

A photograph showing two medical professionals, a woman with glasses and a man, looking at a computer monitor displaying a medical scan. The woman is on the left, and the man is on the right. The background is blurred, showing a clinical setting.

Philips mShield for hospital networks

Securing medical devices

Summary

The proliferation of network-connected medical devices that make use of off-the-shelf, embedded operating systems and the increase of cyberattacks aimed at health care institutions make hospitals increasingly vulnerable to malicious attacks. To protect themselves against such attacks, hospitals need a multi-layer security approach that offers many barriers to intrusion, including patches, anti-malware solutions and firewalls. Philips mShield is a firewall developed for imaging systems that provides an additional layer of security without limiting device functionality. It protects devices so that patients can continue to receive their exams, even if there is malicious activity on the network.

Introduction

Digitization in hospital environments continues to evolve to provide better healthcare for patients and improved workflow for operators. Personal, sensitive and confidential data travels from radiology systems throughout the hospital and back again. Securing this information and protecting it from malicious attacks is as vital as it is difficult.

Use mShield to

- Prevent malware replication over the network
- Ensure equipment availability
- Provide an additional layer of security

Protect your medical device with Philips mShield

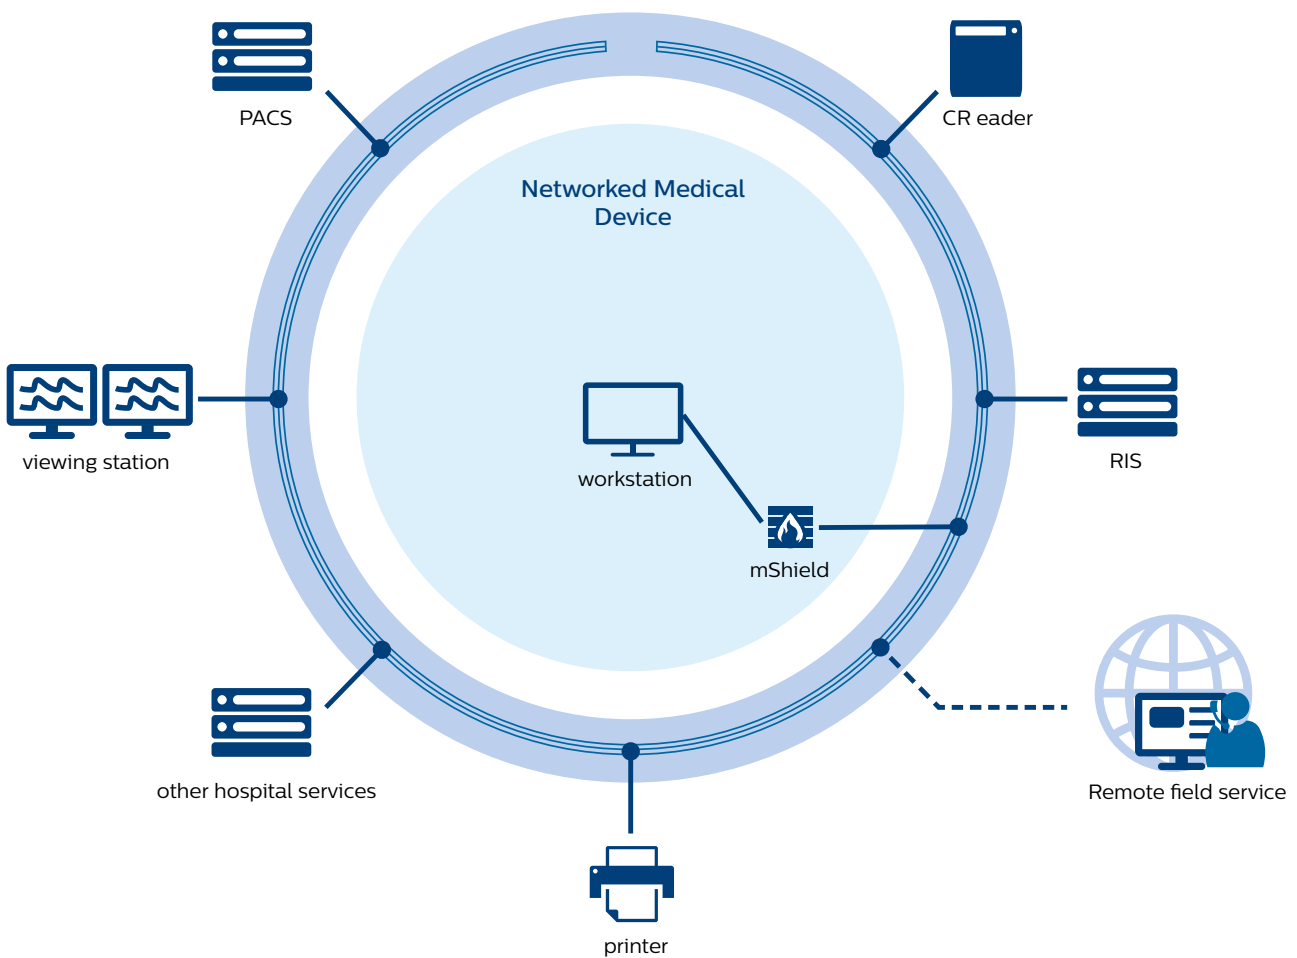
Addressing risk: recommendations and challenges

A multi-layer security concept helps hospitals defend against many threats to their data and systems. It is a common best practice to include:

- Regular essential operating system security patches and anti-malware solutions, to secure against cyberattacks and viruses throughout system lifetime

- Application-level hardening and access controls to mitigate risks by allowing only authorized and recognized data on the system
- Hardware firewalls that only allow authorized traffic

While this list is similar – if not identical – to recommendations for other industries, medical devices have special circumstances that make effective implementation of multi-level security particularly important.



Software patches

Software patches remove vulnerabilities that are discovered after the software has been installed. Regular updates help health care institutions avoid software weaknesses that might provide cyberattacks entry to systems and damage both the systems and data.

Anti-malware solutions

Anti-malware solutions are an important safeguard that should be part of every endpoint security concept. One common anti-malware solution is a virus scanner. To be effective, virus scanners must have up-to-date virus definition files that include new viruses and the latest scanning engine. If medical devices are off-network, even for a short time, they may miss an important update.

A second anti-malware solution is application whitelisting, which blocks all software that is not included on a whitelist from being executed on the system. Application whitelisting effectively “freezes” the system software to a known state. Because it avoids unauthorized modifications even by malware that is not yet known (zero-day exploits), application whitelisting doesn’t require regular updates to be effective.

Hardware firewalls

A third weapon in the battle against attack is hardware firewalls. A typical firewall is designed to establish a barrier between internal and external networks, and uses security rules to determine if network traffic is safe. It also can separate the internal network into subnets and applies firewall filter rules between them. This approach can isolate important nodes on the network until the threat is neutralized.

A dedicated firewall for Philips medical devices: mShield

The special circumstances of medical devices mean that even with outstanding endpoint security, imaging systems are vulnerable. To strengthen hospital security efforts while also protecting the essential healthcare function of Philips medical devices, Philips developed mShield, a dedicated firewall that effectively blocks threats to imaging equipment, protecting imaging systems without limiting their use.

mShield consists of both hardware and software, and is based on the security-oriented operating system OpenBSD¹. It provides network isolation and protection, minimizing the connectivity exposure (“attack surface”) between the medical equipment and the hospital’s network. Each device should have its own mShield, although it is possible to use a single mShield on several connected devices.

Because mShield is so specific, it can use strict rules to evaluate the validity of traffic, and restrict traffic to only authorized devices and specific services. For example, X-ray equipment typically uses DICOM as its main communication protocol and only a few other supporting protocols. With a default-deny-policy and few firewall exceptions, mShield can effectively decouple the modality from the network and hide the modality’s structure, while at the same time maintaining connectivity for medical applications or remote service.

mShield can prevent malware replication over the network, ensure equipment availability, provide an additional layer of security, and offer security if the medical device’s embedded operating system is no longer supported by the operating system’s manufacturer.



Prevents malware replication over the network

mShield blocks virtually all typical network-based replication paths for viruses and worms, thus preventing infections. This approach obviates the biggest problem with virus scanners, which depend on regular and timely updates and can only react when the virus starts to interact with the system.

While it is true that medical equipment could still be compromised by a network protocol accepted by mShield, the risk of infection is low. This is because hospital environments differ in various ways from one another and medical networks are often intensively customized in a way that makes it hard for mass-malware to succeed everywhere. In addition, mShield establishes and enforces trust relations between specific nodes inside a hospital network, further reducing the risk of mass spreading of malware.

Infection is also possible via paths that bypass mShield, such as removable media. The risk depends on the frequency of use, security measures implemented in the medical equipment, and on the hospital's security policy. However, if medical equipment is infected via this pathway, mShield stops the virus from replicating on other machines within the network, because mShield inspects both incoming and outgoing packets.

Ensures availability of the medical equipment

mShield serves as the single point of entry for any network communication coming from or going to the protected medical equipment. Thus, mShield absorbs the impact of the attack for the medical equipment, and even if mShield crashes as a result, the device continues to work.

This makes many types of attacks – as well as accidental network anomalies that could impact the device – impossible or significantly more difficult. The device will work as if disconnected from the network, and will reconnect to transfer images or other data once the network is available.

Provides an additional layer of security

mShield serves to maintain a high level of continuous network security at all times by blocking remote network-based exploitation. In the past, mShield has successfully blocked malware such as the SASSER worm or the WannaCry ransomware.

Offers security for unsupported third-party software

Discontinued, unsupported software also poses challenges for hospitals. Philips supports its equipment for at least 10 years after production has stopped, but third-party software suppliers usually discontinue support for their software, including supplying security patches, much earlier, leading to a support gap. Upgrading to the latest operating system ensures continued support, including receiving security patches. However, when financial or technical limitations make this impossible, mShield adds a layer of protection.

“What we have right now... is still a very fragile healthcare ecosystem, is how I would describe it. And we need to move to a better place where devices are patchable, that they're updateable, that they can stand an exploit, an attack or breach and still function safely and properly, and that the hospital itself and manufacturers and, again, the sector at large is resilient to be able to withstand that and deliver continuity of care.”²

– Susanne Schwartz, MD, acting director, Office of Strategic Partnerships and Technology Innovation, FDA Center for Devices and Radiological Health.

Addendum

How mShield works: technical description, installation and configuration

Philips mShield comprises Layer-2 filtering, including the Address Resolution Protocol (ARP) filter, as well as Layer-3 filtering (state-full packet inspection: SPI) in combination with a stealth mode (also called a bridging mode). As one expert explained, “A firewall is a doorway that everyone knows they want to get through. A packet-filtering bridge is more like a secret agent that picks off the bad guys from the shadows whom nobody can attack back. OpenBSD Packet Filtering (PF) bridges can drastically increase security for any network architecture.”³ The filtering rules are customized for the communication requirements of Philips medical products.

mShield design is in accordance with recommendations of international industry consortia, such as NEMA, COCIR and JIRA. They recommend firewalls as an “effective and flexible tool” and as part of an overall strategy to safeguard the data integrity of medical information systems.⁴

As a tool that is custom-built for protecting systems at the device level, mShield has many benefits over broad-based security measures. It delivers a security solution comparable to host firewall software products, but has advantages in terms of flexibility, service and security.

Acts like an ethernet switch

Typically, firewalls have two or more interfaces, with each interface being configured for a specific subnet. The firewall then does the routing between these subnets in addition to the packet inspection and filtering. Each computer in a subnet has to know the firewall (router) by its IP address (and Medium Access Controller (MAC) address*) in order to send all packets destined for another subnet.

* The MAC address is of the format aa:bb:cc:dd:ee:ff (each colon separated part is a hexadecimal value in the range 00-ff). The MAC address is a unique identifier for a Network Interface Card.

Unlike those common firewalls, mShield is not a router, but instead acts similar to an Ethernet switch, in that mShield does not split up the network into subnets. This provides a service and administration advantage in many cases, including retrofit cases, in which it eliminates the need to re-configure the protected medical devices, as well as the need for a local IT administrator to assign a dedicated subnet for the medical device.

Invisible

mShield is not directly visible in the network – neither by IP nor by MAC addresses. This feature reduces mShield’s attack surface, which enables long duration of maintenance-free operation.

Robust

The mShield operating system – OpenBSD – is known to be one of the most secure network operating systems. As implemented in mShield, this general-purpose operating system is reduced in size and function to only contain the minimum needed. For example, it has only the minimum needed kernel drivers and system tools and it does not allow interactive user logins by default. As well, mShield runs entirely from memory (RAM disk). With the exception of configuration and software updates, there is no write access to the built-in flash disk, which greatly improves uptime and robustness. This architecture allows mShield to survive power outage or manual switch-off at any time without violating the integrity of its internal software. Thus, mShield always reliably starts up in a clean and working state.

Network traffic treatment

Network packet filtering works on different levels of the TCP/IP protocol stack. By default, mShield blocks all packets that are not packaged according to IPv4 or ARP. The IP addresses of all protected hosts are stored in the mShield configuration. This facilitates IP address-based filtering. Furthermore, a specially designed ARP filter prevents ARP spoofing (intentional or accidental e.g. through a duplicate IP address) for the configured addresses.

TCP communication only goes through with a valid set of TCP flags. Any malformed packets are dropped, regardless of their origin (e.g. packets with the SYN flag and the FIN flag set at the same time). UDP and ICMP packets are filtered state-fully as well.

Where it makes sense, the number of source nodes initiating a communication (TCP, UDP or ICMP) is tracked and limited. A maximum packet rate limit is enforced to prevent denial-of-service conditions from affecting the protected medical device. In other words, if a modality is protected by mShield, the worst case during a denial of service condition is the loss of network connectivity as mShield itself becomes saturated, but the modality stays available for local clinical workflows, so patients will still receive their exams. Data transfer occurs when the network and system are secure.

Hardware

The mShield hardware was chosen with security and business continuity in mind and fulfills all needed country-specific regulations, such as CE, UL and CSA. There are no failure-prone mechanical storage devices built in, ensuring long hardware life. It also has no surplus hardware features, like a keyboard or video. The high-quality hardware is certified for extended operating temperatures up to 55 °C. mShield hardware is suitable for upgrading older medical equipment as well as integrating with new medical equipment.

Installation and configuration

To install and configure mShield properly, the service engineer has to know the network configuration of the medical device, all its components and the peer hosts with which the medical device communicates. The use of predefined templates (per system) eases proper mShield setup, and any change of the network parameters can be easily adjusted while preserving network interoperability.

Service engineers use a compatible Philips service tool that contains all necessary information (including version information) to configure mShield, facilitate upgrades, conduct fault-finding and extract log files.

Software updates can be performed locally via the service tools or indirectly via a remote connection to the protected medical device, depending on both the device's ability to support remote software distribution and on market capabilities.

Philips research and development team actively monitors potential and emerging kernel- and application- bugs and uses a world-class quality system to rapidly evaluate and deploy fixes as necessary.

Configuration items

At the time of configuration, local service engineers' main task is to identify all attached devices by their type and network parameters, such as IP address. They must also identify and configure communication relationships between hosts in the network according to the individual situation on site, for example regarding:

- NTP (time synchronization)
- Syslog (central logging / audit trail)
- DICOM or "Secure DICOM" based PACS and RIS systems, Printers, PCR Readers and other devices
- Philips remote service

Logging

The syslog facility, which is built into mShield, writes all log files into the memory (RAM). The log files may be exported prior to loss of power or rebooted to prevent loss. The reason for non-permanent log files within mShield is so that the system can be switched off at any time without violating the integrity of the built-in flash disk.

References

- 1 See <http://openbsd.org/> as well as the respective BSD license found at <http://openbsd.org/policy.html>
- 2 "Medical Device Security: The FDA's View." Careers Info Security. 9 July 2019. <https://www.careersinfosecurity.com/medical-device-security-fdas-view-a-12748>, accessed 5 November 2019.
- 3 Source: George Rosamond "Building a more secure network": <http://www.sans.org/rr/whitepapers/modeling/1415.php>
- 4 Source: "Defending Medical Information Systems Against Malicious Software", Dec. 2003, by NEMA (National Electrical Manufacturers Association-USA), COCIR (European Coordination Committee of the Radiological Electrometrical Industry) and JIRA (Japan Industries Association of Radiological Systems): www.nema.org/prod/med/upload/medical-defending.pdf

