computer emergency response team CERT-EU for the EU institutions, bodies and agencies

CERT-EU Security Whitepaper 17-004

Mitigating Risks Related to Network Devices

Emilien LE JAMTEL
ver. **1.0**
June 22, 2017

TLP: WHITE

Contents

1 Introduction

Network devices – such as routers, switches, firewalls – are essential components of every IT infrastructure. All traffic (encrypted or not) has to go through several such network devices. Compromising network devices allows an adversary to steal sensitive data, corrupt communications, or disrupt activity of the targeted organization.

The range of attacks against network devices has been growing for the past years, from exploitation of undocumented access to development of complex implants modifying the behavior of devices. Known attacks go from passive monitoring to large-scale denial-of-service attacks against big organizations.

The purpose of this document is to provide recommendations on how to assess, prevent, and detect network devices compromise. Additionally, some ideas on how to deal with devices outside of the organization's perimeter are provided.

1.1 Target Audience

This document is aimed at general IT staff that has undertaken the responsibility of mitigating the risk of attacks against network devices in their infrastructure, especially network teams and security teams. This document only provides high-level guidelines. Different approaches are possible and may be valid. It does not supersede any specific applicable policies or procedures, which should be followed if they exist.

In case of doubts or any additional questions about this document, do not hesitate to seek further advice and assistance from your respective authorities or CERT-EU team.

2 Assessment

Identifying the potential attack surface of the infrastructure should be the first step to a proactive defense. This task should be performed by the security team in the organization with the agreement of the network team. To perform these tasks in a timely manner, the security team should have a read access to network device configurations.

2.1 Inventory of Devices

A global Risk Management Plan (RMP) routinely includes an inventory of network devices used in the infrastructure. This inventory should be actively maintained with every change properly and timely updated. The minimal amount of information to keep in the inventory includes:

- device name,
- device description,
- model/manufacturer,
- firmware version,
- operating system version,
- network interfaces,
- physical location,
- device status and last reboot time,
- support information (including planned end-of-support).

Additionally, a history of running configurations of network devices should be part of the inventory process as it enables potential detection of unwanted configuration changes. This part will be discussed later in this document.

The most popular way to remotely manage network devices is SNMP. Use of `SNMPv3` implementation is highly recommended as it corrects a number of security-related issues of earlier versions of SNMP.

2.2 Network Visibility

Network devices and their open services can be reached from outside of the perimeter, as well as from different locations (segments) inside the perimeter. While some services are available from the Internet, others also from different locations on the internal network. Having an accurate picture of the network visibility from different locations is essential to assess the potential attack surface against the organization.

2.2.1 Internet-Facing Devices

To assess the network visibility from outside, one can either perform the scanning or use online (scanning) services.

2.2.1.1 Online Scanning Services

Internet-facing networks are scanned every day by many commercial and private entities to identify and fingerprint available services. Some online services provide openly those scanning results to hackers and security researchers.

Three popular online scanning services include:

Service	Example for CERT-EU network range
https://www.shodan.io/	https://www.shodan.io/search?query=net%3A212.8.189.16%2F28
https://censys.io/	https://censys.io/ipv4?q=212.8.189.16%2F28
https://www.zoomeye.org/	https://www.zoomeye.org/search?q=+cidr%3A212.8.189.16%2F28

Scanning results may differ from one service to another, based on their scanning policy.

2.2.1.2 Nmap

Software solutions to asses network visibility of an IP range are abundant.

One popular tool is Nmap[1] as it performs several tests including host discovery, port scanning, and fingerprinting of operating systems and services. It is distributed under the terms of the GNU General Public License (GPL) as published by the Free Software Foundation and is available for UNIX and Windows systems.

For accurate results, scanning task should be performed from an Internet-facing IP address external to the network being tested.

Example command for full TCP port SYNscan of CERT-EU network range:

```
$ sudo nmap -sS -sV -Pn -p1-65535 -oA CERT-EU 212.8.189.16/28
```

Example command for standard UDP scan of CERT-EU network range:

```
$ sudo nmap -sU -oA CERT-EU-UDP 212.8.189.16/28
```

For detailed explanation of Nmap capabilities, please refer to the manual: https://nmap.org/book/man.html

2.2.2 Internal Network Devices

Network devices can be targeted from internal network as well. If an attacker is able to get a foothold in an internal network using a malware sent via spear-phishing, leaked credentials, insecure remote access, or vulnerabilities on Internet-facing services – he may want to target network devices from there.

Today, as most networks are segmented, it is important to perform reconnaissance from different locations inside of the network to assess the potential attack surface:

- DMZs (especially Internet facing DMZ),
- user LAN,
- admin LAN,
- VPN user LAN.

As a best practice, it is recommended to assess the network segmentation by performing reconnaissance scans between all internal (V)LANs.

[1] http://nmap.org/

2.2.3 Regular Scans

Changes to network topology occur all the time:

- changes to existing devices,
- deployment of new machines/services,
- firewall rules modifications.

Mistakes might happen when these changes occur. Having the capability to detect in a timely manner newly opened services from different locations inside and outside of the perimeter limits the risk of insecure exposed services. Schedule of reconnaissance scanning must be based on a risk model (eg., daily, weekly, or monthly scanning).

Here is a basic risk model example:

criticality of the assets	High	Medium	Low
scan from Internet	daily	daily	daily
scan from Internet DMZ	daily	weekly	weekly
scan from user network	daily	weekly	monthly
scan from admin network	daily	monthly	monthly

Nmap provides the Ndiff tool (https://nmap.org/ndiff/) to aid the comparison of Nmap scans results and detect changes.

A lot of tools (open-source or commercial) exist to perform these tasks (scanning, scheduling, and change tracking) and should be part of every IT security toolset. Most Vulnerability Assessment platforms provide this capability as well.

2.3 Vulnerability Assessment

As any system, network devices can be prone to vulnerabilities or misconfiguration. To identify vulnerabilities that could be exploited remotely or with limited local accounts, regular vulnerability scanning must be performed.

As for the reconnaissance phase, the vulnerability assessment plan must follow a risk model keep track of changes. As vulnerability scanning may be resource consuming, to avoid availability issues it is critical to define time-window for scanning.

To be as exhaustive as possible and to avoid network latency while scanning, scanners must be located close to the assets that are being targeted. Consequently, location of scanners for vulnerability assessment may differ from the reconnaissance phase.

The following simplified schema gives an example of the location for scanners (reconnaissance and vulnerability assessment). Vulnerability scans are made from the closest scanner and reconnaissance scans are used to see available services from different locations in the network.

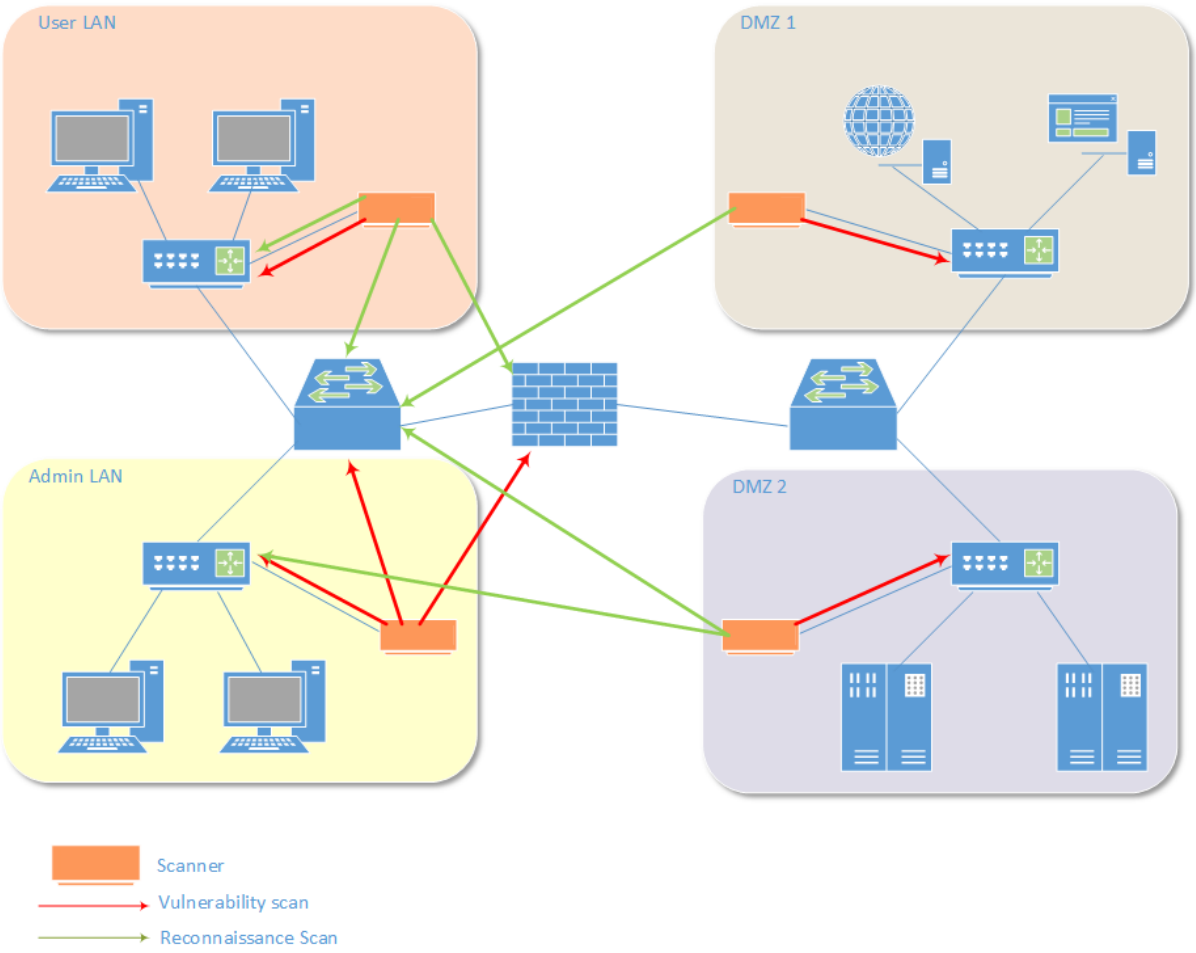

Figure 1: Example locations of scanners

3 Prevention

By enforcing security on network devices, applying strong policies and setting up a state-of-the-art network segmentation, organizations can protect their network against most adversaries.

3.1 Patching

Network devices should be part of the global patch management plan. Security officers must follow the important updates provided by network devices vendors as it may impact the patch management plan.

As explained before, having an exhaustive inventory of network devices is part of the patch management plan.

The following table provides links to security advisories provided by some network devices vendors. The list is non-exhaustive.

Vendor	Links
CISCO	https://tools.cisco.com/security/center/publicationListing.x http://www.cisco.com/go/psirt
Juniper	https: //kb.juniper.net/InfoCenter/index?page=content&channel=SECURITY_ADVISORIES
CheckPoint	https://supportcenter.checkpoint.com/supportcenter/portal?eventSubmit_ doGoviewsecurityalerts=&view=SixMonths#type=Security+Alerts
Fortinet	http://fortiguard.com/rss/ir.xml
Huawei	http://www.huawei.com/en/psirt
HP	http://www8.hp.com/us/en/business-services/it-services/security-vulnerability.html https://h20566.www2.hpe.com/portal/site/hpsc/public/kb/secBullArchive#HF
Palo Alto Networks	https://securityadvisories.paloaltonetworks.com/
NetGear	http://www.netgear.com/about/security/
DELL & SonicWALL	https://support.sonicwall.com/ https://support.software.dell.com/release-notes-product-select
Brocade	http://www.brocade.com/en/support/security-advisories.html
D-Link	http://support.dlink.com/index.aspx
Nokia	https://ps-kb.alcatel-lucent.com/portal/ (needs Authentication)
Ericsson	https://www.ericsson.com/login (needs Authentication)

3.2 Unsupported Devices

As any product, network devices and security appliances reach the end of their Product Life Cycle for a number of reasons. These reasons may be due to market demands, technology innovation and development driving changes in the product, or the products simply mature over time and are replaced by functionally-richer technology.

Customers are usually alerted about the end-of-sale date prior to the end-of-support. As support may be provided for a long time after the end-of-sale notice by the vendor, having a plan for replacement is critical to maintain the security level of an infrastructure.

3.3 Configuration Hardening

Configuration hardening is a vast subject and implementation depends on specific network devices used. However, the security officers should keep in mind the recommendations presented in the following subsections.

3.3.1 Avoid Using Local Accounts

Use of local accounts on network devices (especially when servicing large networks) create issues for network administrators and security managers. By doing so, administrators have to support and maintain a database of local accounts on each device. In most cases credentials will be reused and shared between network administrators, and this may increase the risk of credentials disclosure and compromise.

Remote authentication provides the ability to use individual user accounts for each network administrator and to enforce enterprise password policy. The AAA (Authentication Authorization and Accounting)[2] protocol has become the *de facto* standard for network equipment and has been implemented by most popular vendors. It can even be configured to work with Active Directory [3].

3.3.2 Disable Unused Services

Still today, some network devices come configured with unnecessary services which must be disabled. These unneeded services, especially those that use User Datagram Protocol (UDP), are infrequently used for legitimate purposes, but could be abused to launch attacks against the infrastructure.

3.3.3 Ban Insecure Protocols

A lot of network devices are still managed via `telnet` / `http` . Today most devices support secure protocols implementing encryption that should be used. Insecure protocols should be disabled by default.

In addition, using secure file transfer protocols when moving sensitive data over the network should be enforced. The use of the Secure Copy Protocol (SCP) in place of FTP or TFTP, or the use of `SNMPv3` in place of `SNMPv1` is highly recommended.

3.3.4 Use Dedicated VLAN and workstations for Administration

The administration service (SSH, SNMP ...) should be accessible only to network administrators. Good security practice is to set up administration services on a dedicated network interface connected to a dedicated management VLAN[4]. Also, if possible, prohibit communication between network devices on this management VLAN. Private VLAN[5] and MAC-Access Control List should be implemented to isolate devices in the management VLAN.

[2] https://en.wikipedia.org/wiki/Diameter_(protocol)
[3] http://woshub.com/configuring-network-devices-authentication-using-active-directory/
[4] Out-of-band management network.
[5] https://en.wikipedia.org/wiki/Private_VLAN

Workstations accessing network devices should be dedicated to administrative tasks. Network administrators should use a different machine for activities not related to administration of network devices.

The following simplified schema gives an example of typical out-of-band management network implementation.

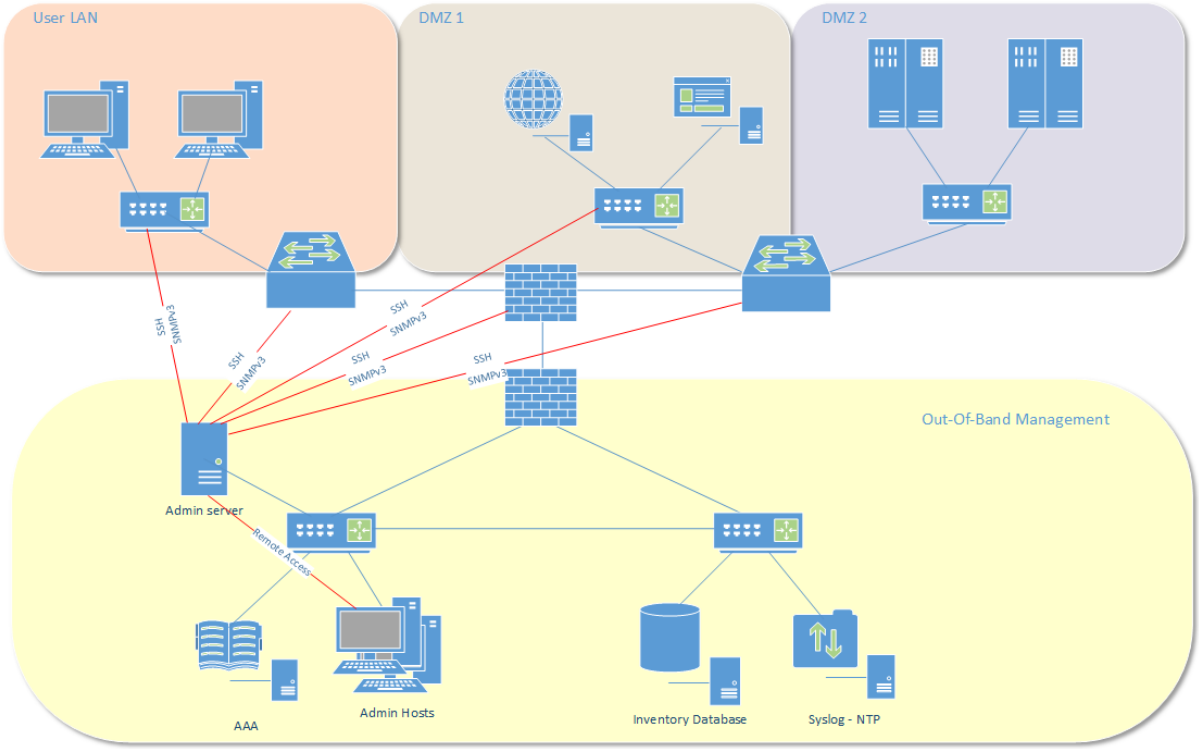

Figure 2: Out-of-band management network implementation

3.3.5 Setup Centralized Logging

Having visibility on existing, emerging, and historic events related to security incidents allows detection and analysis of potential threat to the network infrastructure. All important events from all network devices should be extracted, stored, and correlated in a Security Information and Event Management (SIEM). It provides real-time analysis of security alerts generated by network devices and supports forensics investigation.

To remotely gather logs, two solutions are provided by most network devices:

- `syslog` server logging: Use `syslog` to forward log messages to external `syslog` servers for storage. Keep in mind that `syslog` traffic is not encrypted by default and should be configured to transit through secured networks only.
- `SNMP trap` logging: use `SNMP traps` to send log messages to an external `SNMP` server — less granular in the logging than `syslog`.

Later in this document, information to gather to increase detection capabilities will be discussed.

3.3.6 Apply Password Policy to SNMP Community Strings

With SNMPv1 or SNMPv2, SNMP community strings are used as for authentication and are transmitted in clear text. If SNMP community strings are still used in the infrastructure, they should be considered as *passwords* and the password policy should be applied to them:

- length,
- complexity requirements,
- maximum age.

If a network infrastructure is running for a long period of time, SNMP community strings should be considered as insecure and probably known by threat actors.

3.3.7 Integrity Check Functionalities at Startup

Some vendors provide functionalities to check for the firmware and operating system integrity via digital signatures every time the device is started. Such functionalities will prevent modified devices from running if the signatures are invalid.

4 Detection

Even with strong defenses, some adversaries may find a way to compromise network devices. Organization should be prepared to detect suspicious behavior on such devices and have the capability to investigate those behaviors.

4.1 IPS/IDS and Honeypots

Network devices will be targeted from Internet or from internal network if threat actors have a foothold in it. Most organizations deploy IDS (Intrusion Detection Systems) and/or IPS (Intrusion Prevention System) to detect or prevent attempts against their assets, including network devices. Deploying IPS/IDS on perimeter border is a well-known good practice, but to cover the internal threat as well, security officers should deploy sensors between internal zones.

Events and alerts generated by IPS/IDS should be forwarded to the centralized security information and event management (SIEM) system and monitored by a competent incident-response team. In most cases, same product will be used for event management and log analysis.

Another tools useful for detecting intrusion attempts are network honeypots. Network honeypots are systems impersonating a device or a service and, as they are not supposed to be supporting any real activity, any interaction has to be considered as malicious. By carefully placing network honeypots to sensitive locations, early intrusion attempts against network devices can be detected.

The following table provides a list to honeypot projects used to impersonate network devices or related services. The list is non-exhaustive.

Project	Links
Honeyd	http://www.honeyd.org/
Kippo (SSH)	https://github.com/desaster/kippo
Conpot (ICS)	http://conpot.org/
Bifrozt	https://sourceforge.net/projects/bifrozt/
Kojoney (SSH)	http://kojoney.sourceforge.net/

4.2 Threat Actor Behavior and Suspicious Events

Threat actors may get access to a network device via different ways (0-day, leaked password, weak configuration, etc.), but to gain persistence or conduct operation they may generate some weak signals or cause absence of certain signals, which may help in detection. Without discussing any specific implementation, here is a list of known signals that can be observed after a network device is compromised.

4.2.1 Account Creation

After compromising a device, the attacker may be tempted to create his own privileged account to perform tasks without the risk of losing his access. Also, for accountability, logging may be configured per user so using an *out-of-band* user may be more discreet for the attacker.

An attacker may create a user, set privileges, and then clean his traces by deleting the account so asynchronous log gathering may not be able to detect suspicious activity.

Most network devices have specific log events related to user management, but the following events should always be available:

```
User created
User deleted
User privilege level modified
```

4.2.2 Unexpected Reboots

After compromising a privileged user, the threat actor may try to deploy a modified firmware or operating system. In most cases, device reboot is needed to apply the modifications. The threat actor can wait for scheduled reboot from the victim or do it himself.

On most systems, the uptime (the time a system has been working and available) can be requested via SNMP or via a system command. The last reboot time can also be part of the devices inventory.

Some specific log events may be available:

```
System was rebooted or shutdown
System just started
System is scheduled for reboot
```

4.2.3 Newly Opened Services

A threat actor may open some specific remote services on devices to comply with his operational needs or by mistake, when modifying the running configuration.

Newly opened services on network devices can be detected via network scan or differential configuration assessment.

4.2.4 Extraction of Running Configuration

After compromising a device, the threat actor will extract information about the environment. It usually starts with extraction of the running configuration of the device. The attacker may also use information from one compromised device to retrieve running configuration from other devices. It is indeed not unusual that network devices share SNMP community strings.

Running configuration can be extracted via an interactive command line or via remote services like SNMP. In some cases, the upload of extracted configuration to an external host will be initiated by the network device. Configuration is also stored on disk/memory, so the threat actor may generate dump and exfiltrate them.

IDS should be able to detect unexpected connections from network devices to external host. If deep packet inspection capabilities are available, defining rules to detect configuration extraction via unencrypted data transfer services (FTP/TFTP) with specific keywords (SNMP string, device name convention, internal IP addresses, etc.) may help detect exfiltration of sensitive configuration.

Compromised systems in the infrastructure can also be used as pivot before exfiltration, to avoid suspicious connections from network devices to external hosts.

Some `syslog` events may be triggered when running configuration is extracted or when a remote connection is made:

```
SNMP Write request
Writing configuration on local storage
FTP/TFTP connection from src_ifc:src_ip/src_port to dst_ifc:dst_ip/dst_port
Coredump filesystem image created
Memory Dump created
```

4.2.5 Missing Logs

In order to avoid detection, threat actors will deactivate logging of activity on compromised assets. Depending on the logging policy, regular messages coming from network devices are expected.

Depending on the centralization service (syslog, SIEM, etc.), it may be easy to detect hosts which stopped sending events.

Bellow an example of a search with Splunk for such hosts (not responding since 180 seconds):

```
| metadata index=* type=hosts | eval age = now()-lastTime | where age > 180 | sort age d
  |convert ctime(lastTime) | fields age,host,lastTime
```

Also a statistical analysis on reported events per hour may trigger alerts about strange behavior regarding events sent.

4.2.6 Memory and CPU Threshold

When modifying the behavior of a network device (real-time injection, data exfiltration, etc.), it may impact memory and CPU load. Unexpected memory and CPU threshold alerts may be an indicator that something suspicious is ongoing.

SNMP traps and Remote Monitoring Alarm can be used to receive alerts when a threshold is exceeded. The threshold must be defined based on the normal operation of the network device. Threshold can also be defined for the network usage.

4.2.7 Created Files

To modify the firmware or system image on the device, a threat actor may need to upload a modified image or firmware and store it on the device. In most cases, only full logging (debug mode) will generate alert when a file is created/deleted/modified. The following alerts may be generated:

```
User View file
User Remove file
User Rename file
User Modify file
User Create file
User Create directory
User Remove directory
```

When a suspicious device is analyzed, a good practice is to analyze the available storage by performing a forensics analysis and recover potentially deleted files.

4.2.8 Strange Connections from Network Devices

By modifying the behavior of a device, a threat actor can use it as a *proxy* to attack further into the network. A classic example is to use a compromised device to access SNMP service on other devices. Private VLANs should allow to detect attempted connections initiated by network devices.

Also data can be exfiltrated from the compromised network device. The following simplified schema gives an example of data exfiltration exploiting a compromised network device allowing SNMP requests from Internet.

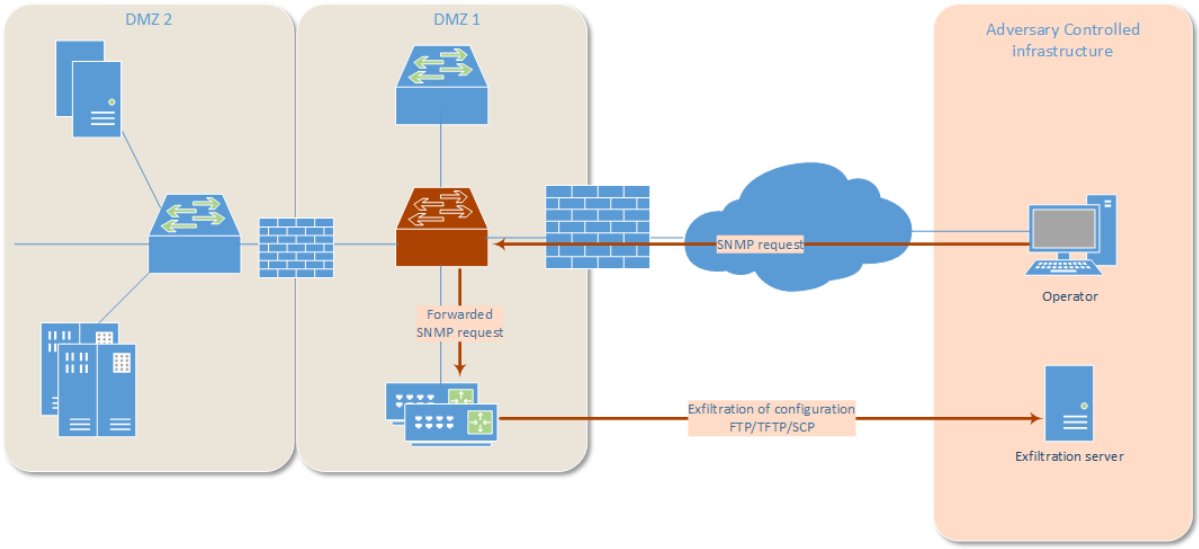

Figure 3: Data exfiltration via proxied SNMP request

Any unusual connection initiated from a network device should be investigated.

4.2.9 Modification of Configuration

A threat actor may try to change the configuration of devices to get access to other parts of the network or to facilitate exfiltration of data. It is critical to be able to detect those changes.

By having an inventory of the configurations, one can easily detect unwanted changes. Another possibility is to create logs when users enter a command or change the configuration, but it may be difficult to distinguish malicious behavior from normal activity of the device administrators.

4.2.10 Suspicious Commands

If the commands history executed on network devices is logged, it is recommended to build for each OS a list of commands generating alerts on a SIEM. Here is a non-exhaustive list of actions which could be abused by threat actors:

```
Update
Writing/Reading firewall configuration
Access to low-level commands (shell, tclsh ...)
booting options
Memory dump
Coredump/snapshots
VPN configuration
Copy from external host (FTP, TFTP, SCP ...)
Debug
Listing admin users
Monitor traffic interface
```

4.3 Integrity Checks

If anything triggers an alert and is considered as potentially suspicious, it is important to be able to perform integrity checks on the devices. Depending on the device, several items have to be checked:

- firmware / BIOS integrity,
- software integrity,
- image integrity (on disk and loaded in memory).

Not all vendors provide public information on how to perform these checks. We will discuss here information publicly available for CISCO devices. For other vendors, it is recommended to request technical details via support services.

4.3.1 Firmware/BIOS Integrity

Firmware can be stored in *read-only* (original firmware) or *read-write* (updated firmware) areas. It is quite unusual that network administrators update the firmware of network devices, so the fact that a firmware is loaded from a *read-write* area may be suspicious.

On CISCO routers, the easiest and most convenient way to detect if the ROMMON (firmware) has been upgraded, is to execute the command `show rom-monitor` from the command line interface (CLI).

```
Router#show rom-monitor
ReadOnly ROMMON version:
System Bootstrap, Version 12.4(1r) [hqluong 1r], RELEASE
SOFTWARE (fc1)
Copyright (c) 2005 by cisco Systems, Inc.
Upgrade ROMMON version:
System Bootstrap, Version 12.4(13r)T, RELEASE SOFTWARE (fc1)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 2006 by cisco Systems, Inc.
Currently running ROMMON from Upgrade region
ROMMON from Upgrade region is selected for next boot
```

The string `Currently running ROMMON from Upgrade region` confirms that the firmware has been upgraded and is loaded from the *read-write* area.

If the device is using the upgraded region and firmware update was not foreseen, one should perform a full coredump of the device and contact CISCO support for further analysis:

Execute and store the output of show commands like:

```
show region
show version
show running-configuration all
show tech-support
```

Create a memory dump, compute hash and store in a secure manner:

```
############## memory dump configuration ###############
router#conf t
ip ftp username Cisco
ip ftp password 7 0321xxxxxxxxx710A1016141D
exception core-file r-router compress timestamp
exception protocol ftp
exception region-size 65536
exception dump ip_address
end
############## memory dump execution ###############
router# write core
```

The content of the ROMMON will be present in the coredump. However, it is difficult to locate. To dump the ROMMON from a CISCO device, you can also boot the device in *priv rommon* mode[6].

4.3.2 Image/Software Integrity

A threat actor can try to modify the operating system or its binaries. Once loaded by the device, the malicious code will be executed. Depending on a vendor, there are several checks one can perform.

On CISCO products, you can either check the image from the running device or check the image offline[7]

To verify the integrity directly from the device:

```
router#verify sup-bootdisk:c7600rsp72043-advipservicesk9-mz.151-3.S3
Verifying file integrity of  sup-bootdisk:c7600rsp72043-advipservicesk9-mz.151-3.S3
.....<output truncated>.....Done!
Embedded Hash  MD5 : FCEBD3E1AF32221091E920D5960CAE45
Computed Hash  MD5 : FCEBD3E1AF32221091E920D5960CAE45
CCO Hash       MD5 : E383BF779E137367839593EFA8F0F725
Signature Verified
router#
```

The *Embedded Hash* should be identical to the *Computed Hash*. The *CCO Hash* should correspond to the value provided in the *Support and Downloads* area on the CISCO website for this image.

However, if the `verify` command is altered to display a wrong hash, it may be recommended to use an offline image file hash. To do that, download the image to another machine and use a tool to calculate the MD5 checksum and the SHA512 checksum and compare it with the value provided by CISCO in the *Support and Downloads* area.

For example on linux, `md5sum` and `sha512sum`:

[6] http://www.bitshift.ch/eng/support/kbase/000002.asp
[7] http://www.cisco.com/c/en/us/about/security-center/integrity-assurance.html

```
$ md5sum 7600rsp72043-advipservicesk9-mz.151-3.S3.bin
e383bf779e137367839593efa8f0f725  7600rsp72043-advipservicesk9-mz.151-3.S3.bin
```

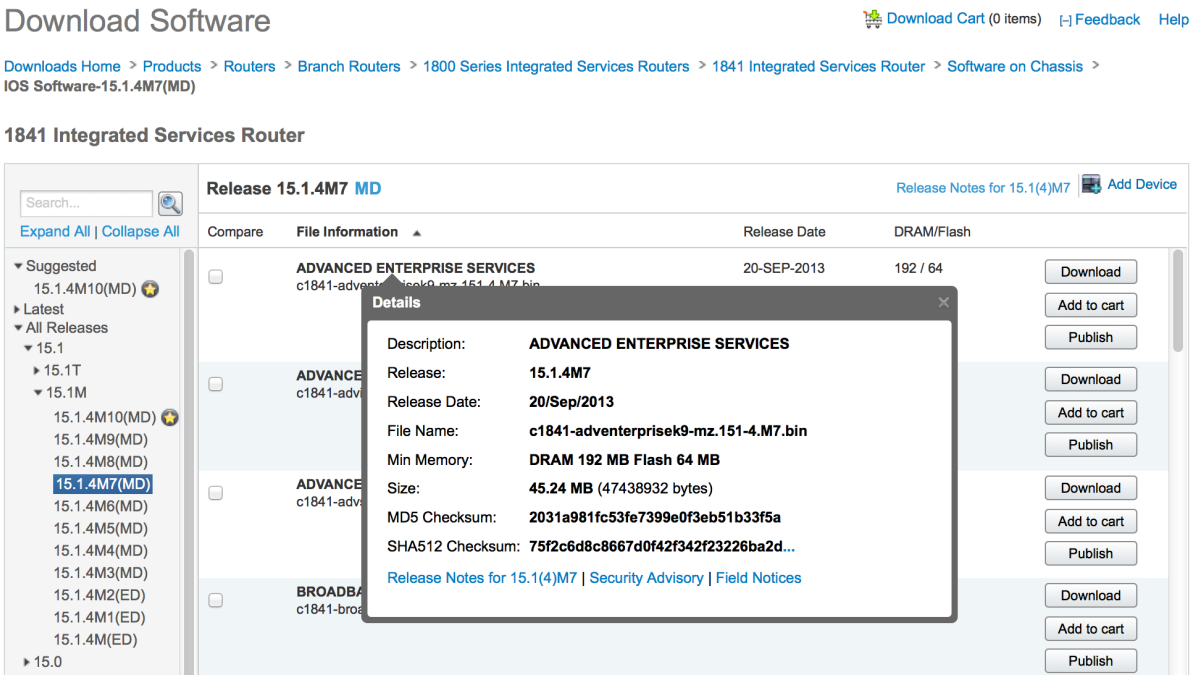

Figure 4: MD5 and SHA512 verification

4.3.3 Loaded Image Integrity

Even if the image on disk is clean, the threat actor may have modified the image loaded in memory. The only way to check for modification is to dump the memory, extract the image loaded and analyze it.

CISCO provides good documentation on how to perform offline analysis of coredump: http://blogs.cisco.com/security/offline-analysis-of-ios-image-integrity

The first step is to generate a coredump:

```
############## memory dump configuration ###############
router#conf t
ip ftp username Cisco
ip ftp password 7 0321xxxxxxxxx710A1016141D
exception core-file r-router compress timestamp
exception protocol ftp
exception region-size 65536
exception dump ip_address
end
############## memory dump execution ###############
router# write core
```

Then one needs to extract the `TEXT` segment from the memory dump. To do that, one needs to locate which part of the dumped memory (output of `write core`) contains the `TEXT` segment. On the device, look at the output of `show region` and `show platform tlb` commands:

```
router# sh region
Region Manager:

    Start         End     Size(b)  Class  Media  Name
    0x08000000  0x0BFFFFFF    67108864  Iomem  R/W    iomem
    0x40000000  0x4BFFFFFF   201326592  Local  R/W    main
    0x40101040  0x42DCFFFF    46985152  IText  R/O    main:text
    0x42DDA370  0x430F534F     3256288  IData  R/W    main:data
    0x430F5350  0x44B156EF    27394976  IBss   R/W    main:bss
    0x44B156F0  0x4BFFFFFF   122595600  Local  R/W    main:heap
    0x50000000  0x5FFF7FFF   268402688  Local  R/W    more heap
    0x80000000  0x8BFFFFFF   201326592  Local  R/W    main:(main_k0)
    0xA0000000  0xABFFFFFF   201326592  Local  R/W    main:(main_k1)

router#sh platform  tlb

    Mistral revision 5
    TLB entries : 45
    Virt Address range      Phy Address range      Attributes
    0x10000000:0x1001FFFF   0x010000000:0x01001FFFF   CacheMode=2, RW, Valid
    0x10020000:0x1003FFFF   0x010020000:0x01003FFFF   CacheMode=2, RW, Valid
    0x10040000:0x1005FFFF   0x010040000:0x01005FFFF   CacheMode=2, RW, Valid
    0x10060000:0x1007FFFF   0x010060000:0x01007FFFF   CacheMode=2, RW, Valid
    0x10080000:0x10087FFF   0x010080000:0x010087FFF   CacheMode=2, RW, Valid
    0x10088000:0x1008FFFF   0x010088000:0x01008FFFF   CacheMode=2, RW, Valid
    0x18000000:0x1801FFFF   0x010000000:0x01001FFFF   CacheMode=0, RW, Valid
    0x19000000:0x1901FFFF   0x010000000:0x01001FFFF   CacheMode=7, RW, Valid
    0x1E000000:0x1E1FFFFF   0x01E000000:0x01E1FFFFF   CacheMode=2, RW, Valid
    0x1E880000:0x1E899FFF   0x01E880000:0x01E899FFF   CacheMode=2, RW, Valid
    0x1FC00000:0x1FC7FFFF   0x01FC00000:0x01FC7FFFF   CacheMode=2, RO, Valid
    0x30000000:0x3001FFFF   0x070000000:0x07001FFFF   CacheMode=2, RW, Valid
    0x40000000:0x41FFFFFF   0x000000000:0x001FFFFFF   CacheMode=3, RO, Valid
    0x42000000:0x427FFFFF   0x002000000:0x0027FFFFF   CacheMode=3, RO, Valid
    0x42800000:0x429FFFFF   0x002800000:0x0029FFFFF   CacheMode=3, RO, Valid
    0x42A00000:0x42BFFFFF   0x002A00000:0x002BFFFFF   CacheMode=3, RO, Valid
    0x42C00000:0x42C7FFFF   0x002C00000:0x002C7FFFF   CacheMode=3, RO, Valid
    0x42C80000:0x42CFFFFF   0x002C80000:0x002CFFFFF   CacheMode=3, RO, Valid
    0x42D00000:0x42D7FFFF   0x002D00000:0x002D7FFFF   CacheMode=3, RO, Valid
    0x42D80000:0x42D9FFFF   0x002D80000:0x002D9FFFF   CacheMode=3, RO, Valid
    0x42DA0000:0x42DBFFFF   0x002DA0000:0x002DBFFFF   CacheMode=3, RO, Valid
    0x42DC0000:0x42DC7FFF   0x002DC0000:0x002DC7FFF   CacheMode=3, RO, Valid
    0x42DC8000:0x42DCFFFF   0x002DC8000:0x002DCFFFF   CacheMode=3, RO, Valid
    0x42DD0000:0x42DD7FFF   0x002DD0000:0x002DD7FFF   CacheMode=3, RW, Valid
    0x42DD8000:0x42DDFFFF   0x002DD8000:0x002DDFFFF   CacheMode=3, RW, Valid
    0x42DE0000:0x42DFFFFF   0x002DE0000:0x002DFFFFF   CacheMode=3, RW, Valid
    0x42E00000:0x42FFFFFF   0x002E00000:0x002FFFFFF   CacheMode=3, RW, Valid
    0x43000000:0x437FFFFF   0x003000000:0x0037FFFFF   CacheMode=3, RW, Valid
    0x43800000:0x43FFFFFF   0x003800000:0x003FFFFFF   CacheMode=3, RW, Valid
    0x44000000:0x45FFFFFF   0x004000000:0x005FFFFFF   CacheMode=3, RW, Valid
    0x46000000:0x47FFFFFF   0x006000000:0x007FFFFFF   CacheMode=3, RW, Valid
    0x48000000:0x49FFFFFF   0x008000000:0x009FFFFFF   CacheMode=3, RW, Valid
    0x4A000000:0x4BFFFFFF   0x00A000000:0x00BFFFFFF   CacheMode=3, RW, Valid
    0x4C000000:0x4DFFFFFF   0x00C000000:0x00DFFFFFF   CacheMode=3, RW, Valid
    0x4E000000:0x4FFFFFFF   0x00E000000:0x00FFFFFFF   CacheMode=3, RW, Valid
    0x08000000:0x09FFFFFF   0x00C000000:0x00DFFFFFF   CacheMode=2, RW, Valid
    0x0A000000:0x0BFFFFFF   0x00E000000:0x00FFFFFFF   CacheMode=2, RW, Valid
    0x50000000:0x51FFFFFF   0x080000000:0x081FFFFFF   CacheMode=3, RW, Valid
    0x52000000:0x53FFFFFF   0x082000000:0x083FFFFFF   CacheMode=3, RW, Valid
    0x54000000:0x55FFFFFF   0x084000000:0x085FFFFFF   CacheMode=3, RW, Valid
    0x56000000:0x57FFFFFF   0x086000000:0x087FFFFFF   CacheMode=3, RW, Valid
    0x58000000:0x59FFFFFF   0x088000000:0x089FFFFFF   CacheMode=3, RW, Valid
    0x5A000000:0x5BFFFFFF   0x08A000000:0x08BFFFFFF   CacheMode=3, RW, Valid
```

```
    0x5C000000:0x5DFFFFFF   0x08C000000:0x08DFFFFFF   CacheMode=3, RW, Valid
    0x5E000000:0x5FFFFFFF   0x08E000000:0x08FFFFFFF   CacheMode=3, RW, Valid
```

In this example the text segment resides between address `0x40101040` and `0x42DDA370`, and it size is 46985152 bytes. The translation between virtual memory and physical memory requires to change the location in our dump by changing the `0x4` by `0x0`.

One can now use `dd` in Linux to extract the loaded image in the coredump and compute the hash of the `TEXT` segment to be compared with the value provided by CISCO in the *Support and Downloads* area.

```
$ dd if=Coredump bs=1 skip=1052736 count=46985152 of=TEXT_segment
$ md5sum TEXT_segment
441ec282be815e24c74ac917f3a42fec  TEXT_segment
```

To perform forensic analysis of a known compromised device, one should use a non-compromised `TEXT` area to locate malicious modifications.

5 Appendix : Mitigating Risks on *Out-of-Perimeter* Devices

5.1 Introduction: ISPs Can/Will Be Compromised

Internet Service Providers (ISPs) are an interesting target for threat actors. They have a lot of clients and compromising their proper devices allow the attackers to act against a target without compromising the target itself.

Most targeted devices are called *Customer Edge Routers*. These devices are owned and managed by the service provider and connect the customer to the MPLS architecture.

Not only most of the customer's Internet traffic will go through these devices, but also communication between various customer sites, so compromising them allows the threat actor to perform a large number of actions against the targets including:

- traffic exfiltration,
- traffic injection,
- DNS tampering,
- user impersonation,
- targeted Denial Of Service (DoS).

The following schema describes quickly the connection between CE (Customer Edge) and PE (Provider Edge) routers:

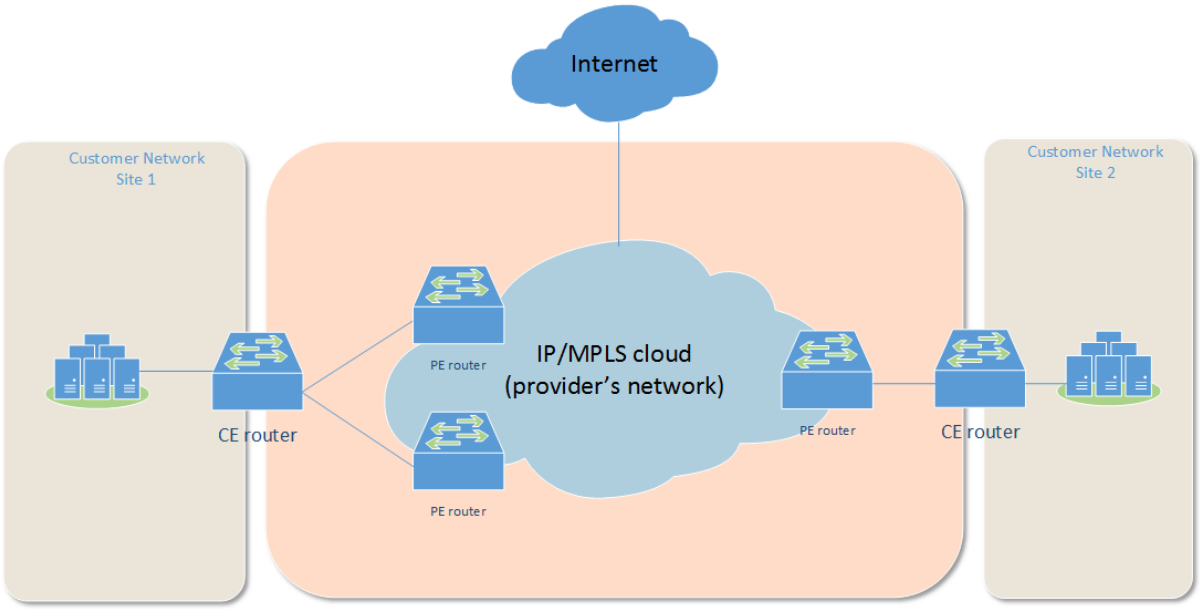

Figure 5: CER and PER

These devices are critical for the perimeter security and are managed by another entity. However, as a customer it is important to enforce some best practices on the ISP side.

5.2 Expectations from the Internet Service Provider

For large organizations, contracts with Internet Service Provider (ISP) are decided through a call-for-tender. The topics presented in the following subsections should be part of the selection process.

5.2.1 Logs, SOC, and Reporting

The ISP has its own logging policy and capabilities to detect potential incidents on devices. Logs should be gathered and stored from all devices involved in the network communications. The level of logging should be detailed by the ISP and validated by the customer IT security team. This should include:

- TACACS logs,
- syslog from all devices,
- traffic logs for Custom Edge Routers.

The logged data should be under review by a dedicated SOC (Security Operation Center), and any incident must be reported to the customer in a timely manner.

Based on the severity of the incident, an Incident Response team should be involved.

Another solution is to request the ISP to forward the logs to a SIEM solution in the customer premises. Then, the SOC of the customer can assess the received data and detect incidents. The reported data can then be correlated with internal logs.

5.2.2 Incident Response Capabilities

The ISP should have a competent and well-sized Incident Response team able to detect and analyze in a timely manner any ongoing incident. It should include the capabilities and procedures to perform integrity checks on devices and analyze logs from a centralized system.

This team should work in cooperation with the Incident Response team of the customer and share results of all investigations.

The customer should include in the agreement with the ISP the requirements for the maximum response time for Incident Responses tasks:

- Incident detection and reporting,
- evidence gathering (logs, memory dumps, forensics images)
- integrity checks results,
- Managerial and technical investigation reports

5.2.3 Security Policies

The security policies applied to Network Devices must be provided by the Service Provider to the Customer. These policies should include:

- description of administration pattern (protocol, authentication, accountability, etc.),
- logging policy,
- upgrade policy and life-cycle of devices,
- network isolation with other customers,
- incident-response capabilities.

5.3 Independent Security Assessment

To ensure that the security is properly handled by the Internet Service Provider, independent security assessment should be performed by a third party. This assessment should be performed in a scheduled manner and include:

- Configuration analysis of network devices,
- Compliance to security policies and best-practices,
- Review of security policies,
- Evaluation of incident response capabilities,
- Review of documented procedures related to device management and incident Responses,
- Physical security assessment.