auto	Enabled	Frame	100Mbps/Full/Enabled (IEEE 802.3x)

* Legal information © 1999 Intel Corporation

Figure 3 - Switch Status Page

Note: Port 9 has been disabled as a new unauthorized device has been connected. In order to restore service, the network administrator will be required to reset the port

3.3. MAC Filtering

The decision to use MAC filtering in two different ways allows for the greatest flexibility yet ensure that unauthorized devices are restricted. The first set of devices are the constantly used machines. These ports (ports 2-7) are configured such that no other devices (tied to MAC address) are able to use these ports at any time. The last four ports are set such that 16 pieces of gear will be allowed on these four ports. Given the total number of possible devices to be configured to use this last set of devices, There is a limitation that is accepted by that lab that there may be a desire to connect two devices that are in one of these last control groups. That is a limitation that was willing to be accepted in order to comply with security directives.

Part of the switch firmware requires that when assigning a MAC for static use it both the MAC and port must b set. If a MAC set for port 16 is connected to port 14, the port will not error, but will simply refuse to allow the network connection. It will appear that the port is in a deactivated state. However immediately upon plugging into the correct port, the connection is allowed to re-establish.

3.3.1. MAC Filtering Defined

Many people refer to MAC filtering as "security through obscurity" ("Security through obscurity: MAC address filtering (Layer 2 filtering)," n.d.). However, modern information security uses a defense in depth approach. In this approach to layered defense, not every control must be absolute. An early SANS whitepaper from Todd McGuiness defines the concept of defense in depth as:

> "… the concept of protecting a computer network with a series of defensive mechanisms such that if one mechanism fails, another will already be in place to thwart an attack. […] Utilizing the strategy of defense in depth will reduce the risk of having a successful and likely very costly attack on a network practices look at security as a process rather than a single tool" (McGuiness, 2001).

One drawback for MAC filtering is that of reliance on the MAC to be unique. This layer of security ignores the possibility of MAC spoofing but relies on other layers for protection against that type of attack.

3.3.2. Enable and Configure MAC Filtering

Start by connecting to the switch using a browser. In this case, the switch is located at 192.168.1.222 and enter the credentials for the device.

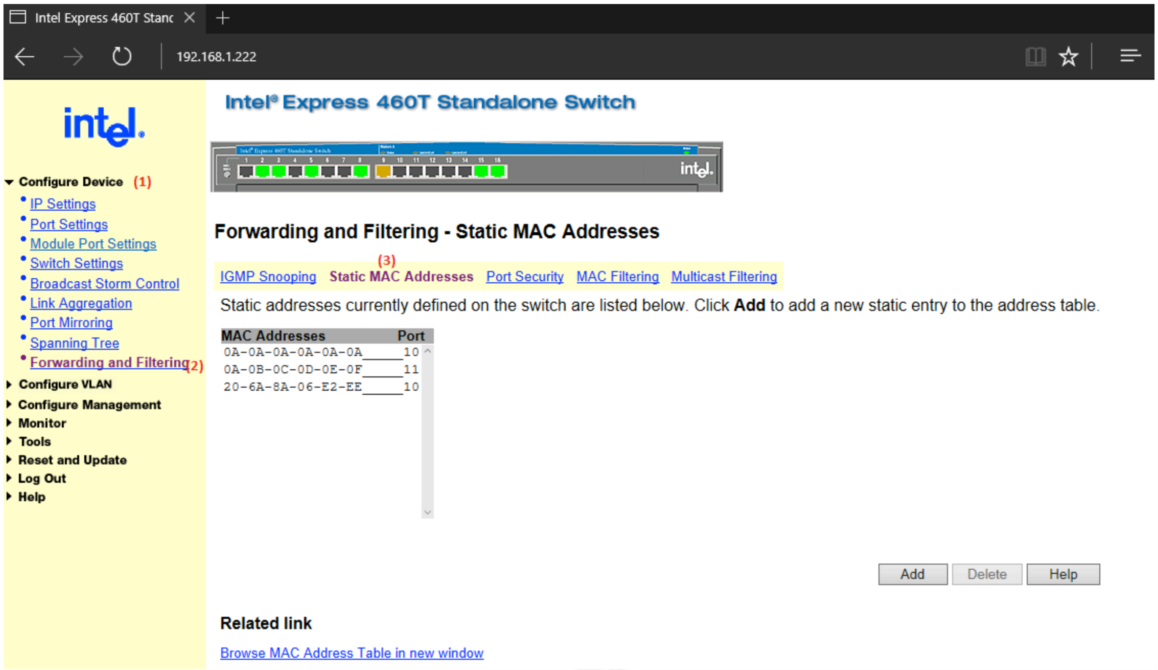

Forwarding and Filtering - Static MAC Addresses

Static addresses currently defined on the switch are listed below. Click **Add** to add a new static entry to the address table.

MAC Addresses	Port
0A-0A-0A-0A-0A-0A	10
0A-0B-0C-0D-0E-0F	11
20-6A-8A-06-E2-EE	10

Figure 4 - Switch Configuration - MAC filtering

To configure MAC filtering within this device, click on **Configure Device** (1), then **Forwarding and Filtering**(2), then on **Static MAC Addresses** (3). From the screen capture above it can be seen that there are currently only 3 devices configured with two on port 10 and 1 on port 11. If there are static settings made and no assigned MAC addresses are assigned a given port, the port will simply not respond and will not be MAC learned by the switch. For this configuration to work as expected

In addition, in order to deactivate MAC Learning, one will need to go to the port security page a Start by connecting to the switch using a browser. In this case, the switch is located at 192.168.1.222 and enter the credentials for the device.

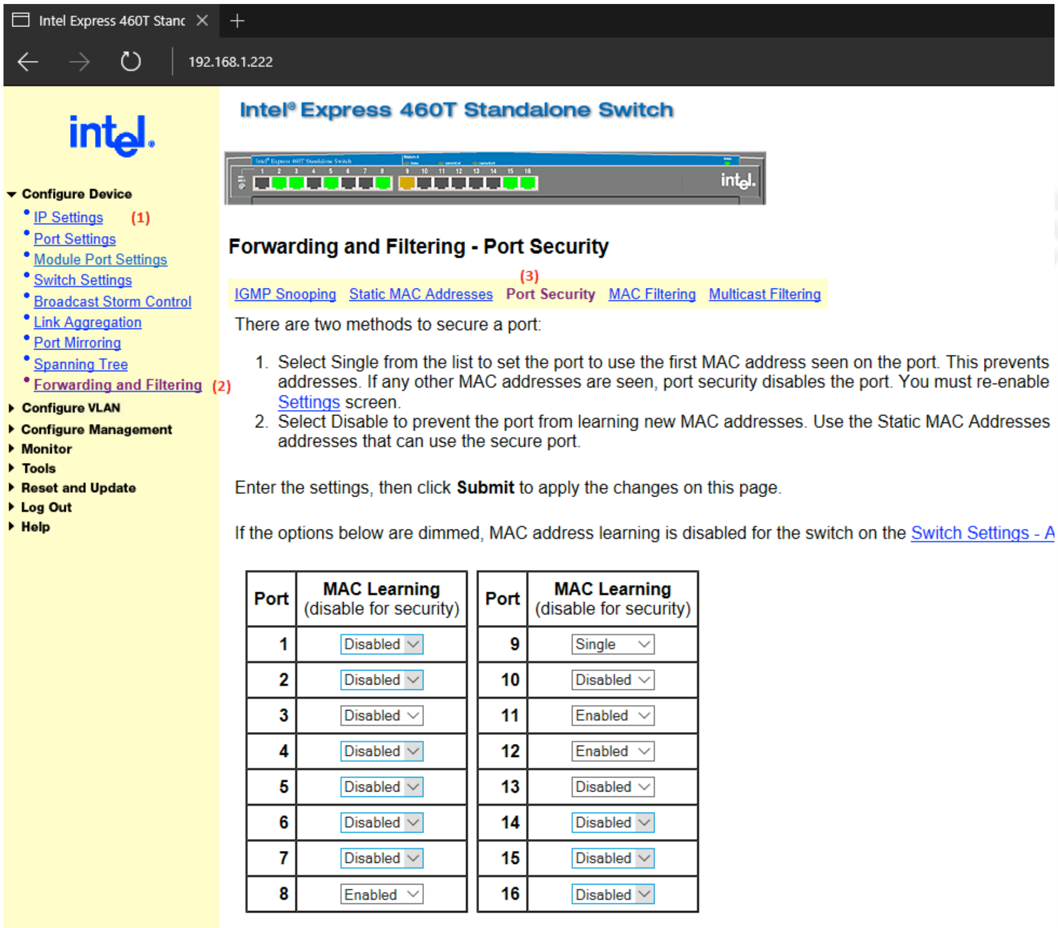

Intel® Express 460T Standalone Switch

Forwarding and Filtering - Port Security

There are two methods to secure a port:

1. Select Single from the list to set the port to use the first MAC address seen on the port. This prevents addresses. If any other MAC addresses are seen, port security disables the port. You must re-enable Settings screen.
2. Select Disable to prevent the port from learning new MAC addresses. Use the Static MAC Addresses addresses that can use the secure port.

Enter the settings, then click **Submit** to apply the changes on this page.

If the options below are dimmed, MAC address learning is disabled for the switch on the Switch Settings - A

Port	MAC Learning (disable for security)	Port	MAC Learning (disable for security)
1	Disabled	9	Single
2	Disabled	10	Disabled
3	Disabled	11	Enabled
4	Disabled	12	Enabled
5	Disabled	13	Disabled
6	Disabled	14	Disabled
7	Disabled	15	Disabled
8	Enabled	16	Disabled

Figure 5 - Switch Configuration Port Security Extended

To configure Port Security within this device, click on **Configure Device** (1), then **Forwarding and Filtering**(2), then on **Port Security** (3). From the screen capture above it can be seen that ports 1-7 and 13-16 have MAC learning disabled (meaning this will use only MAC addresses that are Static and assigned a given port).

3.4. Monitoring and Metrics

One of the unique challenges with monitoring a switch is the methods used for getting data out of the switch in an automated fashion. This particular switch uses Simple Network Management Protocol traps; simply a numeric code generated with events. The SNMP configuration can carry many items about a given device (Name, Manufacturer, time since last boot, IP Address, port configuration, etc.). The challenge with this device is that many of the Intel SNMP traps were not published for easy access. Comparing the entire SNMP configuration meant a line by line comparison between states of several thousand lines of raw text that may be merely strings. Many of the standard settings can be found using various web resources, but ultimately knowing that a port was down or was unplugged was a critical aspect to knowing the real status.

To that end, the switch was configured to send SNMP traps to a dedicated workstation. On that workstation, a copy of SNMP Trapwatcher from BBT Software (http://www.bttsoftware.co.uk/snmptrap.html) was then able to capture unknown traps when the switch ports were set to "deactivated." By cross-referencing the OID for these, it was determined that the object identifier (OID) displayed was "1.3.6.1.4.1.343.6.17.3.1.1.1.9." Using an OID database online this determined was an Intel (1.3.6.1.4.1.343) unit. After a deeper evaluation of the specific OID codes and values, it was determined these values were constant for the various settings and parameters for set ports.

Enabled	Disabled	Unplugged
<intel>.6.17.3.1.1.3.<port>=2	<intel>.6.17.3.1.1.3.<port>=3	<intel>.6.17.3.1.1.3.<port>=3
<intel>.6.17.3.1.1.4.<port>=5	<intel>.6.17.3.1.1.4.<port>=5	<intel>.6.17.3.1.1.4.<port>=2
<intel>.6.17.3.1.1.5.<port>=3	<intel>.6.17.3.1.1.5.<port>=3	<intel>.6.17.3.1.1.5.<port>=4

Figure 6 - SNMP Query Findings

So if one were to use software such as SNMPGET queries OIDs 1.3.6.1.4.1.343.6.17.3.1.1.3.1, 1.3.6.1.4.1.343.6.17.3.1.1.4.1, and 1.3.6.1.4.1.343.6.17.3.1.1.5.1 with a resultant set of 253, then the port is enabled. If the result codes were 353, the port was disabled and 324 meant the port was unplugged or the device on that port was powered down. The BATCH commands for this example:

```
snmpget -r:192.168.1.222 -o:.1.3.6.1.4.1.343.6.17.3.1.1.3.9 -q >>result.txt

snmpget -r:192.168.1.222 -o:.1.3.6.1.4.1.343.6.17.3.1.1.4.9 -q >>result.txt

snmpget -r:192.168.1.222 -o:.1.3.6.1.4.1.343.6.17.3.1.1.5.9 -q >>result.txt
```

Parsing through the result.txt file then yielded the 253, 353, or 324 based upon the port status.

The above demonstrated how to pull the data from the switch, the process to pull the error codes, convert this to a useful, human-readable output would require some custom coding. The complexity was in finding where the data was stored, but the retrieval of said data seemed simple. Below are examples of the batch processes used for this effort. To ensure ease of use, each port to be queried was given its own batch file. Each of those batch files was then called by the master file to launch each of the port query commands.

3.4.1. Batch Files

MASTER.bat

```
IF EXIST portstatus.txt del /F portstatus.txt
call test8.bat
call test9.bat
call test10.bat
call test11.bat
call test12.bat
```

TEST8.bat

```
REM Clean temp file
IF EXIST tempresult.txt del /F tempresult.txt

REM SNMPGet commands to txt
snmpget -r:192.168.1.222 -o:.1.3.6.1.4.1.343.6.17.3.1.1.3.8 -q
>>tempresult.txt
snmpget -r:192.168.1.222 -o:.1.3.6.1.4.1.343.6.17.3.1.1.4.8 -q
>>tempresult.txt
snmpget -r:192.168.1.222 -o:.1.3.6.1.4.1.343.6.17.3.1.1.5.8 -q
>>tempresult.txt

REM Cycle txt to create single 3 digit output
SetLocal EnableDelayedExpansion
set content=
for /F "delims=" %%i in (tempresult.txt) do set
content=!content!%%i
echo %content% > 111.txt
EndLocal
set /p ercode=<111.txt

REM - compare errcode to known values output to a status file
if %ercode%==353 goto start1
if %ercode%==253 goto start2
if %ercode%==324 goto start3
goto end

:start1
echo portstatus8=Disabled >>portstatus.txt
goto end
:start2
echo portstatus8=Enabled >>portstatus.txt
goto end

:start3
echo portstatus8=Unplugged >>portstatus.txt
goto end

:end
```

The above will output the status to a file called portstatus.txt. The results of that file looks like:

```
portstatus8=Unplugged
portstatus9=Disabled
portstatus10=Enabled
portstatus11=Enabled
portstatus12=Enabled
```

4. Conclusion

The problem that led to the ACME compromise was clearly due to a lack of controls. While there were multiple problems within a multitude of control areas, the primary issue was that of access control and the failure of Control #1. The immediate reaction of the company to address the risk was rigid and had a severe business impact. In order to reduce the impact and return to a more normal operational status, management determined that the new solution needed to address several areas. The first requirement was to allow DHCP where needed, but control access to the network through MAC address filtering and port security. The primary goal of both MAC Filtering and Port Security is to restrict access to known devices. Since financial issues are always a concern and the company has an open source mandate, the solution also needed use existing hardware and software or open source tools. MAC Filtering and Port Security met this requirement, as these functions are part of the existing software for this switch. Another requirement was that this new solution would employ an automated function to shut down a given port if an unauthorized device was connected. Port Security by its definition provides this functionality. The only area in which the new solution did not meet the requirements was that of notification. This particular switch does use SNMP traps to send change notices, but given that the SNMP Trap was very much a generic and undocumented code, there was room for improvement. There may yet be information about this particular switch that details this setting, but that IT staff was unable to find it at this time. The attached script files are only one way in which this information is available. By using this data, one has the opportunity to establish a process to query the switch and then compare outputs to determine if a change in status has taken place. Performing all of the above actions not only secured the weaknesses that allowed the compromise to happen but also allowed for management to learn more about this area, how it operates and what could be done to streamline those operations.

References

Center for Internet Security. (n.d.). Retrieved from https://www.cisecurity.org/critical-controls/faq/#faq-who-created

ISO/IEC 27002 code of practice. (n.d.). Retrieved from http://www.iso27001security.com/html/27002.html

Knaffl, W. (2016, May 3). *Case study: How CIS controls can limit the cascading failures during an attack.* Retrieved from https://www.sans.org/reading-room/whitepapers/casestudies/case-study-cis-controls-limit-cascading-failures-attack-36957

McGuiness, T. (2001, November 11). *Defense in depth.* Retrieved from https://www.sans.org/reading-room/whitepapers/basics/defense-in-depth-525

National Institute of Standards and Technology. (n.d.). NIST Manuscript Publication Search. Retrieved from http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r4.pdf

SANS Critical security controls training course | 20 critical controls | SEC566. (n.d.). Retrieved from https://www.sans.org/course/implementing-auditing-critical-security-controls?msc=cishp

SANS Institute - CIS Critical Security Controls: A Brief History. (n.d.). Retrieved from https://www.sans.org/critical-security-controls/history

Security through obscurity: MAC address filtering (Layer 2 filtering). (n.d.). Retrieved from http://www.cyberciti.biz/tips/linux-unix-bsd-mac-filtering.html

Switchport Security Configuration. (n.d.). Retrieved from https://www.pluralsight.com/blog/it-ops/switchport-security-configuration

Taralla, J. (n.d.). *SANS SEC566: Implementing and auditing the critical security – in depth*.